Documentation 8.2

ZABBIX

05.11.2026

Contents

User manual **5**

Welcome to Zabbix Documentation **5**

Get started with Zabbix . . . 5
Zabbix Cloud . . . 5
Mobile app . . . 5
Developer center . . . 5
Community & other resources . . . 5
Installation and first steps . . . 6
Web interface installation . . . 6
Quickstart guides . . . 11
Installation: additional guides . . . 85
Upgrading . . . 165
Known issues . . . 181
Requirements . . . 189
What's new and upgrade notes: 8.2 . . . 199
Upgrade notes . . . 199
Breaking changes . . . 199
Other . . . 199
What's new and upgrade notes: 8.2.x . . . 199
Zabbix appliance . . . 199
Processes . . . 203
Server . . . 205
Proxy . . . 231
Agent . . . 275
Agent 2 . . . 304
Java gateway . . . 365
Sender . . . 369
Get . . . 371
JS . . . 372
Web service . . . 373
Bridge adapter . . . 376
MCP server . . . 379
Environment variables . . . 386
Configuration . . . 388
Hosts and host groups . . . 399
Items . . . 417
Problem detection with triggers . . . 635
Visualization . . . 686
Templates and template groups . . . 710
Templates out of the box . . . 710
Notifications upon events . . . 724
Macros . . . 775
Users and user groups . . . 789
Scheduled reports . . . 800
Data export . . . 804
Configuration export/import . . . 812
Regular expressions . . . 877
Events and resolution . . . 883
Problem acknowledgment and suppression . . . 886
Manual closing of problems . . . 889

Maintenance periods . . . 890
Tagging . . . 895
Complex event processing . . . 898
Event correlation . . . 902
Discovery . . . 909
Network discovery . . . 909
Active agent autoregistration . . . 918
Low-level discovery . . . 922
Virtual machine monitoring . . . 944
VMware monitoring item keys . . . 946
Virtual machine discovery key fields . . . 963
JSON examples for VMware items . . . 967
VMware monitoring setup example . . . 972
Web monitoring . . . 977
Web monitoring items . . . 986
Real-life scenario . . . 987
Service monitoring . . . 994
Service tree . . . 995
SLA . . . 999
Setup example . . . 1000
Service monitoring upgrade . . . 1004
Encryption and security . . . 1005
Encryption using certificates . . . 1013
Encryption using pre-shared keys . . . 1020
Secure connection to the database . . . 1023
Secure connection between frontend and server . . . 1035
Access control . . . 1039
Cryptography . . . 1043
Securing web server . . . 1044
Storage of secrets . . . 1047
Google Chrome TLS certificate trust . . . 1055
Security best practices . . . 1057
Web interface . . . 1058
Menu . . . 1058
Frontend sections . . . 1066
User settings . . . 1332
Global search . . . 1337
Frontend maintenance mode . . . 1339
Page parameters . . . 1340
Definitions . . . 1343
Additional frontend languages . . . 1343
Creating your own theme . . . 1344
Debug mode . . . 1345
Cookies used by Zabbix . . . 1346
Time zones . . . 1347
Resetting password . . . 1348
Time period and host selectors . . . 1348
Configuration parameters: Zabbix frontend . . . 1350
API . . . 1356
Method reference . . . 1361
Appendix 1. Reference commentary . . . 2227
Appendix 2. Changes from 8.0 to 8.2 . . . 2234
Appendix 3. Changes in 8.2 . . . 2234
Extensions . . . 2234
Loadable modules . . . 2236
Frontend modules . . . 2245
Appendixes . . . 2246
Installation and setup . . . 2246
Items . . . 2292
Supported functions . . . 2332
Macros . . . 2332
Version compatibility . . . 2366
Time period syntax . . . 2367

Unit symbols . . . 2368
Command execution . . . 2369
Escaping examples . . . 2370
What is Zabbix . . . 2372

Zabbix Cloud **2376**

Zabbix Cloud **2376**
Set up and manage Zabbix Cloud . . . 2376
Learn about Zabbix Cloud . . . 2376
Deploy Zabbix in the cloud . . . 2376
Node configuration . . . 2386
Overview . . . 2387
Overview tab . . . 2388
Access filters tab . . . 2390
Encryption tab . . . 2391
Backups tab . . . 2391
History tab . . . 2392
Maintenance tab . . . 2393
Upgrade tab . . . 2393
Adding users . . . 2395
Introduction . . . 2395
Adding users to organizations . . . 2395
Adding users to nodes . . . 2396
Key differences of Zabbix Cloud . . . 2397
Introduction . . . 2397
Audit log . . . 2399
Set up a Zabbix Cloud proxy . . . 2401

Mobile app **2409**

Mobile app **2409**
Set up Zabbix for the app . . . 2409
Learn about the app . . . 2409
Set up Zabbix for the app . . . 2410
App overview . . . 2411
Secure mobile access . . . 2419

Developer center **2420**

Developer center **2420**
Build on Zabbix . . . 2420
Creating a plugin . . . 2420
Plugin interfaces . . . 2425
Python library for Zabbix . . . 2428
Installation . . . 2428
Quickstart guide . . . 2428
Use Zabbix API . . . 2432
Collect data from Zabbix agent . . . 2435
Send data to Zabbix server or proxy . . . 2436
Debug logging . . . 2440
Modules . . . 2440
Module file structure . . . 2441
Widgets . . . 2448
Tutorials . . . 2456
Examples . . . 2483
Changes to extension development . . . 2483

Zabbix manpages **2484**
zabbix_agent2 . . . 2484
zabbix_agentd . . . 2485
zabbix_get . . . 2487
zabbix_js . . . 2489
zabbix_proxy . . . 2490
zabbix_sender . . . 2492

zabbix_server 2497
zabbix_web_service 2499

Copyright notice **2500**

User manual

Welcome to Zabbix Documentation

Your go-to resource for working with Zabbix monitoring, from basic setups to advanced configurations. This manual covers everything needed to install, configure, and operate Zabbix.

What's new in Zabbix 8.2

Get started with Zabbix

Installation Step-by-step instructions to install Zabbix on your preferred platform, covering various operating systems and database configurations.

Requirements A list of supported platforms and software prerequisites to help you prepare your environment for a successful Zabbix deployment.

Quickstart guides Concise, task-oriented how-tos that walk you through the basics—from your first configuration steps to receiving your first problem alert.

Zabbix Cloud

Start using Zabbix Cloud Guided instructions to get your Zabbix Cloud instance up and running quickly, including initial configuration and connecting your monitored devices.

Zabbix Cloud vs. on-premises Compare the differences in management, scalability, and maintenance between the two deployments, and learn which can support your workflow most seamlessly.

Explore Zabbix Cloud Get an overview of Zabbix Cloud, including its key features, perks, and what comes with a subscription. For common questions on setup and usage, see the FAQ.

Mobile app

Get the app Install Zabbix mobile app (coming soon) on your device, and start monitoring your Zabbix environment on the go.

Set up Zabbix for the app Install Zabbix bridge adapter to let mobile devices connect to Zabbix. Then add your device and turn on push notifications.

App overview Learn the basics of using Zabbix mobile app.

Developer center

Frontend modules Guides and references for building custom modules that extend or modify the Zabbix frontend to suit specific use cases.

Widgets A breakdown of widget structure and logic, with instructions for creating custom dashboard elements tailored to your needs.

Plugins An overview on how to develop and manage Zabbix plugins to extend functionality or integrate with external systems.

Community & other resources

Zabbix Forums A space to exchange ideas, find solutions, and share experiences across all levels of Zabbix expertise.

Zabbix Blog News, tutorials, and case studies curated by the Zabbix team and community contributors.

How-to videos A video library featuring demonstrations, discussions, and tips to help you make the most of Zabbix.

Installation and first steps

The following sequence outlines the Zabbix setup process:

1. Verify that the required software is installed and running.
2. Install Zabbix.
3. Install the Zabbix web interface.
4. Start monitoring.

Details for each step are provided below.

Required software Zabbix works well with the following software (often called the LAMP stack):

- **Linux operating system:** Ubuntu, Debian, CentOS, or others
- **Web server:** Apache or NGINX
- **Database server:** MySQL/MariaDB or PostgreSQL
- **PHP**

On most modern Linux systems, PHP may already be installed. In that case, you may only need to set up the database server and web server.

For more detailed information, such as software supported versions, see Required software.

Install Zabbix The **easiest** way of installing Zabbix is from **distribution packages** (alternatively, from containers). The packages install minor software dependencies.

Precise installation instructions, tailored for your operating system, are available on the **Zabbix website**.

Select your platform and follow the instructions provided. It is recommended to install a Zabbix server, a Zabbix agent 2 and the frontend (web interface) for a start. You may add Zabbix proxies and more agents later.

Part of the installation process is creating a database (for keeping configuration and history data) and importing initial data into it.

Install Zabbix web interface Complete the steps of the web interface **installation wizard**.

You will be asked for the database password (created in the previous step), so that the frontend can communicate with the database.

Start monitoring Log in to the Zabbix web interface.

Use **quickstart guides** to start monitoring popular resources (Linux, Apache, MySQL, or other). These quick guides explain which templates you want to work with to start monitoring popular resources.

You may want to start with the quickstart guide for monitoring a Linux server.

Web interface installation

This section provides instructions for installing the Zabbix web interface. Follow these steps after installing the Zabbix backend.

The Zabbix frontend is written in PHP and requires a supported web server such as Apache or NGINX with PHP configured. See the frontend requirements for information about the required software and libraries.

To use a language other than English in the Zabbix web interface, its locale must be installed on the web server, along with the PHP gettext extension. Please note that changing the language may change the time/date format.

To set up SSL for the Zabbix frontend, see Setting up SSL for Zabbix frontend.

Installation wizard

To install the web interface, open the Zabbix frontend URL in a browser:

- For Apache: `http://<server_ip_or_name>/zabbix`
- For NGINX: `http://<server_ip_or_name>`

Welcome screen

The frontend installation wizard opens on the Welcome screen.

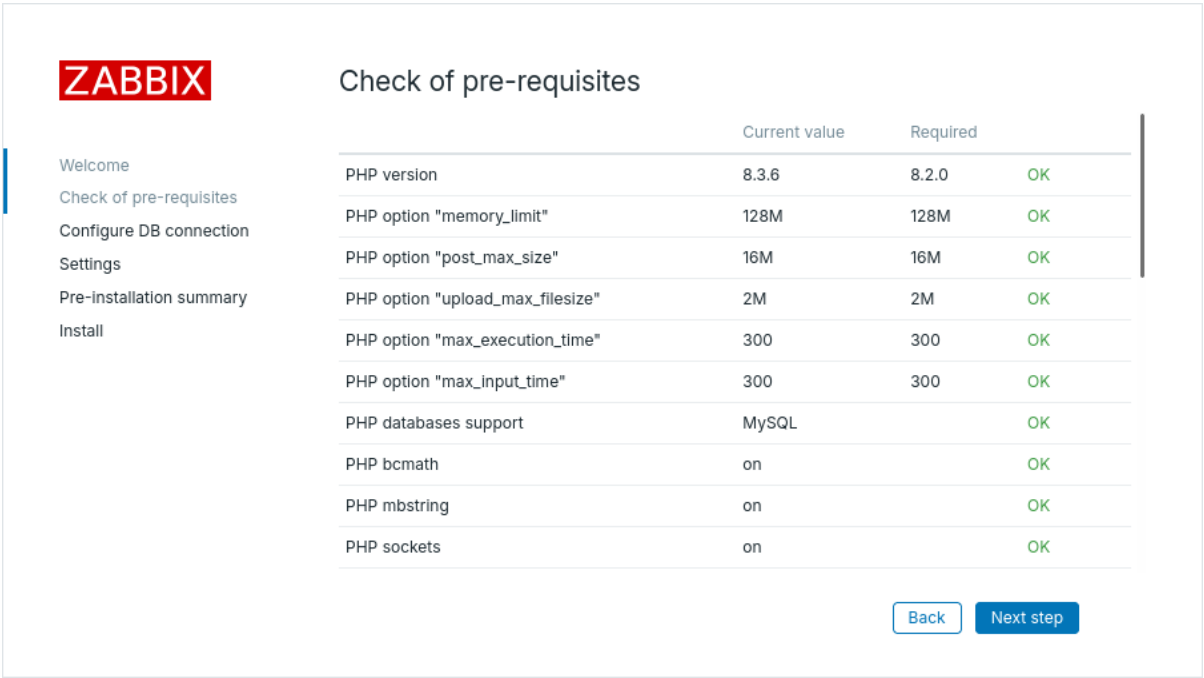


The system-wide default language is English (en_US), with the US time/date format. To change the system-wide default language, select the desired language from the *Default language* drop-down menu.

Click the *Next step* button to proceed.

Check of pre-requisites

Make sure that all mandatory Zabbix frontend pre-requisites are met.



See Requirements for a list of required and optional PHP extensions, the required PHP version, and supported databases. Note that failed optional pre-requisites are displayed with a *Warning* status, and the setup process can proceed even if they are not met.

Required values for PHP configuration directives in the `php.ini` file used by the web server:

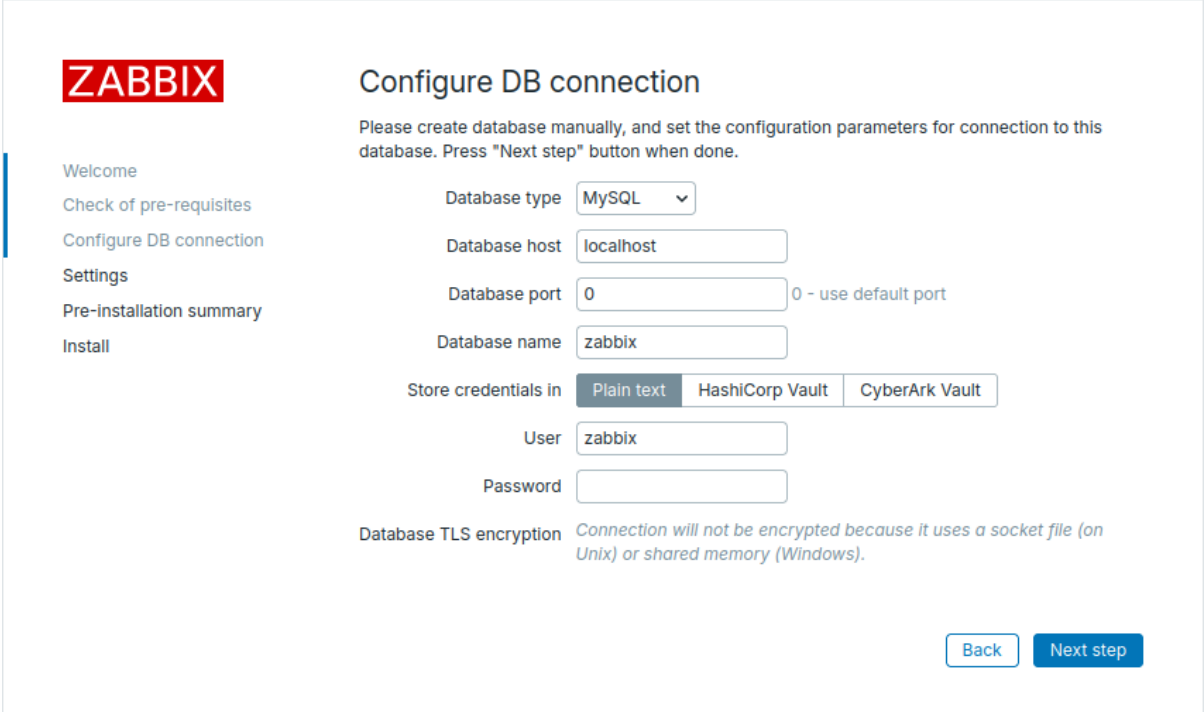
```
memory_limit = 128M ; Minimum value
post_max_size = 16M ; Minimum value
upload_max_filesize = 2M ; Minimum value
max_execution_time = 300 ; Minimum value, values 0 and -1 are also allowed
max_input_time = 300 ; Minimum value, values 0 and -1 are also allowed
session.auto_start = 0 ; Must be disabled
arg_separator.output = "&" ; "&amp;" is also allowed
```

Attention:
If the web server user or group needs to be changed, verify that the new user or group has the required permissions for the PHP session directory; otherwise, Zabbix setup may be unable to continue.

If all mandatory pre-requisites are met, proceed to the next step.

Configure DB connection

Enter the details for connecting to the database.



If the *Database host* field is left empty, the default Unix-domain socket is used.

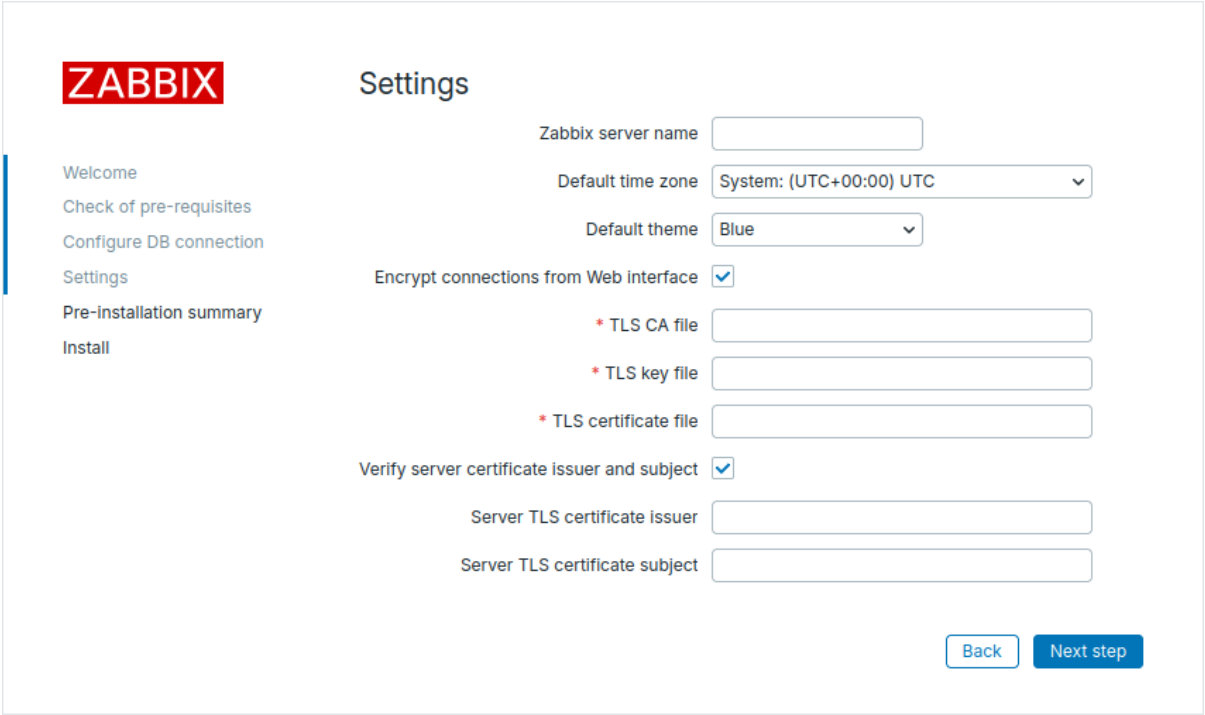
For MySQL, entering `localhost` in the *Database host* field also results in a connection via the default Unix-domain socket. The wizard does not provide an option to configure the database socket. To use a custom socket, configure it in the Zabbix server settings (for example, with the DBSocket= parameter in the Zabbix server configuration file (`zabbix_server.conf`)). This keeps the frontend consistent with the server-to-database connection settings.

For PostgreSQL, the *Database host* field can contain multiple `host:port` entries separated by commas. The string is split into parts and each port is validated in the same way as the *Database port* field. After successful validation, the string is parsed into hosts and ports to construct the connection string. If a port is not defined for a host in the comma-separated list, the value of *Database port* is used as the default port. If a socket path is specified (for example, `/var/run/pgbouncer`), that Unix-domain socket is used instead.

To configure a secure connection to the database, select the *Database TLS encryption* checkbox and proceed to Frontend to the database to configure the fields. The *Database TLS encryption* checkbox is disabled in the following cases: - For MySQL, *Database host* is set to *localhost*. - For PostgreSQL, the value of the *Database host* field begins with a slash, or the field is empty.

To configure Zabbix to retrieve secrets from HashiCorp Vault, select *HashiCorp Vault* for *Store credentials in* and proceed to HashiCorp configuration to fill in the expanded fields. To configure Zabbix to retrieve secrets from CyberArk Vault, select *CyberArk Vault* for *Store credentials in* and proceed to CyberArk configuration to fill in the expanded fields.

Settings



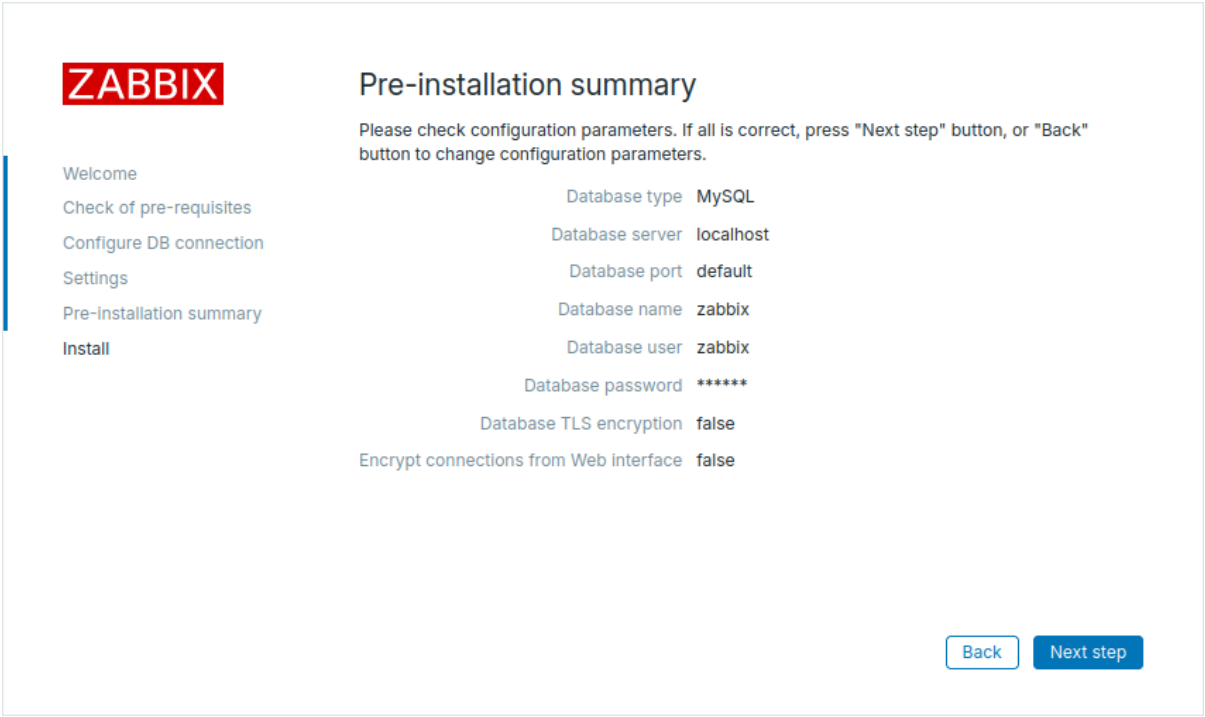
Entering a name for the Zabbix server is optional; however, if submitted, it will be displayed in the menu bar and page titles.

Set the default time zone and theme for the frontend.

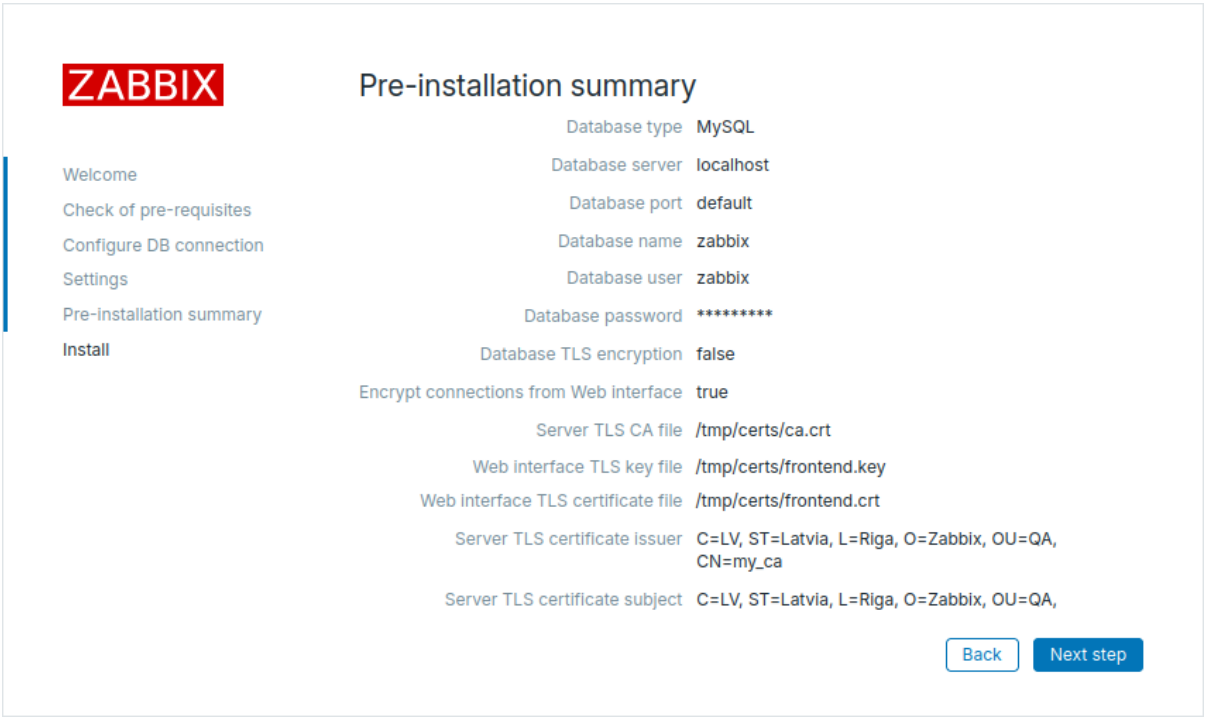
To configure TLS connections between Zabbix frontend and Zabbix server, select the *Encrypt connections from Web interface* checkbox and proceed to configuring the TLS connection to fill in the expanded fields.

Pre-installation summary

Review a summary of settings.



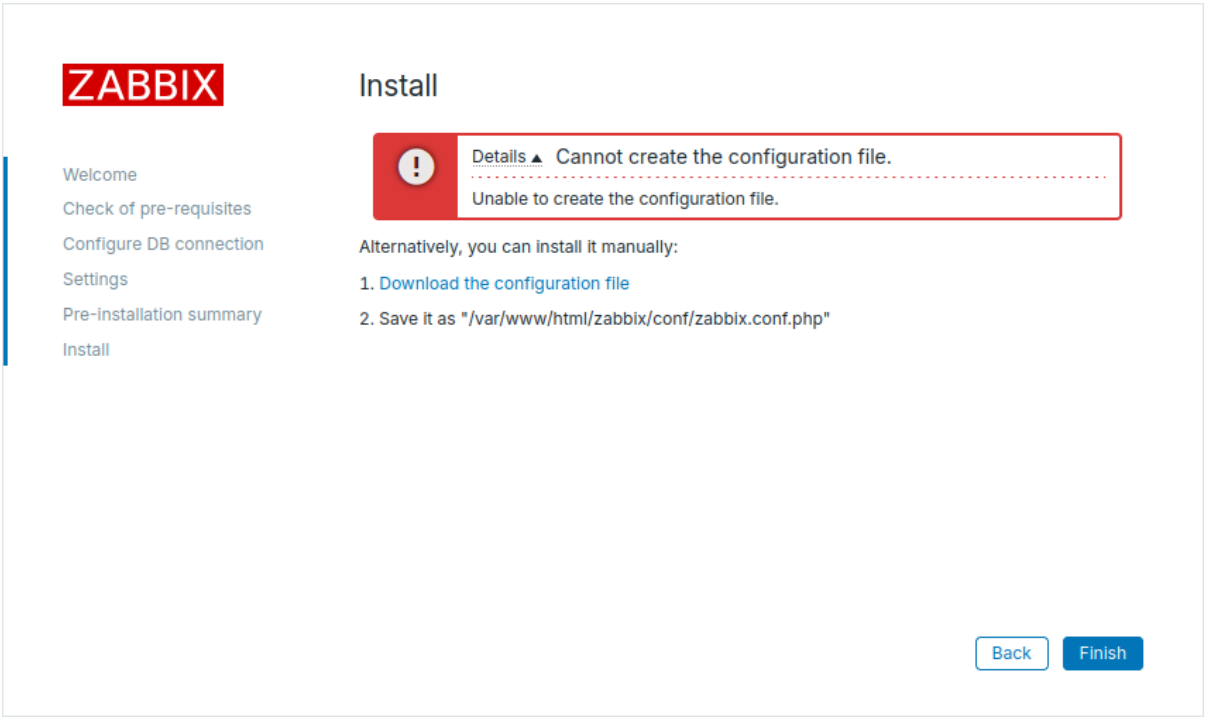
If TLS is configured, the wizard will display its details.



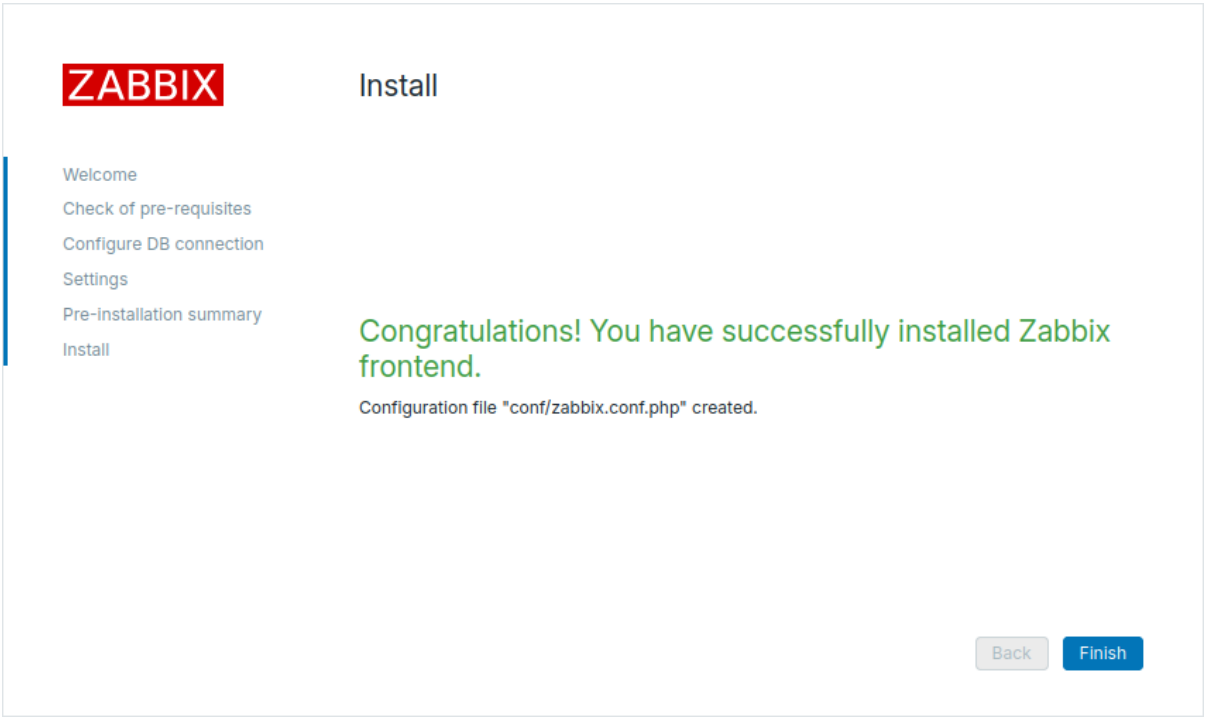
Install

If the web server user has write access to the `conf/` directory, the configuration file will be saved automatically, and you can proceed to the next step right away.

If the web server user does not have write access to the `conf/` directory, you must download the configuration file and place it under the `conf/` directory in the web server's HTML documents directory where you copied the Zabbix PHP files.

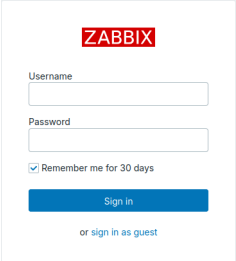


Finish the installation.



Log in

Login with the default credentials: - User name: *Admin* - Password: *zabbix*



Frontend configuration file

After installation, you can reconfigure the Zabbix frontend by updating parameters in the `zabbix.conf.php` frontend configuration file.

The default path to the configuration file depends on the operating system and Zabbix installation method:

- For installation from packages: `/usr/share/zabbix/ui/conf/zabbix.conf.php`
- For installation from sources: `/var/www/html/zabbix/conf/zabbix.conf.php`
- For appliance: `/etc/zabbix/web/zabbix.conf.php`

See the configuration file options for details on configuring the Zabbix frontend.

Quickstart guides

After you have installed Zabbix and its web interface, it is time to start monitoring.

Use these quick guides to start monitoring popular resources:

- Linux server monitoring
- Apache web server monitoring
- MySQL server monitoring

Quickstart guides focus on using official Zabbix templates for monitoring. A template is a pre-configured set of monitoring elements. A template allows you to scale monitoring—you can test a template on a single monitoring target (device, app, etc), tweak it, and then apply it to multiple similar monitoring targets.

While templates are fine for out-of-the-box monitoring, you may also be interested in customizing its elements.

To better understand the basic elements that make up a template, consult Basic configuration.

Basic configuration

These guides introduce basic elements of Zabbix configuration:

- Login and configuring user
- New host
- New item
- New trigger
- Receiving problem notification
- New template

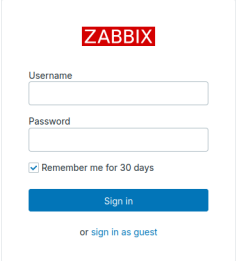
Login and configuring user

Overview

In this section, you will learn how to log in and set up a system user in Zabbix.

Login

Open the Zabbix frontend URL in a browser: - For Apache: http://<server_ip_or_name>/zabbix - For NGINX: http://<server_ip_or_name>



This is the Zabbix welcome screen. Enter the user name **Admin** with password **zabbix** to log in as a Zabbix superuser. Access to all menu sections will be granted.

For security reasons, it is strongly recommended to change the default password for the Admin account immediately after the first login.

Persistent login

To stay logged in for up to 30 days, select *Remember for 30 days* before clicking *Sign in*.

Remember for 30 days enabled:

- Your session remains active for 30 days.
- *Auto-logout* is overridden, keeping you signed in until the period ends.
- You will be auto-logged in on future visits within 30 days without re-entering credentials.

Remember for 30 days disabled:

- Any previously enabled auto-login is cleared.
- Session will expire according to the configured *Auto-logout* interval.

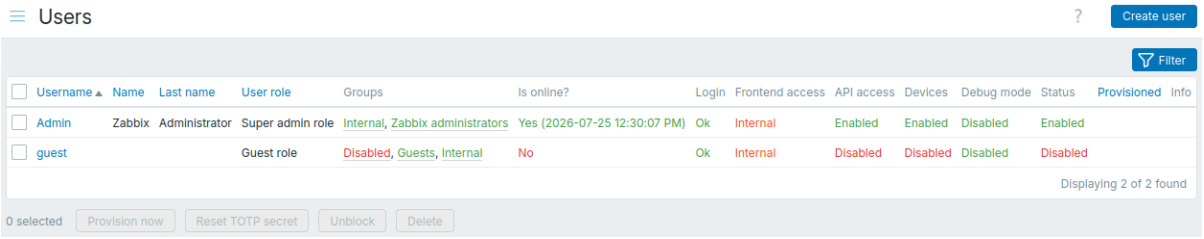
Protection against brute force attacks

In case of five consecutive failed login attempts, Zabbix interface will pause for 30 seconds in order to prevent brute force and dictionary attacks.

The IP address of a failed login attempt will be displayed after a successful login.

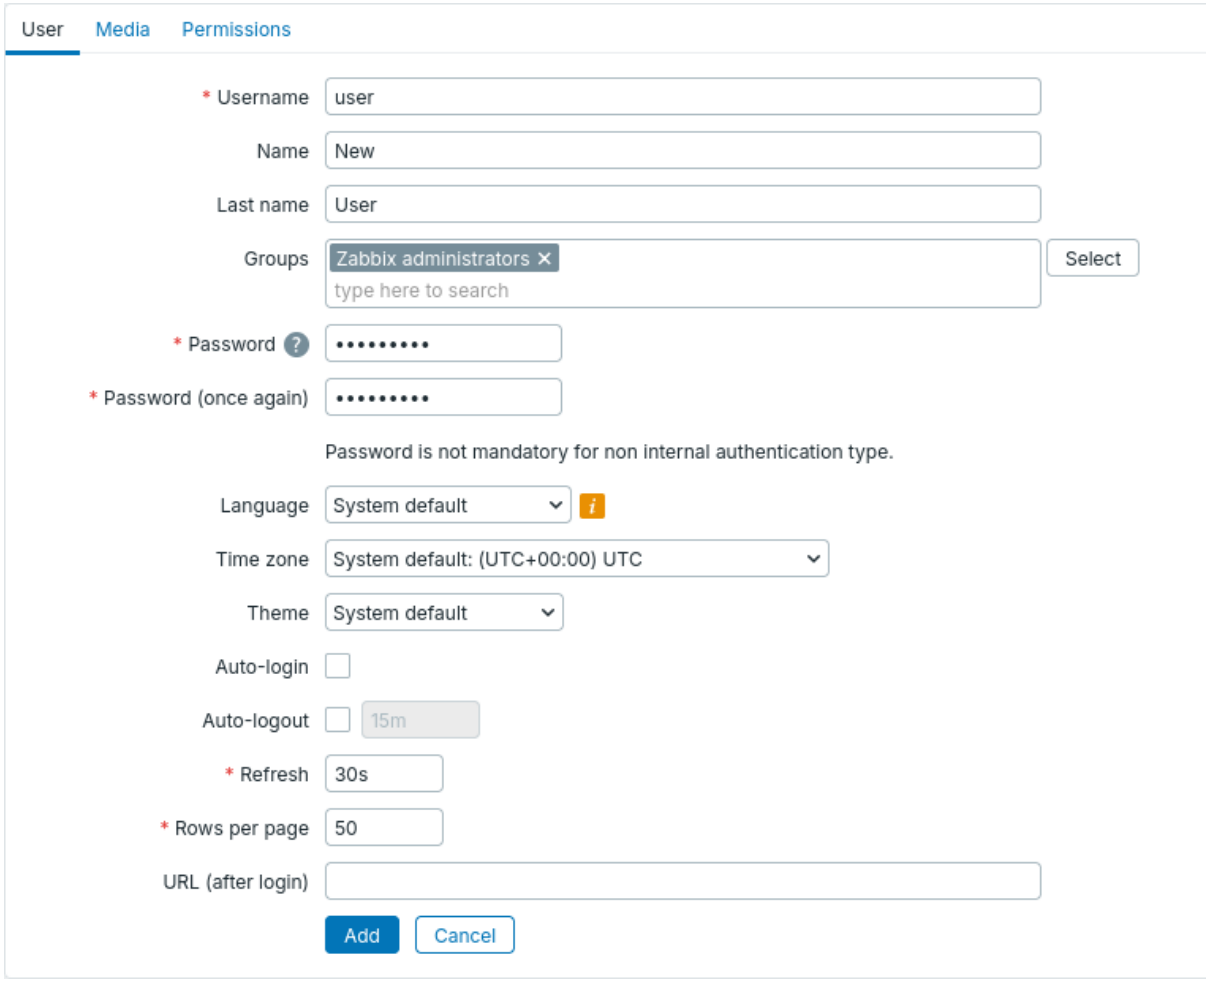
Adding user

To view information about users, go to *Users* > *Users* in the sidebar's vertical menu.



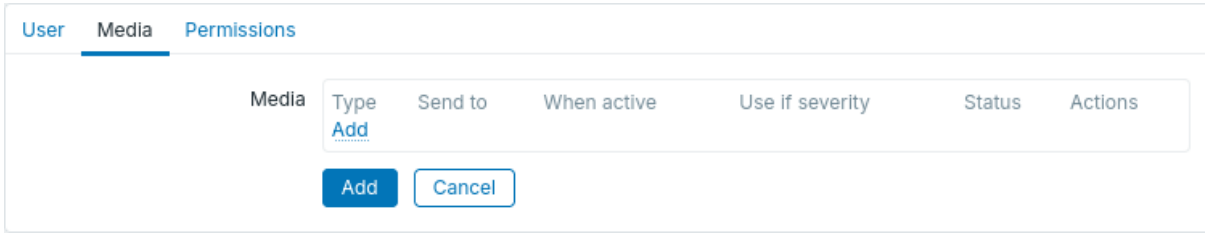
To add a new user, select *Create user* in the upper-right corner.

In the new user form, make sure to add your user to one of the existing user groups, for example 'Zabbix administrators'.



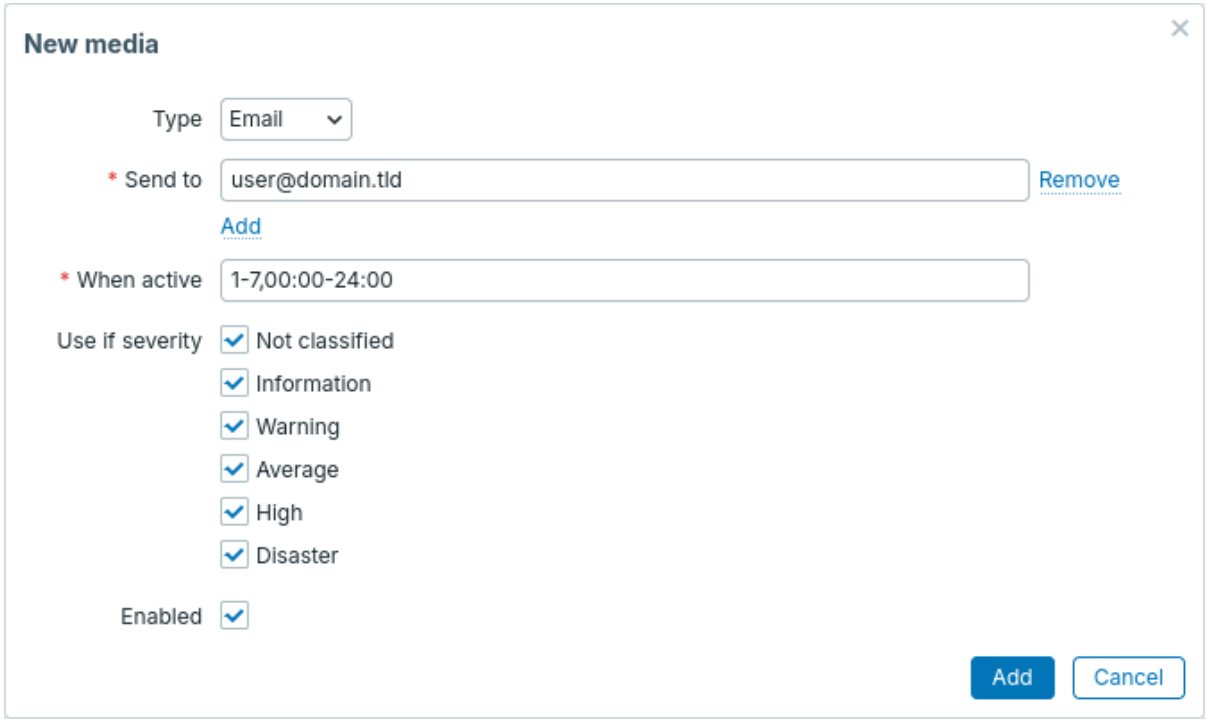
All mandatory input fields are marked with a red asterisk. For details about input fields on this configuration form, refer to the User settings page.

By default, new users have no media (notification delivery methods) defined for them. To create one, go to the *Media* tab and click *Add*.



In the pop-up, enter the user's email address.
 All media types are disabled by default and can be enabled in *Alerts* > *Media types*.

You can specify a time period when the medium will be active (see Time period specification page for a description of the format). By default, a medium is always active. You can also customize trigger severity levels for which the medium will be active, but leave all of them enabled for now.



Click *Add* to save the medium, then go to the *Permissions* tab.

Permissions tab has a mandatory field *Role*. The role determines which frontend elements the user can view and which actions he is allowed to perform. Click *Select* and select one of the roles from the list. For example, select *Admin role* to allow access to all Zabbix frontend sections, except *Administration*. Later on, you can modify permissions or create more user roles. Upon selecting a role, permissions will appear in the same tab:

Users

User | Media | Permissions

* Role: Admin role ✕ | Select

User type: Admin

Permissions

User group permissions

Group	Type	Permissions
All groups	Hosts	None
All groups	Templates	None

Proxies: None

Proxy groups: None

Permissions can be assigned for user groups only.

Access to UI elements

Dashboards: Dashboards

Monitoring: Problems | Hosts | Latest data | Maps | Discovery

APM: Traces | Metrics | Logs

Services: Services | SLA | SLA report

Inventory: Overview | Hosts

Reports: System information | Scheduled reports | Availability report | Top 100 triggers | Notifications

Data collection: Template groups | Host groups | Templates | Hosts | Maintenance | Discovery

Alerts: Trigger actions | Service actions | Discovery actions | Autoregistration actions | Internal actions

User profile settings

Redirect URL after login

Access to services

Read-write access to services: None

Read-only access to services: All

Access to modules

Action log | Clock | Discovery status | Favorite graphs | Favorite maps | Gauge | Geomap | Graph | Graph (classic) | Graph prototype | Honeycomb | Host availability | Host card | Host navigator | Item card | Item history | Item navigator | Item value | Map | Map navigation tree | Pie chart | Problem hosts | Problems | Problems by severity | Scatter plot | SLA report | System information | Top hosts | Top items | Top triggers | Trigger overview | URL | Web monitoring

Access to API

Disabled

Devices

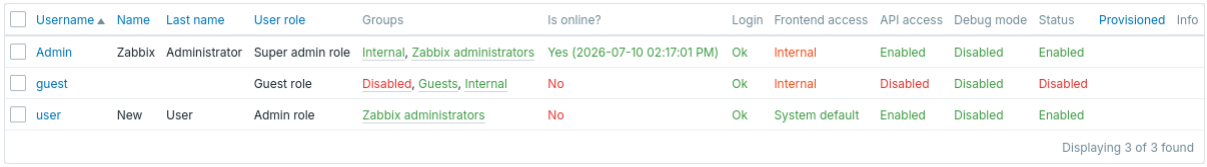
Enabled | Manage own devices

Access to actions

Create and edit dashboards | Create and edit maps | Create and edit maintenance | Add problem comments | Change severity | Acknowledge problems | Suppress problems | Close problems | Execute scripts | Manage API tokens | Manage scheduled reports | Manage SLA | Invoke "Execute now" on read-only hosts | Change problem ranking | Create and edit own media | Select "Server" for monitoring and discovery

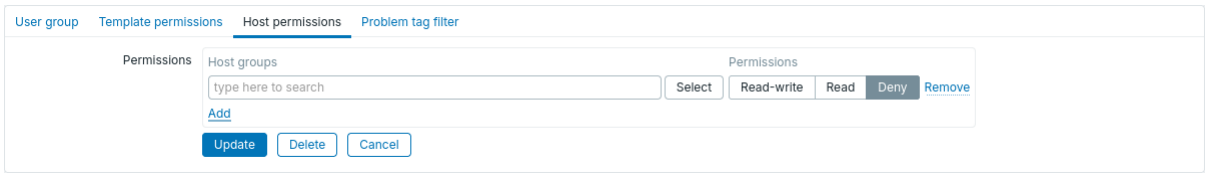
Add | Cancel

Click *Add* in the user properties form to save the user. The new user appears in the userlist.

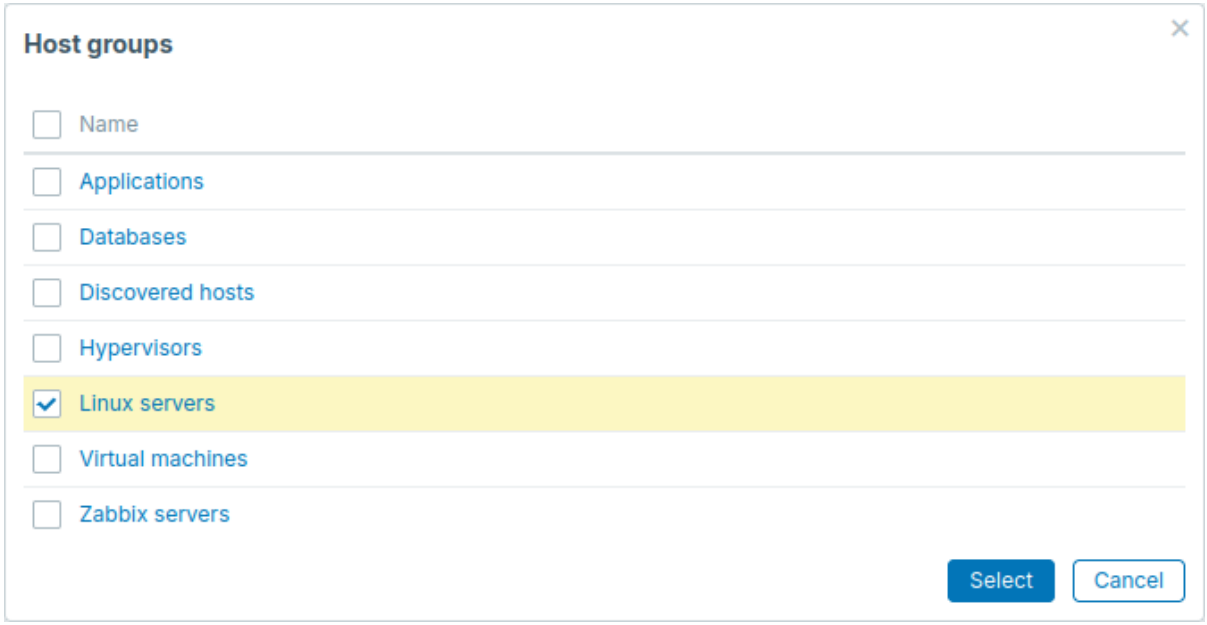


Adding permissions

By default, a new user has no permissions to access hosts and templates. To grant the user rights, click the group of the user in the *Groups* column (in this case - *Zabbix administrators*). In the *User groups* properties form, go to the *Host permissions* tab to assign permissions to host groups. Click Add for the host group selection field to be displayed:



Then click *Select* next to the field to see the list of the host groups. This user is to have read-only access to *Linux servers* group, so mark the appropriate checkbox in the list and click *Select* to confirm your choice.



Click the *Read* button to set the permission level and then *Update* to save the changes made to the user group configuration.

To grant permissions to templates, you will need to switch to the *Template permissions* tab and specify template groups. The steps are identical to assigning permissions to host groups. An overview of templates is available in the New template section of this Quickstart.

Attention:
In Zabbix, access rights to hosts and templates are assigned to user groups, not individual users.

Done! You may try to log in using the credentials of the new user.

New host

Overview

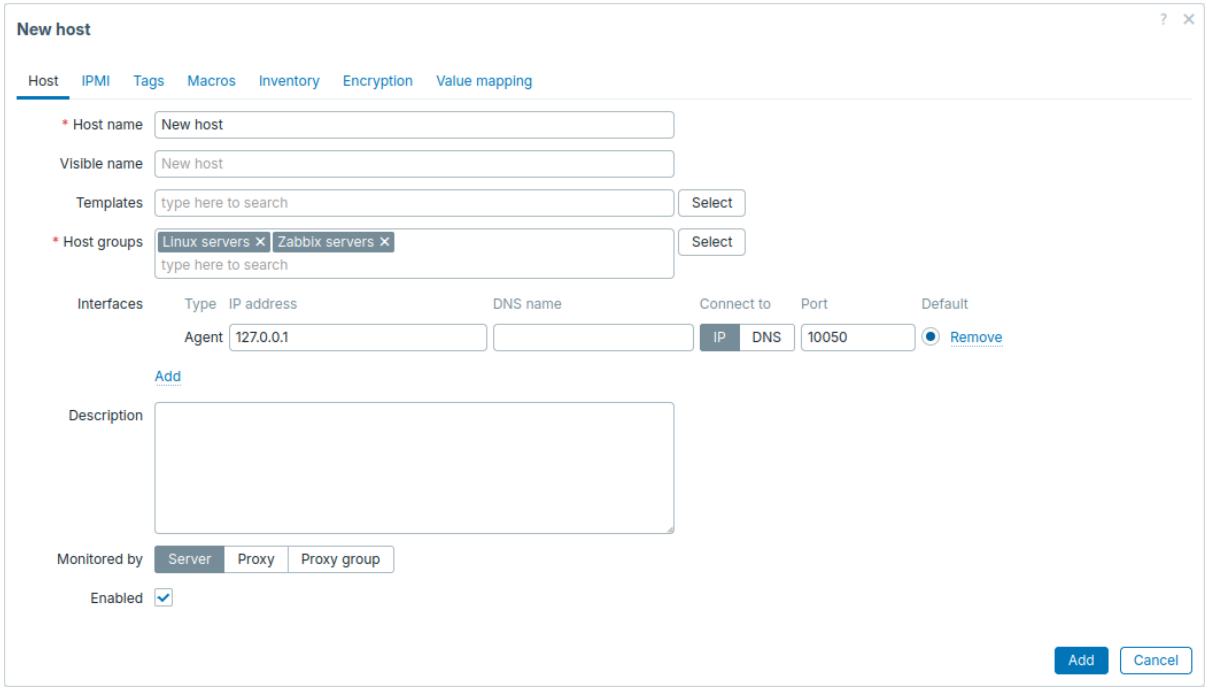
In this section you will learn how to set up a new host.

A host in Zabbix is a networked entity (physical, virtual) that you wish to monitor. The definition of what can be a ”host” in Zabbix is quite flexible. It can be a physical server, a network switch, a virtual machine or some application.

Adding host

Information about configured hosts in Zabbix is available in the *Data collection* > *Hosts* menu section, as well as the *Monitoring* > *Hosts* menu section. There is already one pre-defined host, called ”Zabbix server”, but we want to learn adding another.

To add a new host, click *Create host*. A host configuration form is opened.



All mandatory input fields are marked with a red asterisk.

The bare minimum to enter here is:

Host name

- Enter a host name. Alphanumerics, spaces, dots, dashes and underscores are allowed.

Host groups

- Select one or several existing groups by clicking the *Select* button or enter a non-existing group name to create a new group.

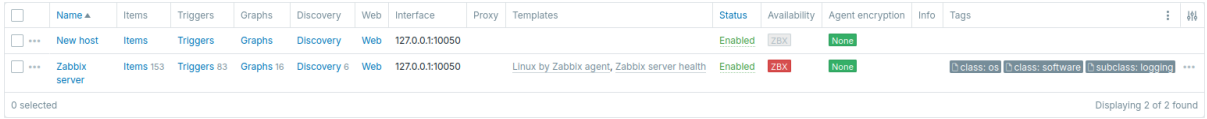
Note:
All access permissions are assigned to host groups, not individual hosts. That is why a host must belong to at least one group.

Interfaces: IP address

- Although not a required field technically, a host interface is necessary for collecting certain metrics. To use Zabbix agent passive checks, specify the agent's IP or DNS in this field. Note that you should also specify Zabbix server's IP or DNS in the Zabbix agent configuration file 'Server' directive. If Zabbix agent and Zabbix server are installed on the same machine, you need to specify the same IP/DNS in both places.

Other options may be left unchanged, using their default values.

When done, click *Add*. Your new host should be visible in the host list.



The *Availability* column contains indicators of host availability per each interface. When a Zabbix agent interface is defined, we can use the agent availability icon (labeled *ZBX*) to determine the host's availability:

- ZBX - host status has not been established; no metric check has happened yet.
- ZBX - host is available, a metric check has been successful.
- ZBX - host is unavailable, a metric check has failed (move your mouse cursor over the icon to see the error message). There might be some error with communication, possibly caused by incorrect interface credentials. Check that Zabbix server is running, and try refreshing the page later as well.

New item

Overview

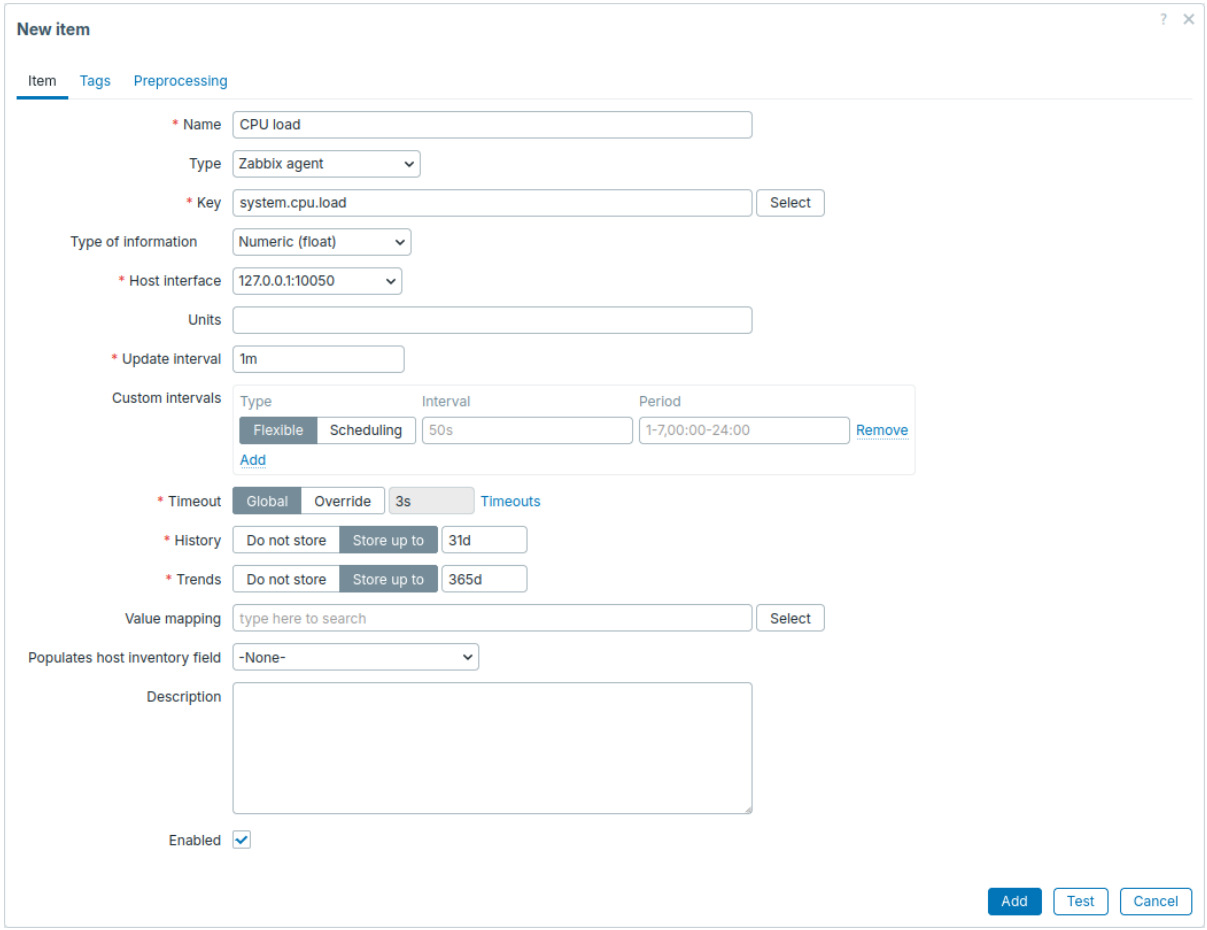
In this section, you will learn how to set up an item.

Items are the basis of gathering data in Zabbix. Without items, there is no data - because only an item defines a single metric or what kind of data to collect from a host.

Adding item

All items are grouped around hosts. That is why to configure a sample item we go to *Data collection* > *Hosts* and find the "New host" we have created.

Click the *Items* link in the row of "New host", and then click *Create item*. An item configuration form is opened.



All mandatory input fields are marked with a red asterisk.

For our sample item, the essential information to enter is:

Name

- Enter *CPU load* as the value. This will be the item name displayed in lists and elsewhere.

Key

- Manually enter *system.cpu.load* as the value. This is the technical name of an item that identifies the type of information that will be gathered. The particular key is just one of pre-defined keys that come with Zabbix agent.

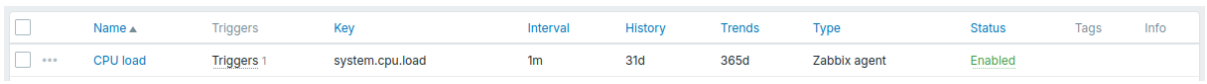
Type of information

- This attribute defines the format of the expected data. For the *system.cpu.load* key, this field will be automatically set to *Numeric (float)*.

Note:
You may also want to reduce the number of days item history will be kept, to 7 or 14. This is good practice to relieve the database from keeping lots of historical values.

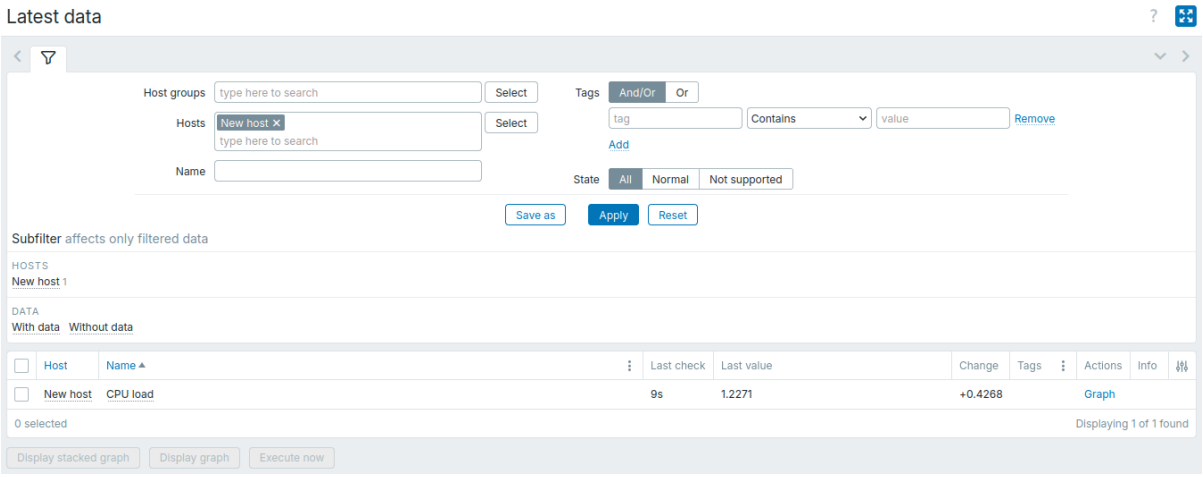
Other options may be left unchanged, using their default values.

When done, click *Add*. The new item should appear in the item list.



Viewing data

With an item defined, you might be curious if it is actually gathering data. For that, go to *Monitoring* > *Latest data*, select 'New host' in the filter and click *Apply*.



With that said, it may take up to 60 seconds for the first data to arrive. That, by default, is how often the server reads configuration changes and picks up new items to execute.

If you see no value in the 'Change' column, maybe only one value has been received so far. Wait 30 seconds for another value to arrive.

If you do not see information about the item as in the screenshot, make sure that:

- you have filled out the item 'Key' and 'Type of information' fields exactly as in the screenshot;
- both the agent and the server are running;
- host status is 'Monitored' and its availability icon is green;
- the host selected in the host filter is correct;
- the item is enabled.

Graphs

With the item working for a while, it might be time to see something visual. Simple graphs are available for any monitored numeric item without any additional configuration. These graphs are generated on runtime.

To view the graph, go to *Monitoring* > *Latest data* and click the *Graph* link next to the item.



New trigger

Overview

In this section you will learn how to set up a trigger.

Items only collect data. To automatically evaluate incoming data we need to define triggers. A trigger contains an expression that defines a threshold of what is an acceptable level for the data.

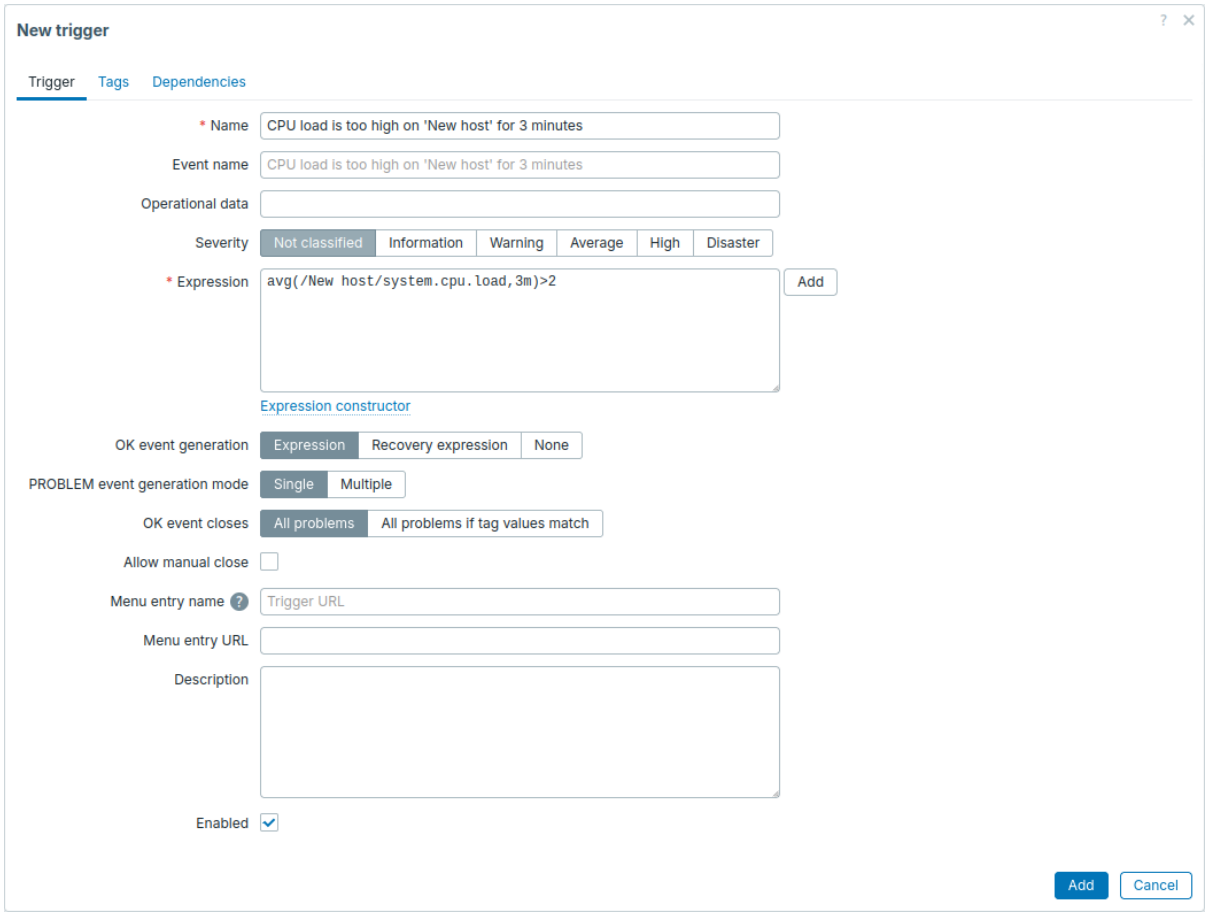
If that level is surpassed by the incoming data, a trigger will "fire" or go into a 'Problem' state - indicating that something has happened that may require attention. If the level is acceptable again, trigger returns to an 'Ok' state.

Adding trigger

To configure a trigger for our item:

1. Go to *Data collection* > *Hosts*.
2. Click the *Triggers* link in the row of 'New host'.
3. Click *Create trigger*.

A trigger configuration form is opened.



For our trigger, the essential information to enter here is:

Name

- Enter *CPU load is too high on 'New host' for 3 minutes* as the value. This will be the trigger name displayed in lists and elsewhere.

Expression

- Enter: `avg(/New host/system.cpu.load,3m)>2`

This is the trigger expression. Make sure that the expression is entered right, down to the last symbol. The item key here (`system.cpu.load`) is used to refer to the item. This particular expression basically says that the problem threshold is exceeded when the CPU load average value for 3 minutes is over 2. You can learn more about the syntax of trigger expressions.

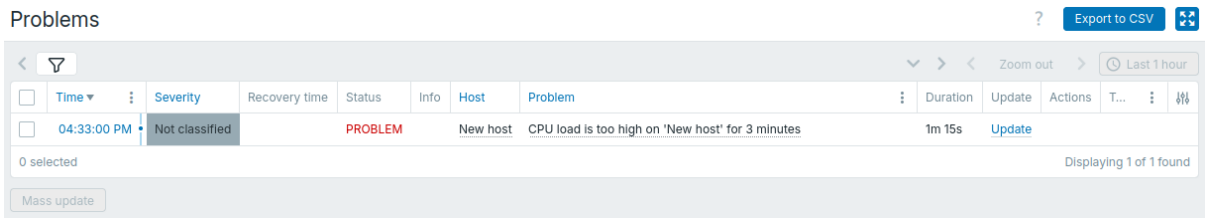
Other options may be left unchanged, using their default values.

When done, click *Add*. The new trigger should appear in the trigger list.

Displaying trigger state

With a trigger defined, you might be interested to see its state.

If the CPU load has exceeded the threshold level you defined in the trigger, the problem will be displayed in *Monitoring* > *Problems*.



The flashing in the status column indicates a recent change of trigger status, one that has taken place in the last 30 minutes.

Receiving problem notification

Overview

In this section you will learn how to set up alerting in the form of notifications in Zabbix.

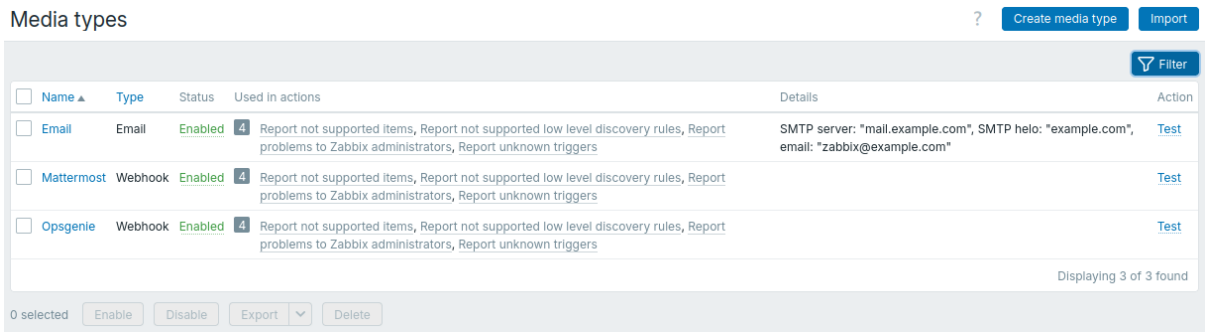
With items collecting data and triggers designed to ”fire” upon problem situations, it would also be useful to have some alerting mechanism in place to notify about important events even when Zabbix frontend is not being actively checked.

This is what notifications do. Email being the most popular delivery method for problem notifications, we will learn how to set up an email notification.

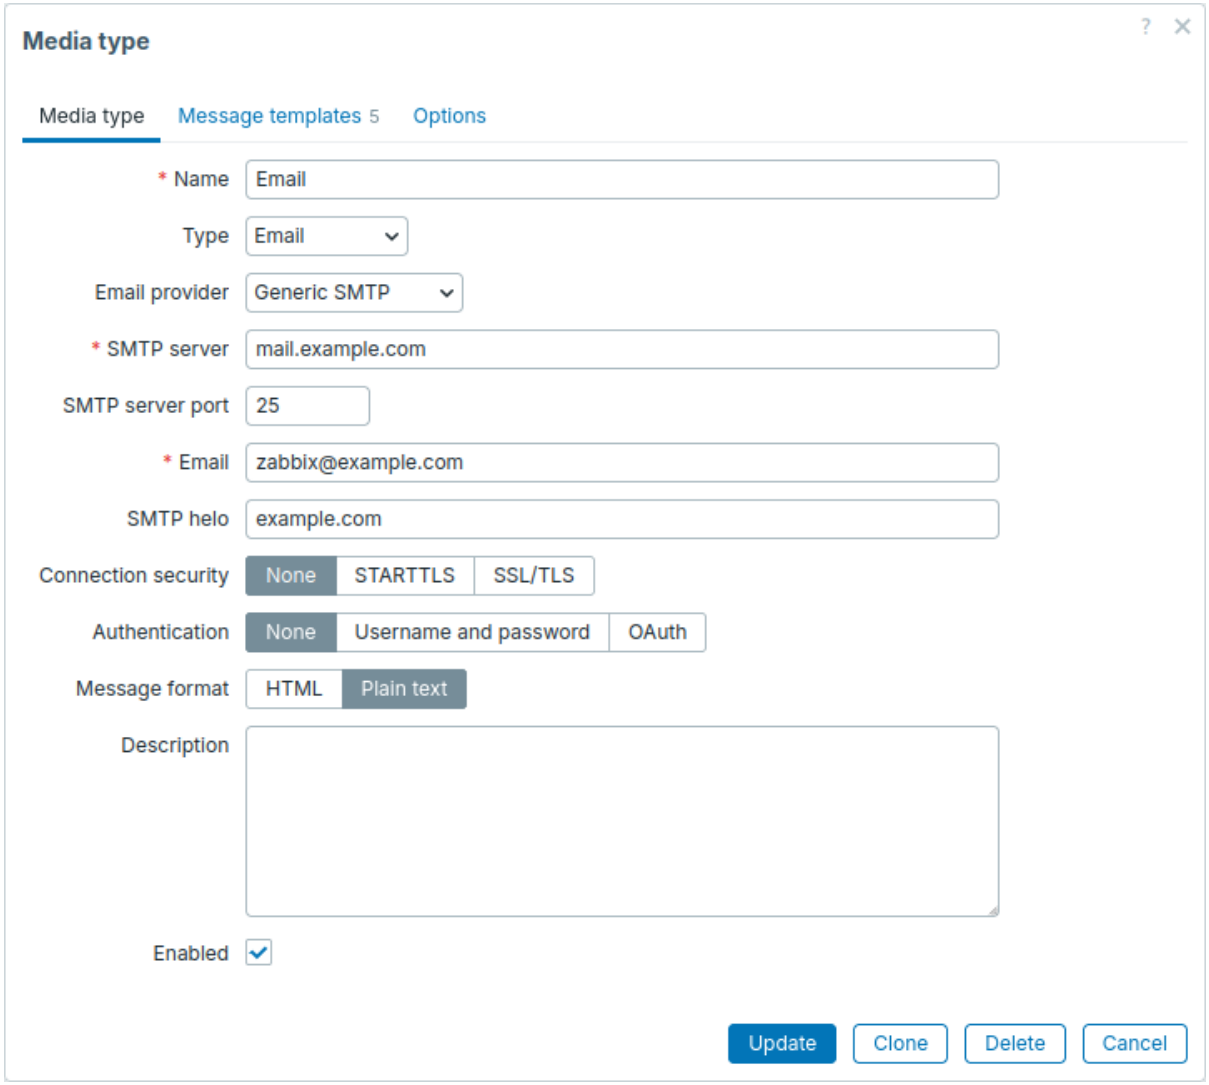
Email settings

Initially there are several predefined notification delivery methods in Zabbix. Email is one of those.

To configure email settings, go to *Alerts* > *Media types* and click *Email* in the list of pre-defined media types.



An email settings configuration form is opened.



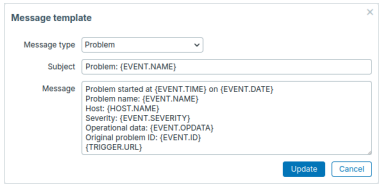
All mandatory input fields are marked with a red asterisk.

In the *Media type* tab, set the values of *SMTP server*, *SMTP helo*, and *Email* to the appropriate for your environment.

Note:
The value in the *Email* field will be used as the 'From' address for the notifications sent from Zabbix.

Next, it is required to define the content of the problem message. The content is defined by means of a message template, configured in the *Message templates* tab.

Click *Add* to create a message template, and select *Problem* as the message type.

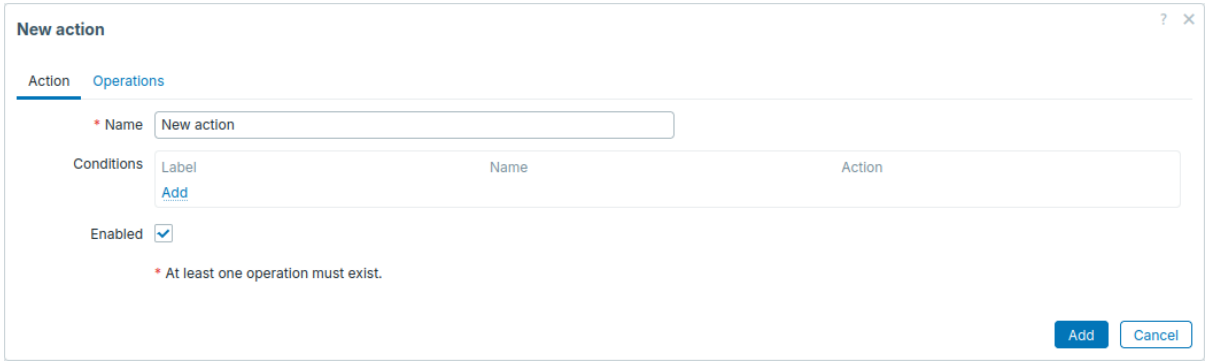


Click *Add* when ready and save the form.

Now you have configured *Email* as a working media type. The media type must also be linked to users by defining specific delivery addresses (like we did when configuring a new user); otherwise, it will not be used.

New action

Delivering notifications is one of the things actions do in Zabbix. Therefore, to set up a notification, go to *Alerts* > *Actions* > *Trigger actions* and click *Create action*.

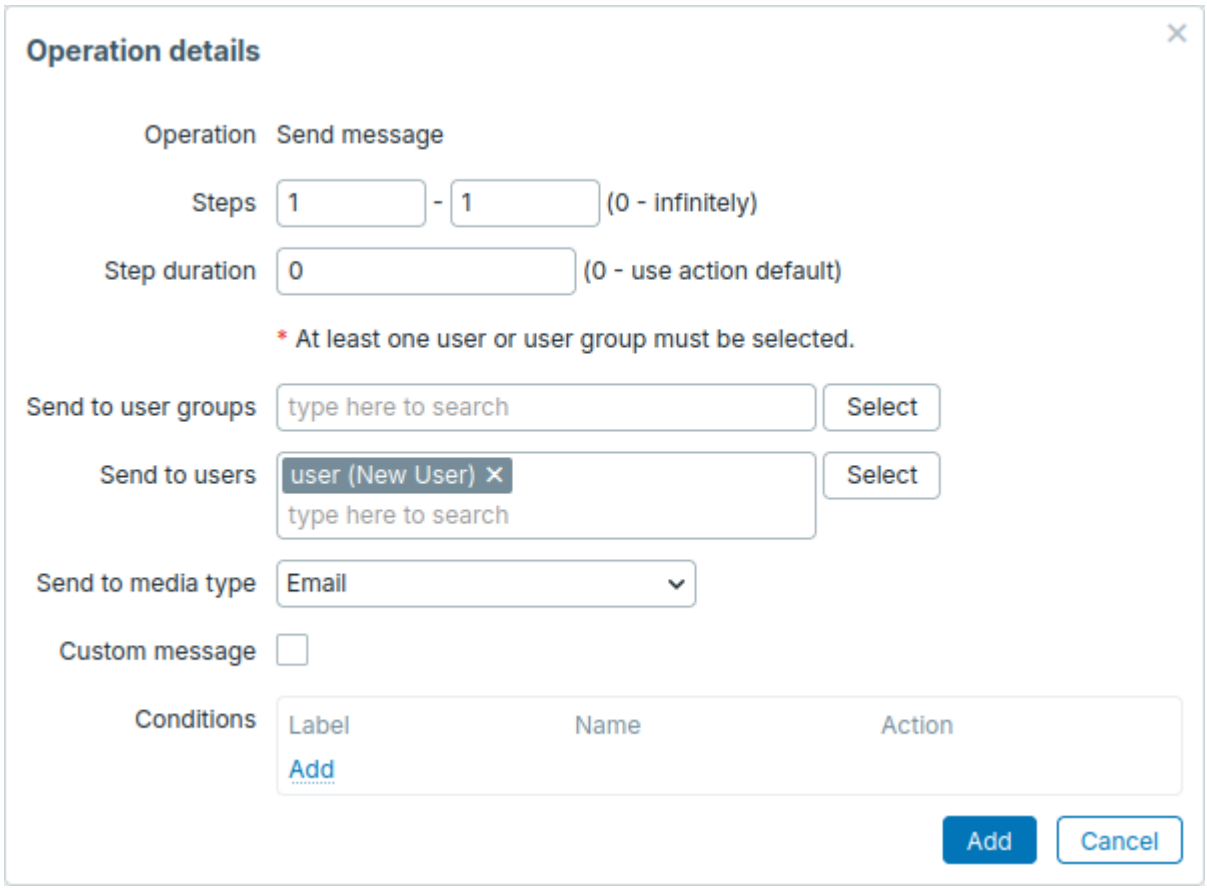


All mandatory input fields are marked with a red asterisk.

In this form, enter a name for the action.

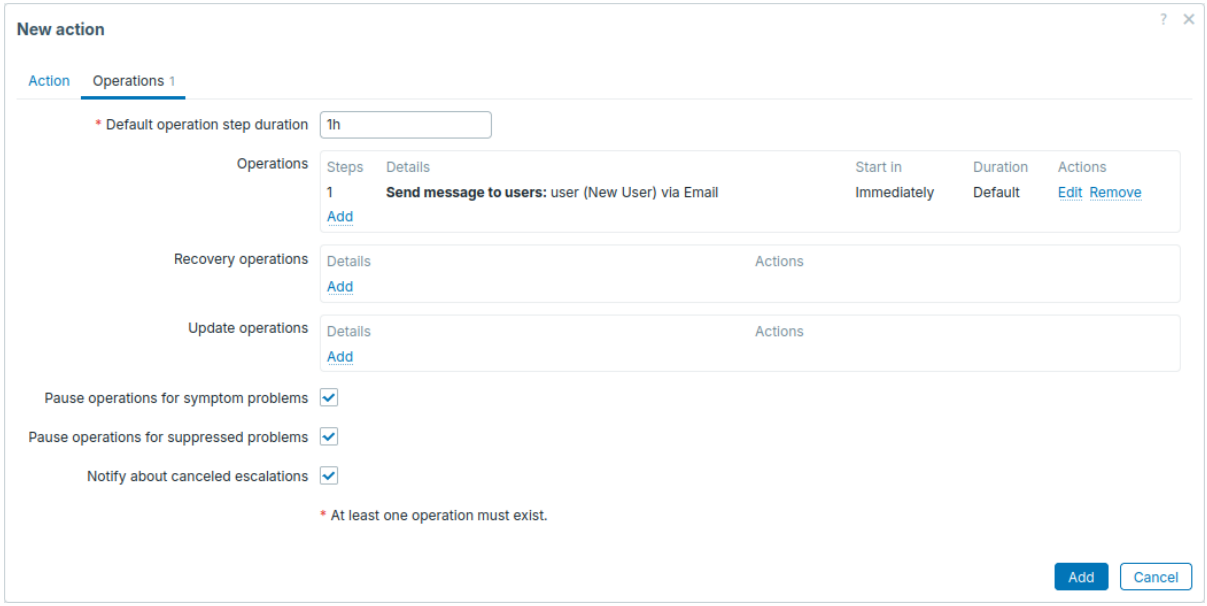
In the most simple case, if we do not add any more specific conditions, the action will be taken upon any trigger change from 'Ok' to 'Problem'.

We still should define what the action should do - and that is done in the *Operations* tab. Click *Add* in the *Operations* block, which opens a new operation form.



All mandatory input fields are marked with a red asterisk.

Here, click *Select* in the *Send to Users* block and select the user ('user') we have defined. Select "Email" as the value of *Send to media type*. When done with this, click *Add*, and the operation should be added:



That is all for a simple action configuration, so click *Add* in the action form.

Receiving notification

Now, with delivering notifications configured, it would be fun to actually receive one. To help with that, we might on purpose increase the load on our host - so that our trigger "fires" and we receive a problem notification.

Open the console on your host and run:

```
cat /dev/urandom | md5sum
```

You may run one or several of these processes.

Now go to *Monitoring* > *Latest data* and see how the values of 'CPU Load' have increased.Remember, for our trigger to "fire", the 'CPU Load' value has to go over '2' for 3 minutes running. Once it does:

- in *Monitoring* > *Problems* you should see the trigger with a flashing `Problem` state;
- you should receive a problem notification in your email.

Attention:
If notifications do not work:

- verify once again that both the email settings and the action have been configured properly
- make sure the user you created has at least read permissions on the host which generated the event, as noted in the *Adding user* step. The user, being part of the 'Zabbix administrators' user group must have at least read access to 'Linux servers' host group that our host belongs to.
- Additionally, you can check out the action log by going to *Reports* > *Action log*.

New template

Overview

In this section you will learn how to set up a template.

Previously we learned how to set up an item, a trigger and how to get a problem notification for the host.

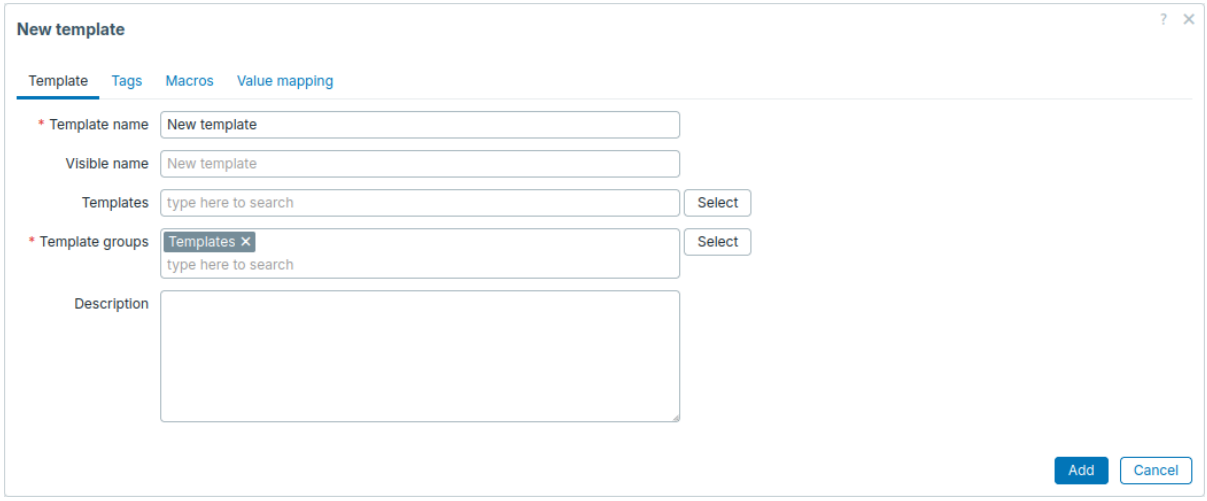
While all of these steps offer a great deal of flexibility in themselves, it may appear like a lot of steps to take if needed for, say, a thousand hosts. Some automation would be handy.

This is where templates come to help. Templates allow you to group useful items, triggers and other entities so that those can be reused again and again by applying to hosts in a single step.

When a template is linked to a host, the host inherits all entities of the template. So, basically a pre-prepared bunch of checks can be applied very quickly.

Adding template

To start working with templates, we must first create one. To do that, in *Data collection* > *Templates* click *Create template*. This will bring up a template configuration form.



All mandatory input fields are marked with a red asterisk.

The required parameters to enter here are:

Template name

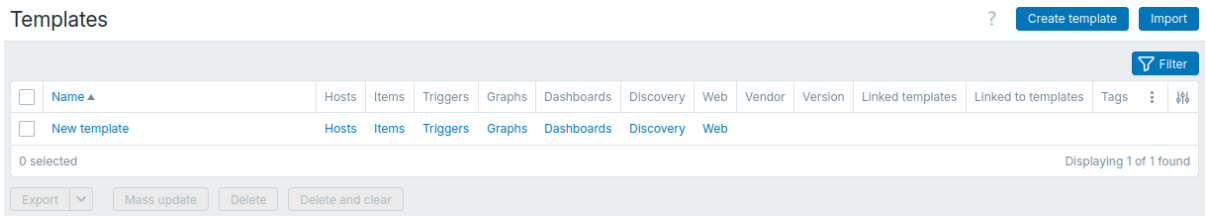
- Enter a template name. Alpha-numericals, spaces and underscores are allowed.

Template groups

- Select one or several groups by clicking *Select* button. The template must belong to a group.

Note:
Access permissions to template groups are assigned in the user group configuration on the **Template permissions** tab in the same way as host permissions. All access permissions are assigned to groups, not individual templates, that's why including the template into at least one group is mandatory.

When done, click *Add*. Your new template should be visible in the list of templates. You can also use the filter to find your template.

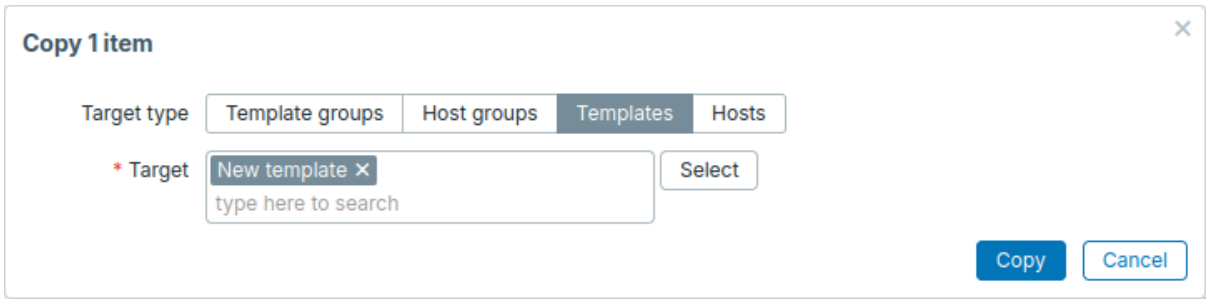


As you may see, the template is there, but it holds nothing in it - no items, triggers or other entities.

Adding item to template

To add an item to the template:

1. Go to *Data collection* > *Hosts*.
2. Locate 'New host' in the list and click the *Items* hyperlink.
3. Mark the checkbox of the 'CPU Load' item in the list.
4. Click *Copy* below the list.
5. Select the *Templates* tab.
6. Select the template to copy the item to.



All mandatory input fields are marked with a red asterisk. 7. Click *Copy*.

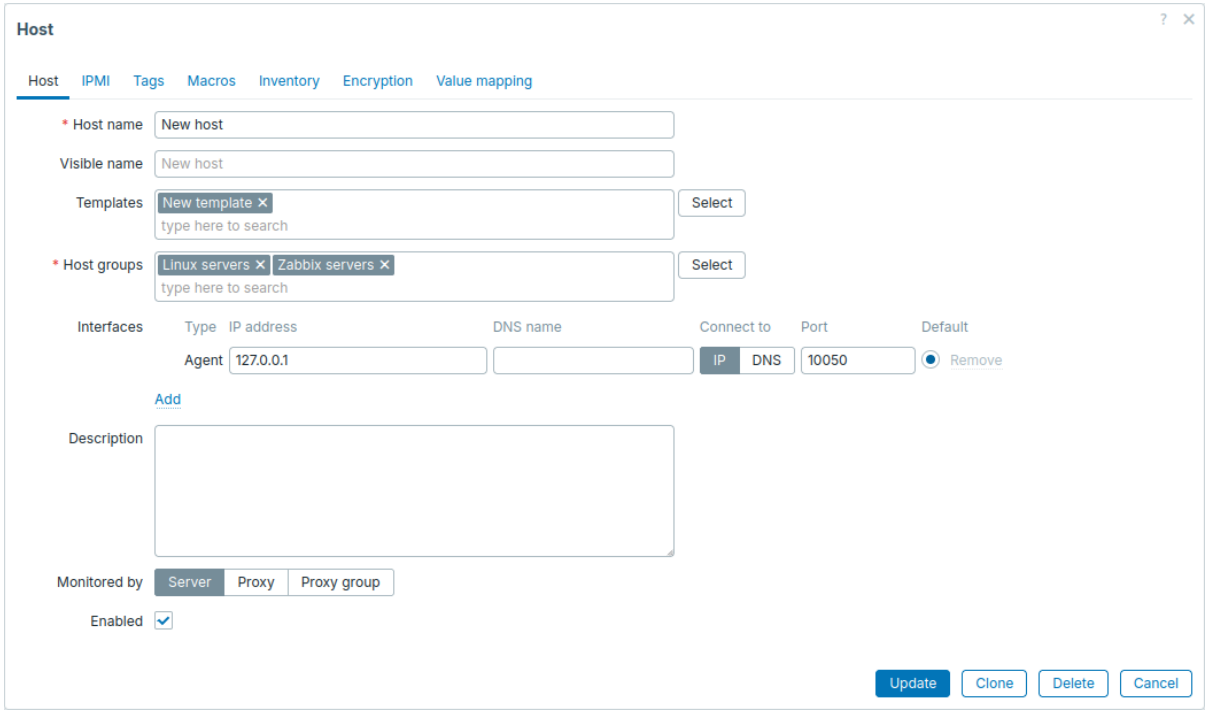
If you now go to *Data collection* > *Templates*, 'New template' should have one new item in it.

We will stop at one item only for now, but similarly you can add any other items, triggers or other entities to the template until it's a fairly complete set of entities for given purpose (monitoring OS, monitoring single application).

Linking template to host

With a template ready, it only remains to add it to a host. To do so:

1. Go to *Data collection* > *Hosts*.
2. Click 'New host' to open its configuration form and find the *Templates* field.
3. Start typing *New template* in the *Templates* field. The name of the newly created template should appear in the dropdown list.
4. Scroll down to select. See that it appears in the *Templates* field.



5. Click *Update* in the form to save the changes.

The template is now added to the host, with all entities that it holds.

This way it can be applied to any other host as well. Any changes to the items, triggers and other entities at the template level will propagate to the hosts the template is linked to.

Linking pre-defined templates to hosts

As you may have noticed, Zabbix comes with a set of predefined templates for various OS, devices and applications. To get started with monitoring very quickly, you may link the appropriate one of them to a host, but beware that these templates need to be fine-tuned for your environment. Some checks may not be needed, and polling intervals may be way too frequent.

More information about templates is available.

Apache web server monitoring

To start monitoring the Apache web server, follow this process:

1. Ensure access to Apache statistics.
2. Apply Apache monitoring template.
3. Look at latest data.

To use this guide you must have Zabbix server and the web interface installed. You must be logged in to Zabbix as the default *Admin* user.

Access to Apache stats

The goal of this step is to make sure that Apache web server statistics can be read by Zabbix.

In Apache, real-time information about the server's performance and activity is provided by the built-in `mod_status` module. When enabled, it creates a web-accessible page (usually at `/server-status`) showing performance and activity data.

To verify that `mod_status` is enabled for the Apache web server, run:

On Debian/Ubuntu	On RHEL-based systems
`apache2ctl -M | grep status`	`httpd -M | grep status`

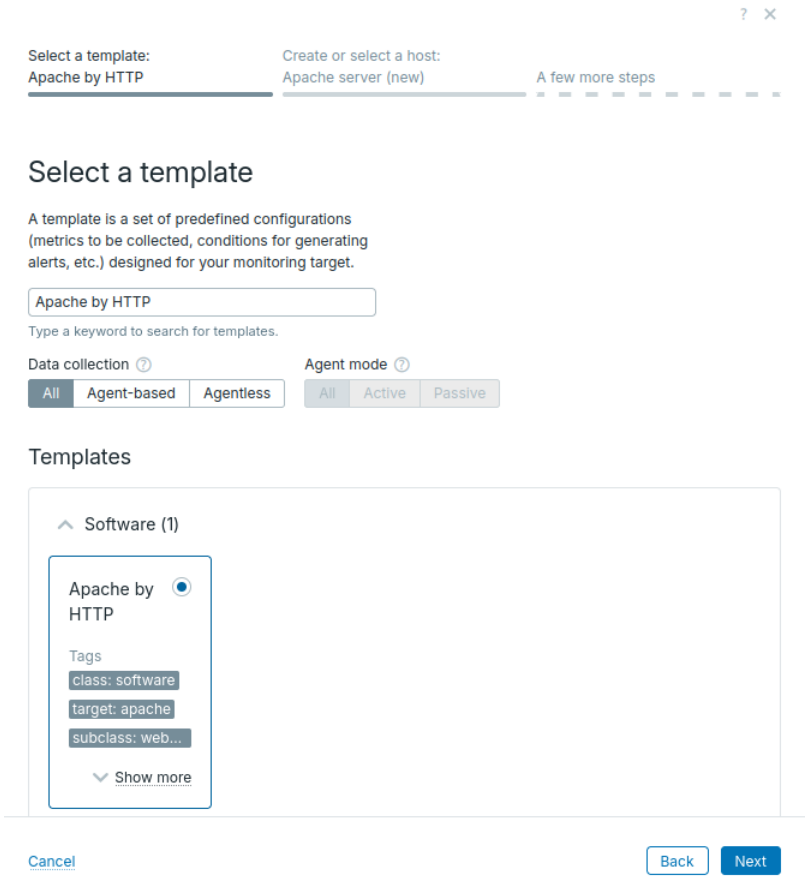
If you see "`status_module (shared)`" in the output then the status module is enabled already. If not, then you need to enable the status module.

Apply Apache monitoring template

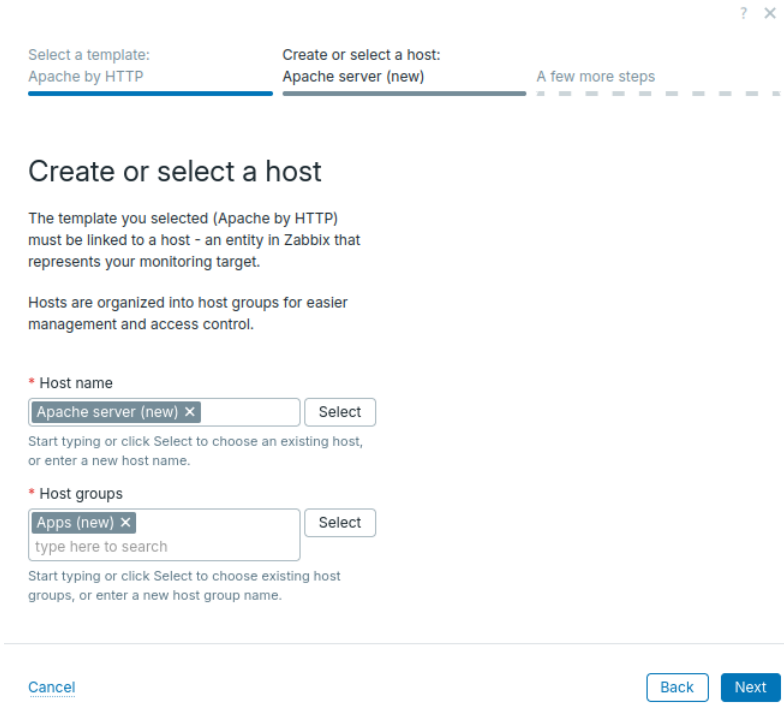
In Zabbix this template is ready to use. It contains pre-configured elements ready for data collection, visualisation and analysis.

To apply the template:

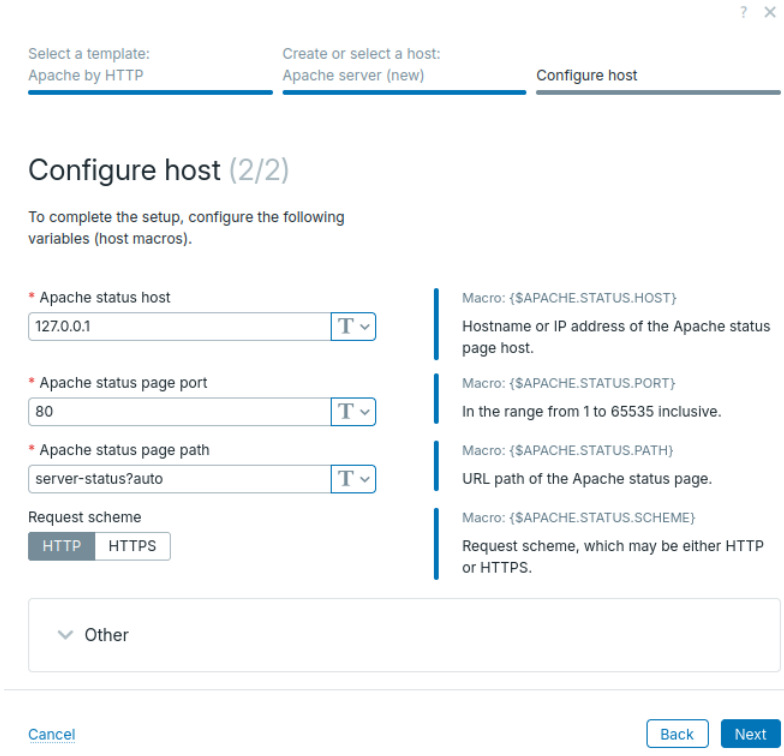
1. Start the Host Wizard in *Data collection* > *Hosts*.
2. Select the *Apache by HTTP* template.



3. Create an Apache host, add a group for it. This virtual host will contain Apache-related data.



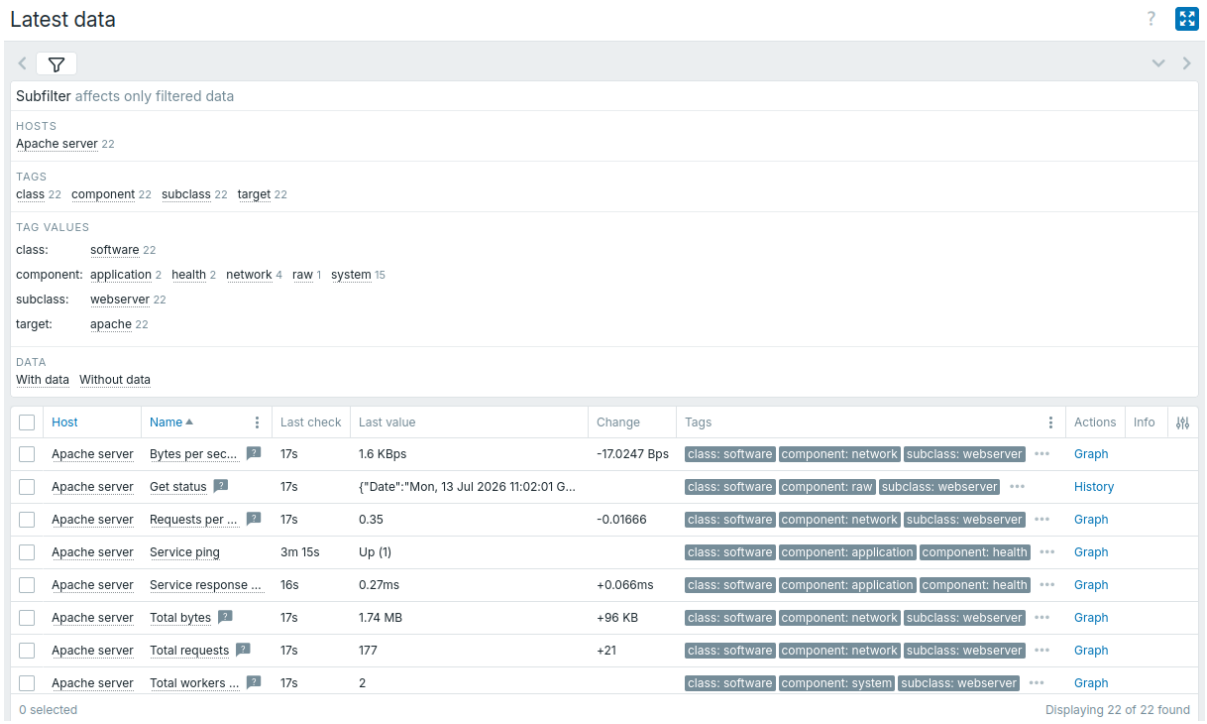
4. Configure the host. Importantly, enter the Apache web server address (127.0.0.1 on a local machine, IP address on a remote machine).



You can leave other settings to their defaults.

5. To complete the Host Wizard, click *Create* and *Finish*.

Look at latest data



Congratulations, Apache web server monitoring has been started!

Zabbix provides visualisation options and delivers alerts if any problems arise. To configure alerts in the email, see Receiving problem notification.

Linux server monitoring

To start monitoring a Linux server, follow this process:

1. Apply the Linux monitoring template.
2. Install Zabbix agent on the server.
3. Look at latest data.

To start monitoring a non-local Linux server you need to know its IP address, and the IP address of the Zabbix server.

A Zabbix agent will be used on the Linux server to monitor it. Instructions for installing the agent will be provided in the process.

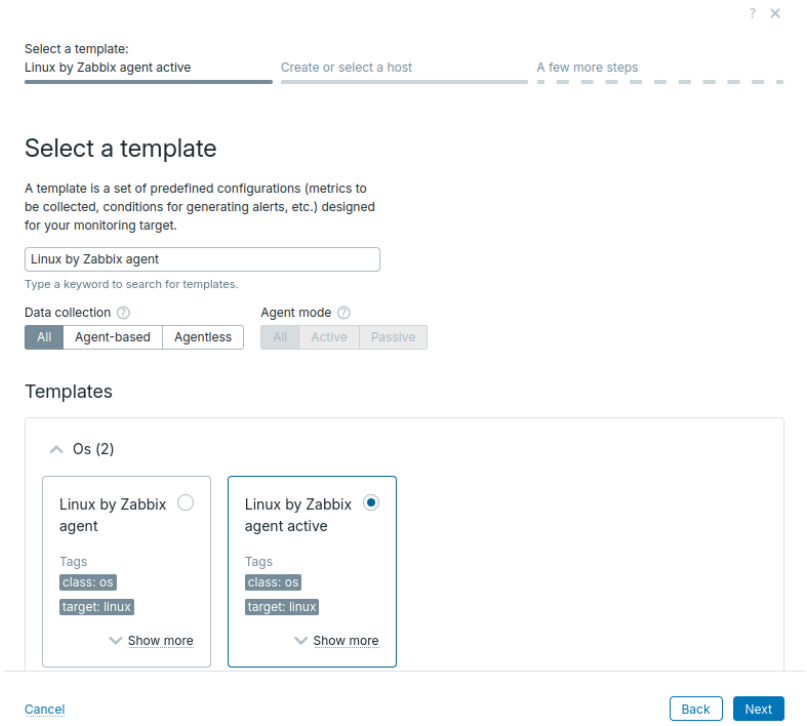
To use this guide you must have Zabbix server and the web interface installed. You must be logged in to Zabbix as the default *Admin* user.

Apply Linux monitoring template

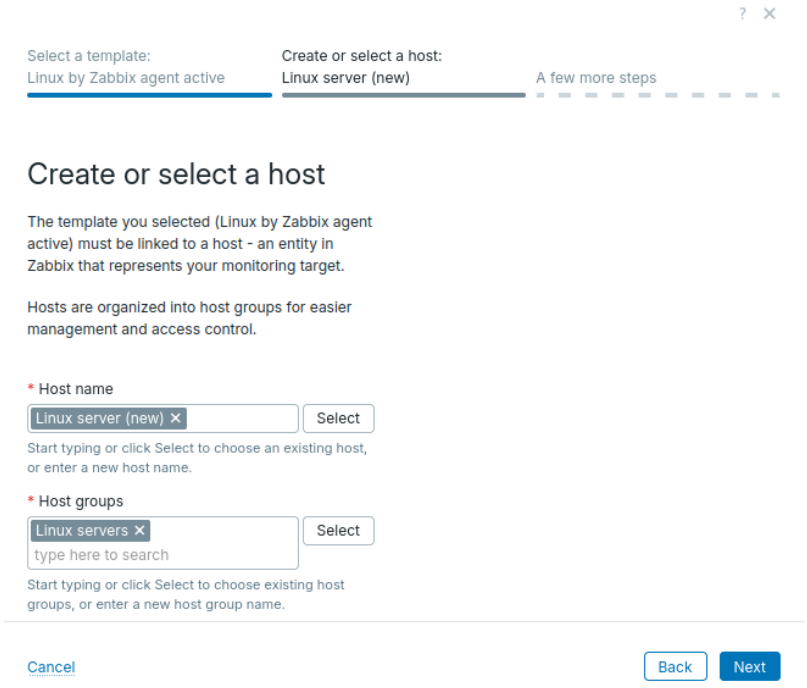
In Zabbix this template is ready to use. It contains pre-configured elements ready for data collection, visualisation and analysis.

To apply the template:

1. Start the Host Wizard in *Data collection* > *Hosts*.
2. Select the *Linux by Zabbix agent active* template.



3. Create a host for the Linux server, add a group for it. This virtual host will contain all data related to the Linux server monitoring.



4. Configure and install the agent.

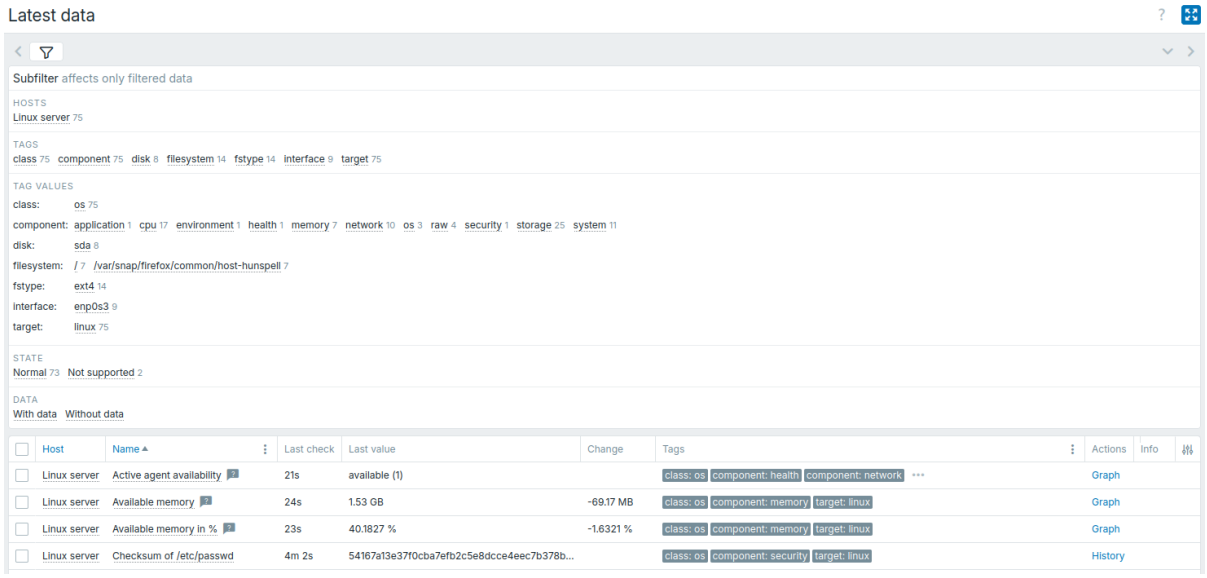


To install the agent, execute the command provided, on the Linux server. It will install the agent, and configure the agent's connection to Zabbix server with encryption.
 If you already have the agent installed, you can skip this step (click *Next*). However, you will be asked to provide the pre-shared key identity and the pre-shared key to connect to the agent.

You can leave other settings to their defaults.

5. Finish configuring the host: click *Create* and *Finish*.

Look at latest data



Congratulations, Linux server monitoring has been started!

Zabbix provides visualisation options and delivers alerts if any problems arise. To configure alerts in the email, see Receiving problem notification.

MySQL server monitoring

To start monitoring a MySQL server, follow this process:

1. Create limited MySQL user.
2. Apply MySQL monitoring template.
3. Look at latest data.

To start monitoring a MySQL server you need to know its DSN, and the IP address of the Zabbix server.

A Zabbix agent will be used to monitor the MySQL server. Instructions for installing the agent will be provided in the process.

To use this guide you must have Zabbix server and the web interface installed. You must be logged in to Zabbix as the default *Admin* user.

Create limited MySQL user

It is recommended to create a MySQL user, limited in permissions only to those needed for monitoring MySQL. Replace `<password>` with a strong, unique password.

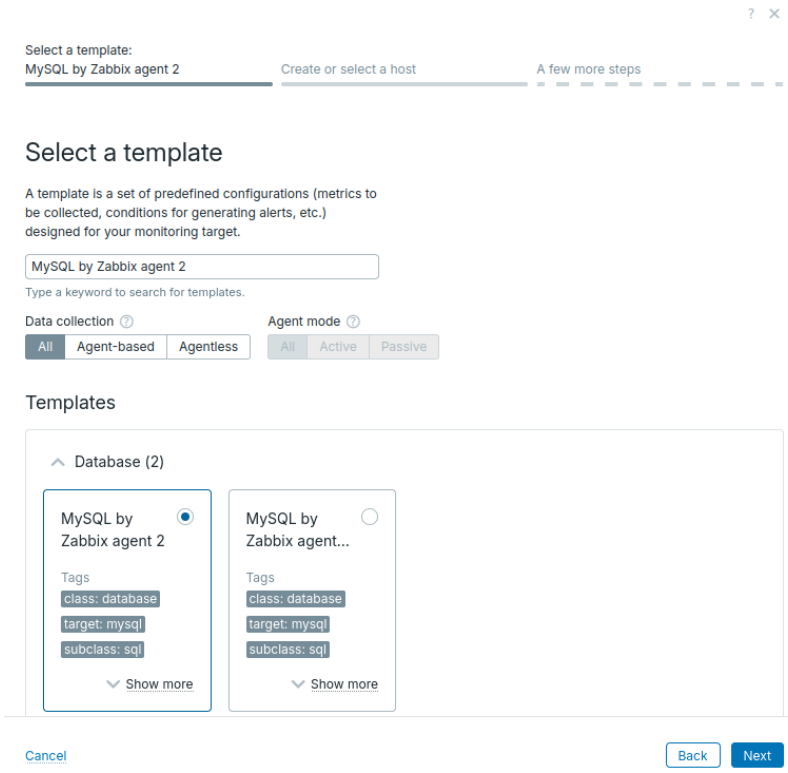
```
mysql> CREATE USER 'zbx_monitor'@'%' IDENTIFIED BY '<password>';
mysql> GRANT REPLICATION CLIENT,PROCESS,SHOW DATABASES,SHOW VIEW ON *.* TO 'zbx_monitor'@'%';
```

Apply MySQL monitoring template

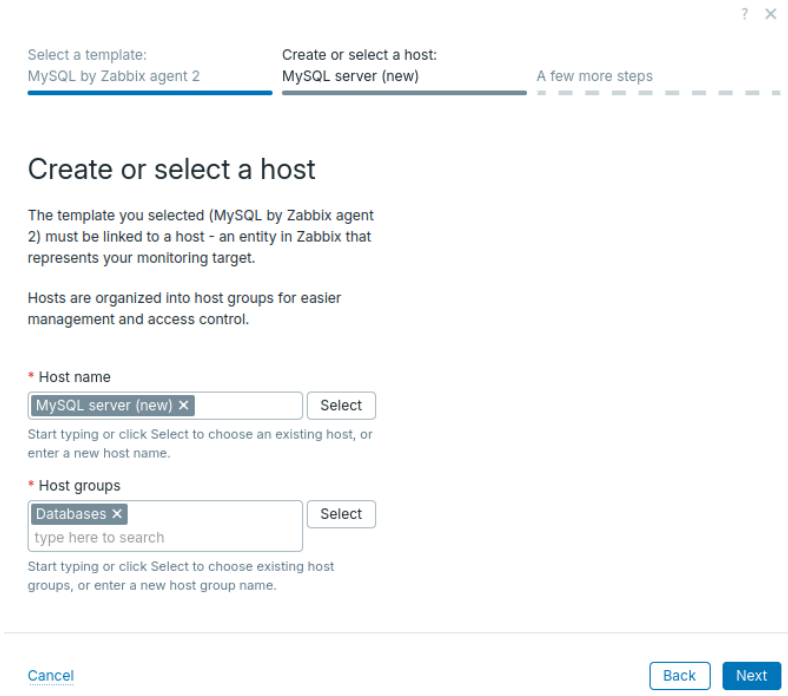
In Zabbix this template is ready to use. It contains pre-configured elements ready for data collection, visualisation and analysis.

To apply the template:

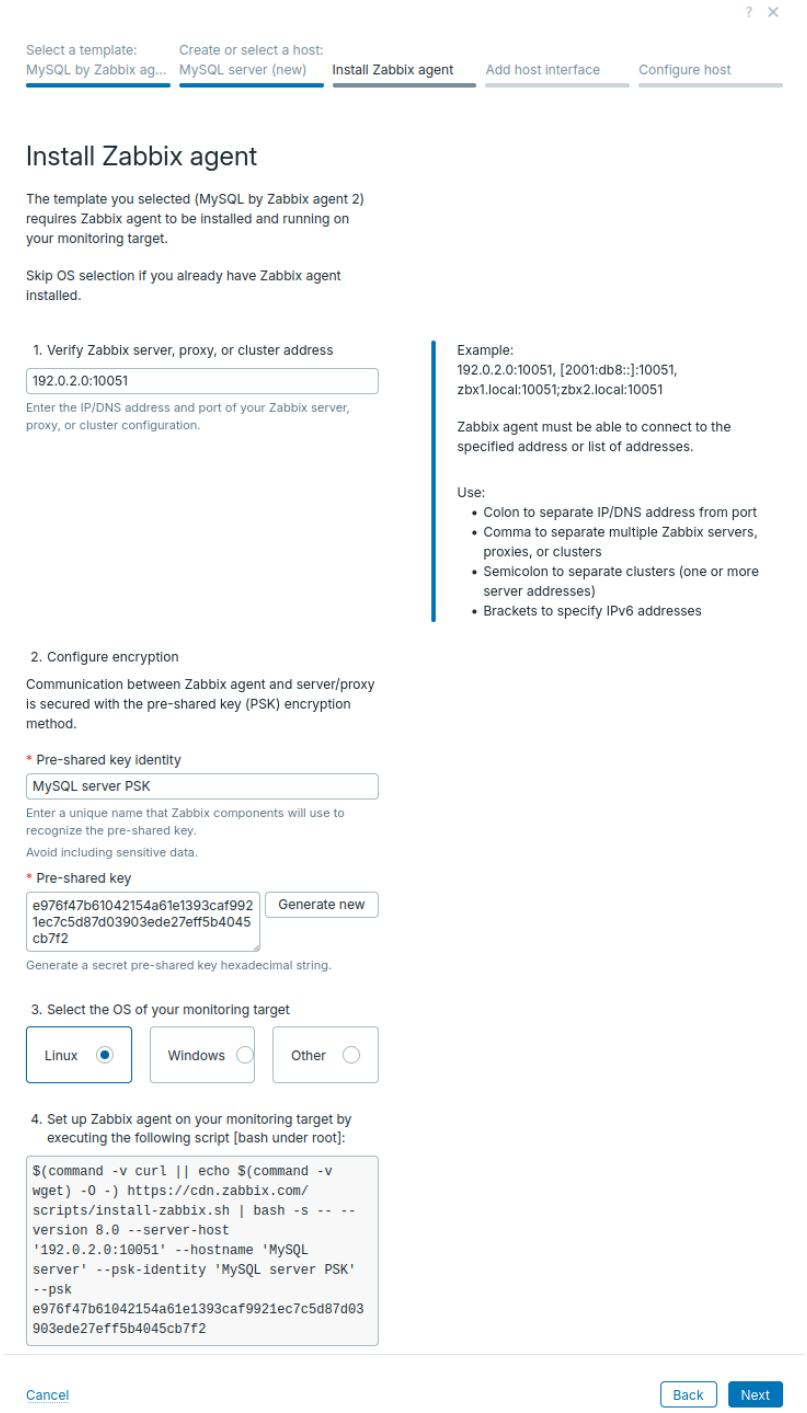
1. Start the Host Wizard in *Data collection* > *Hosts*.
2. Select the *MySQL by Zabbix agent 2* template.



3. Create a MySQL server host, add a group for it. This virtual host will contain all data related to the MySQL server monitoring.

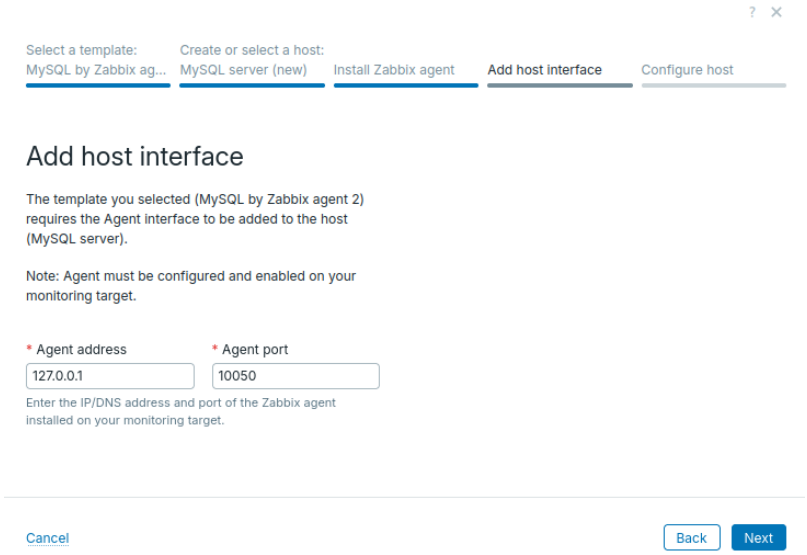


4. Configure and install agent 2 (Zabbix agent 2 provides MySQL monitoring out of the box).

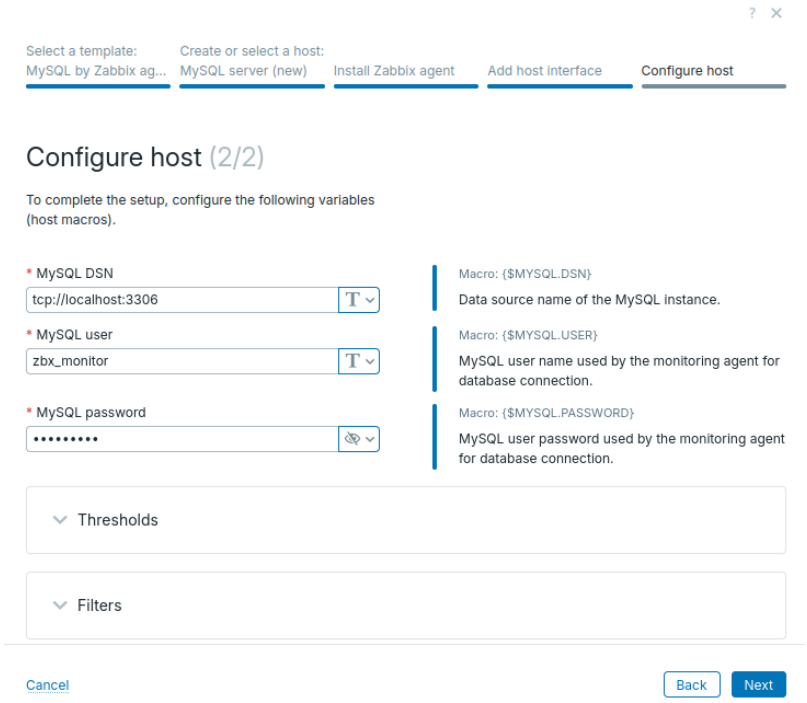


To install the agent, execute the command provided, on the machine hosting the MySQL server. It will install the agent, and configure the agent's connection to Zabbix server with encryption.
 If you already have agent 2 installed, you can skip this step (click *Next*). However, you will be asked to provide the pre-shared key identity and the pre-shared key to connect to the agent.

5. Enter the agent (Linux server) IP address and port.



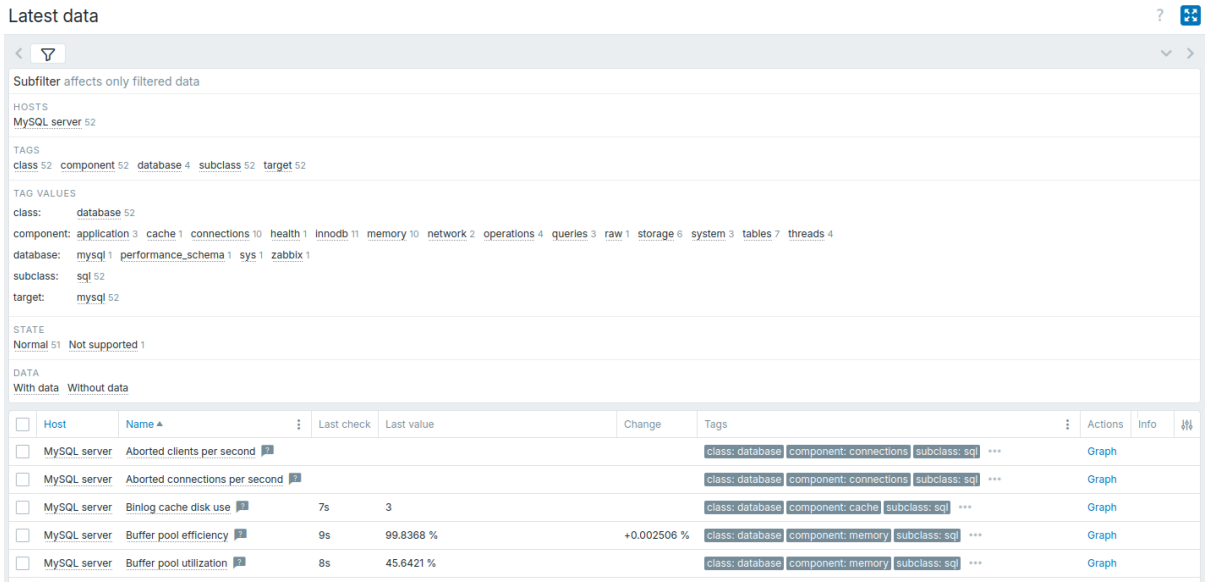
6. Finish configuring the host. Importantly, you will be asked to enter the MySQL server DSN, MySQL user and password. Provide the `zbx_monitor` user with the password that you created.



You can leave other settings to their defaults.

6. To complete the Host Wizard, click *Create* and *Finish*.

Look at latest data



Congratulations, MySQL server monitoring has been started!

Zabbix provides visualisation options and delivers alerts if any problems arise. To configure alerts in the email, see Receiving problem notification.

More guides

This documentation section contains quick recipes for setting up Zabbix for some commonly required monitoring goals.

It is designed with the new Zabbix user in mind and can be used as a navigator through other documentation sections that contain information required for resolving the task.

The following quick reference guides are available:

- Monitor Linux with Zabbix agent
- Monitor Windows with Zabbix agent
- Monitor Apache via HTTP
- Monitor MySQL with Zabbix agent 2
- Monitor VMware with Zabbix
- Monitor network traffic with Zabbix
- Monitor network traffic with Zabbix using active checks
- Monitor websites with Browser items
- Monitor website certificates with Zabbix agent 2 (passive)
- Monitor a network switch or router with Zabbix
- Monitor Windows event log using active checks

Monitor Linux with Zabbix agent

Introduction

This page walks you through the steps required to start basic monitoring of Linux machines with Zabbix. The steps described in this tutorial can be applied to any Linux-based operating system.

Who this guide is for

This guide is designed for new Zabbix users and contains the minimum set of steps required to enable basic monitoring of your Linux machine. If you are looking for deep customization options or require more advanced configuration, see Configuration section of Zabbix manual.

Prerequisites

Before proceeding with this guide, you need to download and install Zabbix server and Zabbix frontend according to the instructions for your OS.

Install Zabbix agent

Zabbix agent is the process responsible for gathering data.

Check your Zabbix server version:

```
zabbix_server -V
```

Install Zabbix agent of the same version (recommended) on the Linux machine that you want to monitor. Based on your monitoring needs, it may be the same machine, where Zabbix server is installed, or a completely different machine.

Choose the most suitable installation method:

- Run as a Docker container - see the list of available images in Zabbix Docker repository.
- Install from Zabbix packages (available for Alma Linux, CentOS, Debian, Oracle Linux, Raspberry Pi OS, RHEL, Rocky Linux, SUSE Linux Enterprise Server, Ubuntu).
- Compile from sources.

Configure Zabbix for monitoring

Zabbix agent can collect metrics in active or passive mode (simultaneously).

> **Note:**
> A passive check is a simple data request. Zabbix server or proxy asks for some data (for example, CPU load) and Zabbix agent sends back the result to the server. Active checks require more complex processing. The agent must first retrieve from the server(s) a list of items for independent processing and then bulk send the data back. See Passive and active agent checks for more info.

Monitoring templates provided by Zabbix usually offer two alternatives - a template for Zabbix agent and a template for Zabbix agent (active). With the first option, the agent will collect metrics in passive mode. Such templates will deliver identical monitoring results, but using different communication protocols.

Further Zabbix configuration depends on whether you select a template for active or passive Zabbix agent checks.

Passive checks

Zabbix agent

1. Open the agent configuration file on the machine where the agent is installed.

```
sudo vi /etc/zabbix/zabbix_agentd.conf
```

2. Add the IP address or DNS name (and optional port) of your Zabbix server to the server parameter. For example:

```
Server=192.0.2.1:10051
```

Zabbix agent will use this address to accept incoming connections only from the specified Zabbix servers or proxies to retrieve data.

3. Restart Zabbix agent.

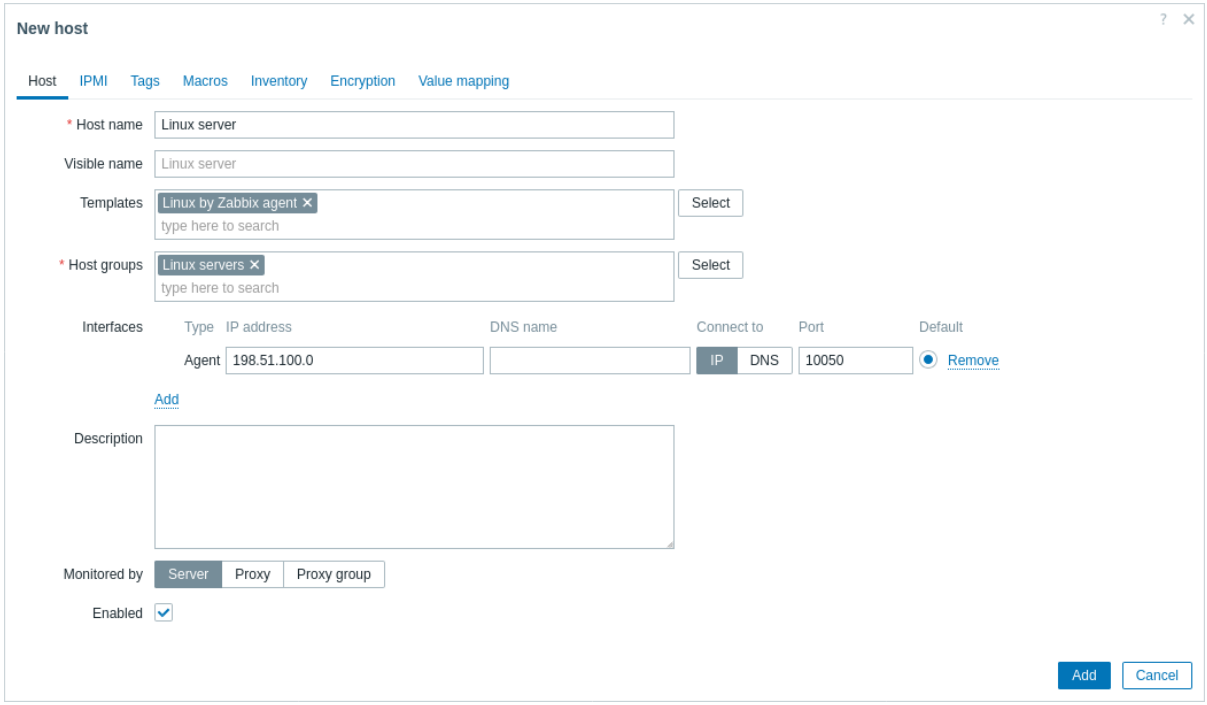
```
systemctl restart zabbix-agent
```

Zabbix frontend

1. Log in to Zabbix frontend.

2. Create a host in Zabbix web interface.

- In the *Host name* field, enter a host name (e.g., ”Linux server”).
- In the *Templates* field, type or select the template ”Linux by Zabbix agent” that will be linked to the host.
- In the *Host groups* field, type or select a host group (e.g., ”Linux servers”).
- In the *Interfaces* parameter, add *Agent* interface and specify the IP address or DNS name of the Linux machine where the agent is installed.



3. Click *Add* to add the host. This host will represent the monitored Linux machine.

Active checks

Zabbix agent

1. Open the agent configuration file on the machine where the agent is installed:

```
sudo vi /etc/zabbix/zabbix_agentd.conf
```

2. Add the IP address or DNS name (and optional port) of your Zabbix server to the ServerActive parameter. For example:

```
ServerActive=192.0.2.1:10051
```

Zabbix agent will use this address to connect to Zabbix server's trapper port (default: 10051) and request the configuration data for active checks.

3. Define the Hostname parameter, which must match the host name that will be defined in Zabbix frontend. In our example, it is:

```
Hostname=Linux server
```

The Hostname value must match because, for active checks, Zabbix agent uses it to retrieve the correct host configuration from the server. Specifically, the agent initiates a connection to the server and identifies itself using the Hostname value. The server then provides the monitoring configuration for that host. If these values differ, the agent will not receive the appropriate configuration, resulting in missing metrics or monitoring issues.

4. Restart Zabbix agent.

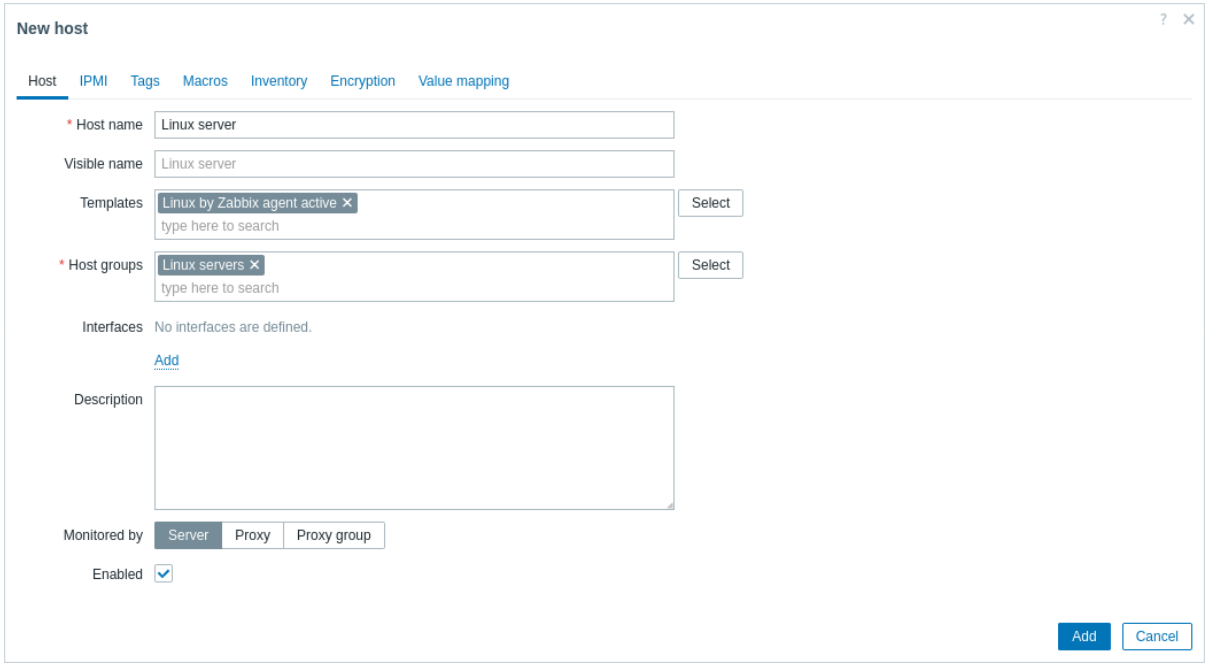
```
systemctl restart zabbix-agent
```

Zabbix frontend

1. Log in to Zabbix frontend.

2. Create a host in Zabbix web interface.

- In the *Host name* field, enter a host name (e.g., "Linux server") that matches the *Hostname* parameter value defined earlier in the agent configuration file.
- In the *Templates* field, type or select the template "Linux by Zabbix agent active" that will be linked to the host.
- In the *Host groups* field, type or select a host group (e.g., "Linux servers").
- Leave *Interfaces* undefined; an interface is not necessary for active checks because the agent initiates the connection to the server instead of listening for a connection from the server.

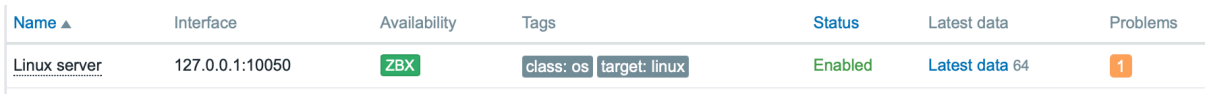


3. Click *Add* to add the host. This host will represent the monitored Linux machine.

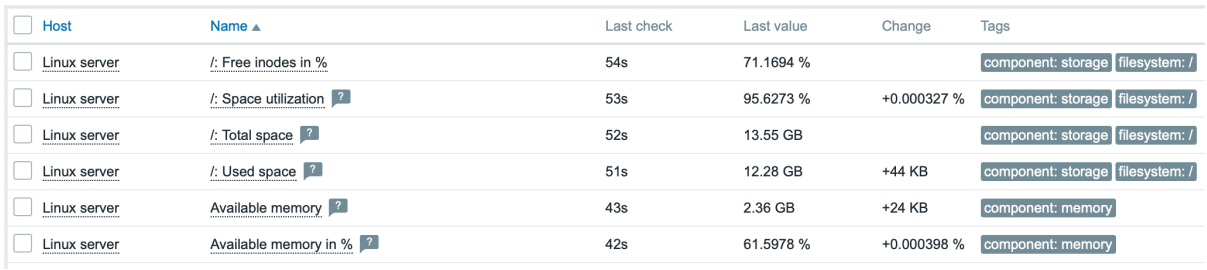
View collected metrics

Congratulations! At this point, Zabbix is already monitoring your Linux machine.

To view collected metrics, open the *Monitoring* > *Hosts* menu section and click the *Latest data* next to the host.



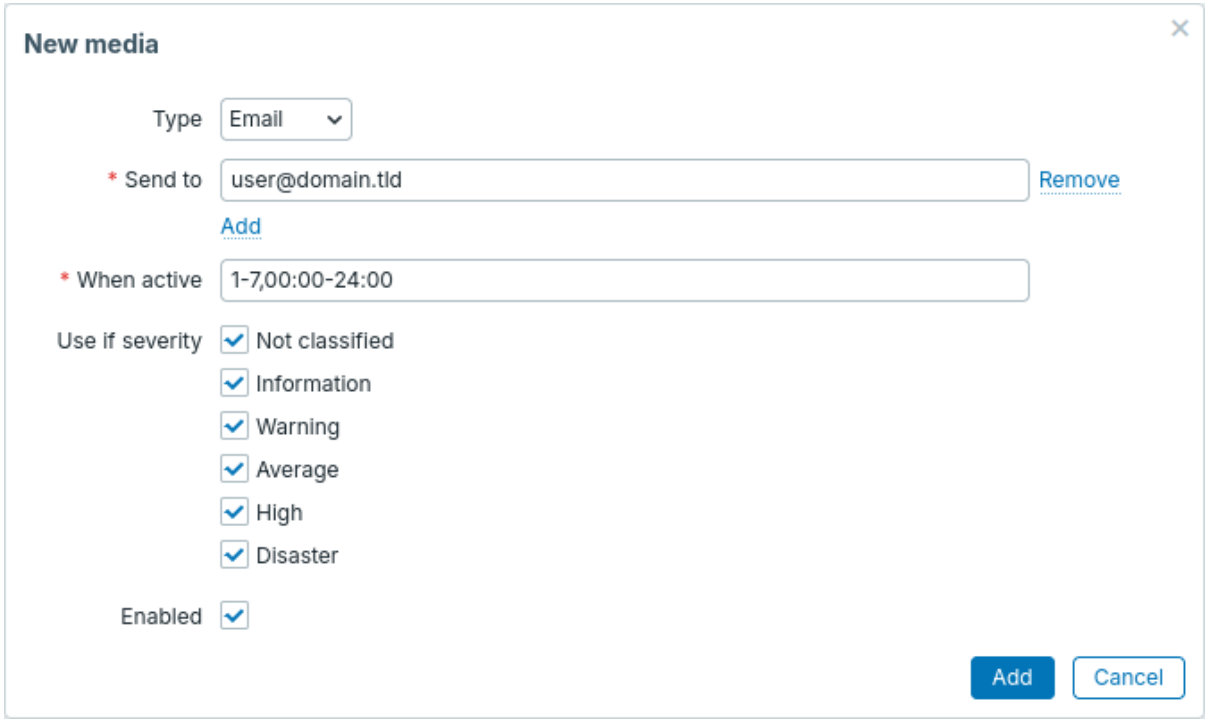
This action will open a list of all the latest metrics collected from Linux server host.



Set up problem alerts

Zabbix can notify you about a problem with your infrastructure using a variety of methods. This guide provides configuration steps for sending email alerts.

1. Go to the *User settings* > *Profile*, switch to the tab *Media* and add your email.



2. Follow the guide for Receiving problem notification.

Next time, when Zabbix detects a problem you should receive an alert via email.

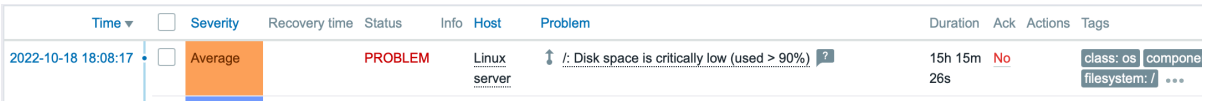
Test your configuration

On Linux, you can simulate high CPU load and as a result receive a problem alert by running:

```
cat /dev/urandom | md5sum
```

You may need to run several md5sum processes for CPU load to exceed the threshold.

When Zabbix detects the problem, it will appear in the *Monitoring* > *Problems* section.



If the alerts are configured, you will also receive the problem notification.

See also:

- Creating an item - how to start monitoring additional metrics (custom monitoring without templates).
- Zabbix agent items, Zabbix agent items for Windows - full list of metrics you can monitor using Zabbix agent on Windows.
- Problem escalations - how to create multi-step alert scenarios (e.g., first send message to the system administrator, then, if a problem is not resolved in 45 minutes, send message to the data center manager).
- Installation from packages - how to install Zabbix components using official RPM and DEB packages for various Linux distributions, ensuring access to the latest features and bug fixes.

Monitor Windows with Zabbix agent

Introduction

This page walks you through the steps required to start basic monitoring of Windows machines with Zabbix.

Who this guide is for

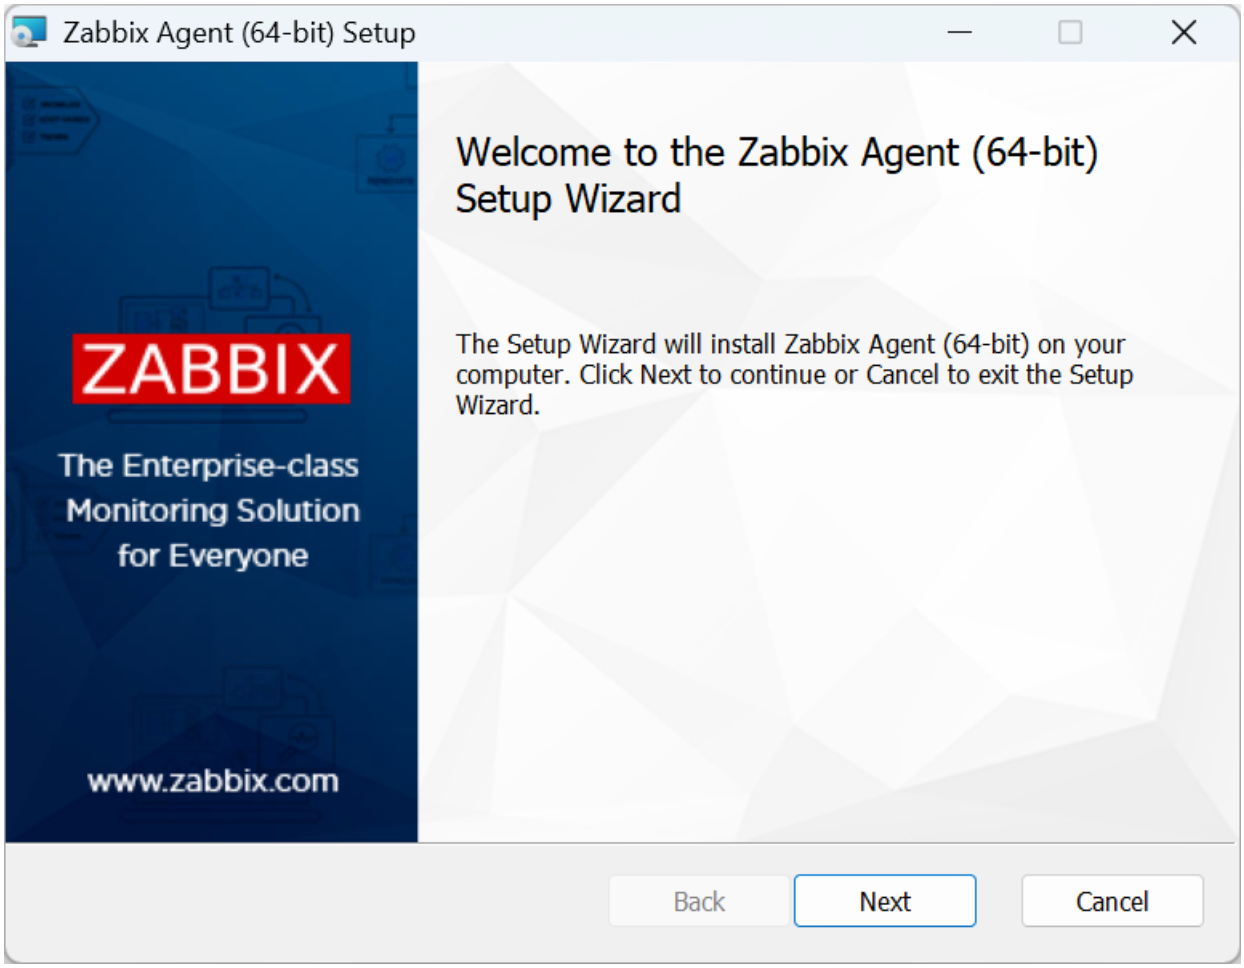
This guide is designed for new Zabbix users and contains the minimum set of steps required to enable basic monitoring of your Windows machine. If you are looking for deep customization options or require more advanced configuration, see Configuration section of Zabbix manual.

Prerequisites

Before proceeding with this installation guide, you must download and install Zabbix server and Zabbix frontend according to instructions for your OS.

Install Zabbix agent

Zabbix agent is the process responsible for gathering data. You need to install it on the Windows machine that you want to monitor. Follow Zabbix agent installation instructions for Windows.



Configure Zabbix for monitoring

Zabbix agent can collect metrics in active or passive mode (simultaneously).

Note:
A passive check is a simple data request. Zabbix server or proxy asks for some data (for example, CPU load) and Zabbix agent sends back the result to the server. Active checks require more complex processing. The agent must first retrieve from the server(s) a list of items for independent processing and then bulk send the data back. See Passive and active agent checks for more info.

Monitoring templates provided by Zabbix usually offer two alternatives - a template for Zabbix agent and a template for Zabbix agent (active). With the first option, the agent will collect metrics in passive mode. Such templates will deliver identical monitoring results, but using different communication protocols.

Further Zabbix configuration depends on whether you select a template for active or passive Zabbix agent checks.

Passive checks

Zabbix frontend

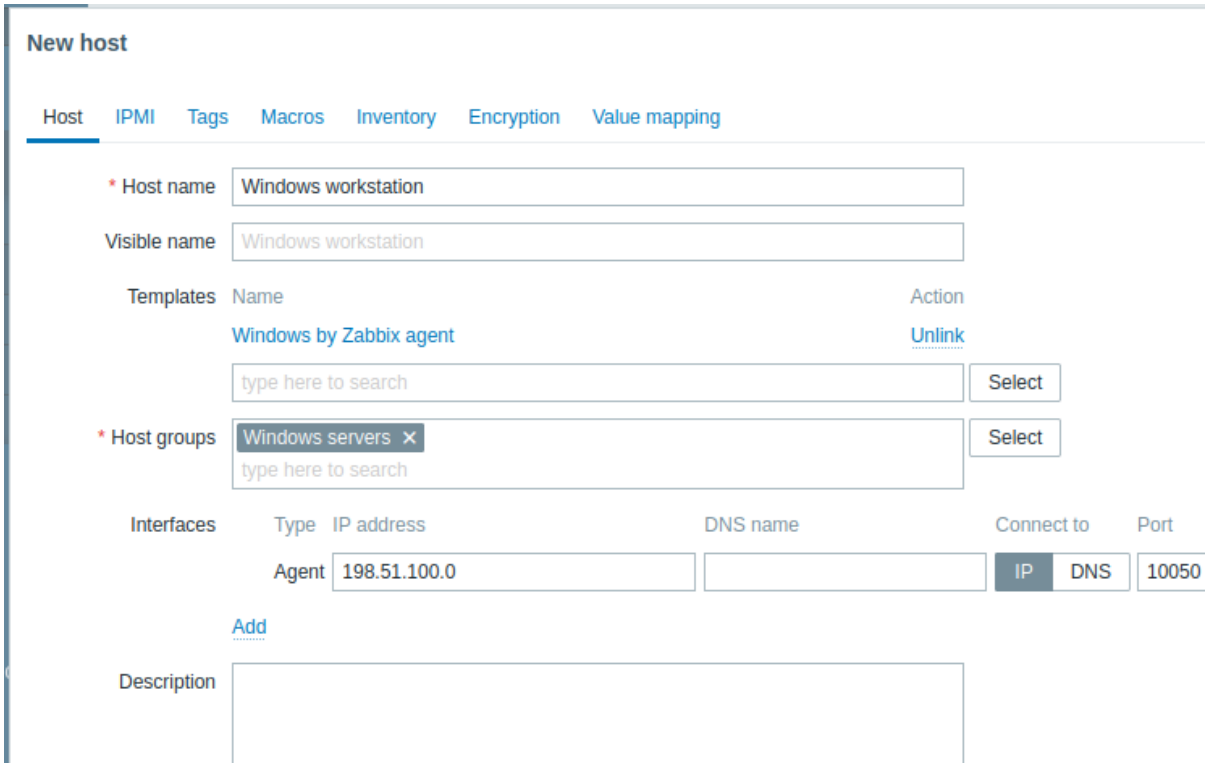
1. Log into Zabbix frontend.

2. Create a host in Zabbix web interface.

This host will represent your Windows machine.

3. In the *Interfaces* parameter, add *Agent* interface and specify the IP address or DNS name of the Windows machine where the agent is installed.

4. In the *Templates* parameter, type or select *Windows by Zabbix agent*.



Zabbix agent

For passive checks the Zabbix agent needs to know the IP address or DNS name of the Zabbix server. If you have provided the correct information during the agent installation, the configuration file has already been updated. Otherwise, you need to specify it manually. Open the agent configuration file, and add the IP/DNS of your Zabbix server to the `Server` parameter.

Example:

```
Server=192.0.2.22
```

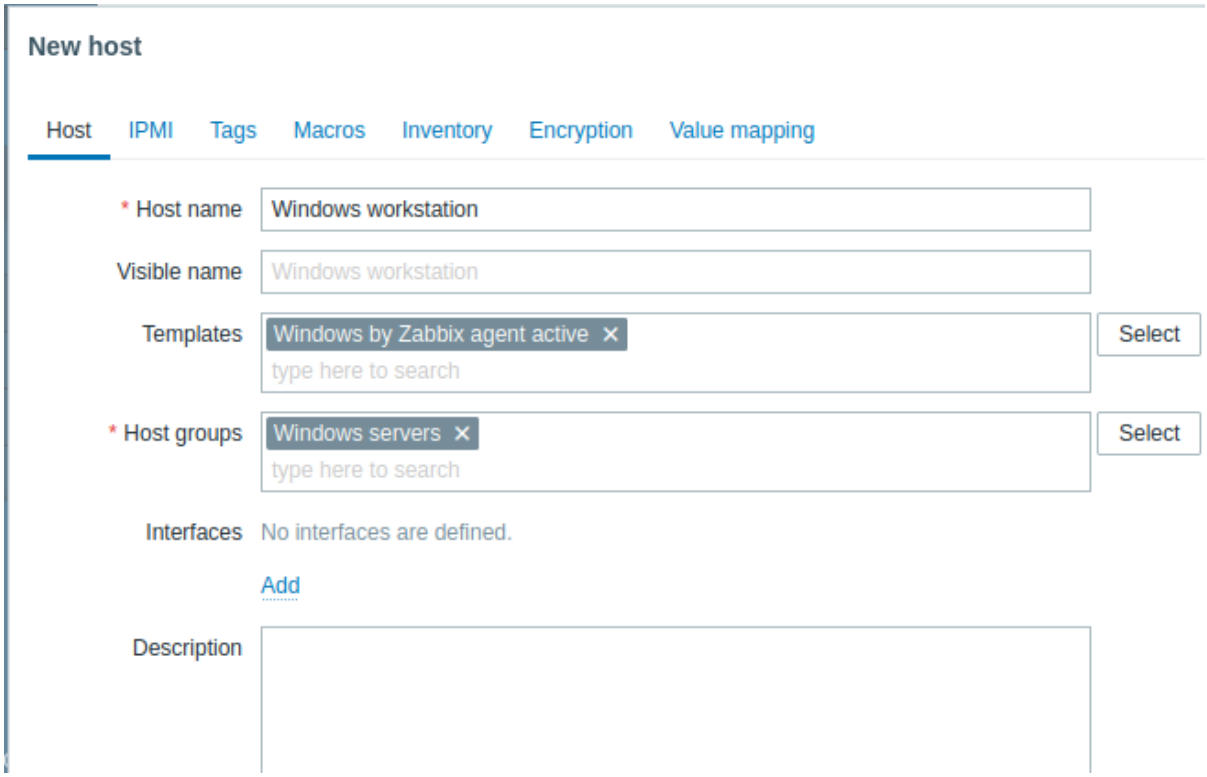
Active checks

Zabbix frontend

1. Log in to Zabbix frontend.
2. Create a host in Zabbix web interface.

This host will represent your Windows machine.

3. In the *Templates* parameter, type or select *Windows by Zabbix agent active*.



Zabbix agent

In the `C:\Program files\Zabbix Agent` folder open the file `zabbix_agentd.conf` and add:

- The name of the host you created in Zabbix web interface to the *Hostname* parameter.
- The IP address or DNS name of your Zabbix server machine to the *ServerActive* parameter (might be prefilled if you have provided it during Zabbix agent setup).

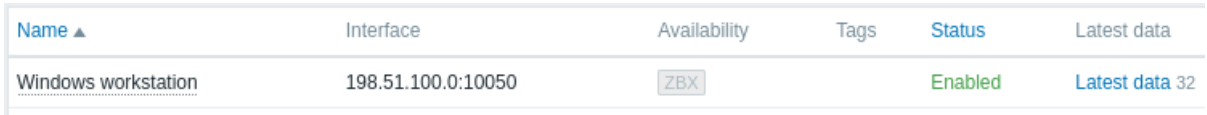
Example:

```
ServerActive= 192.0.2.22
Hostname=Windows workstation
```

View collected metrics

Congratulations! At this point, Zabbix is already monitoring your Windows machine.

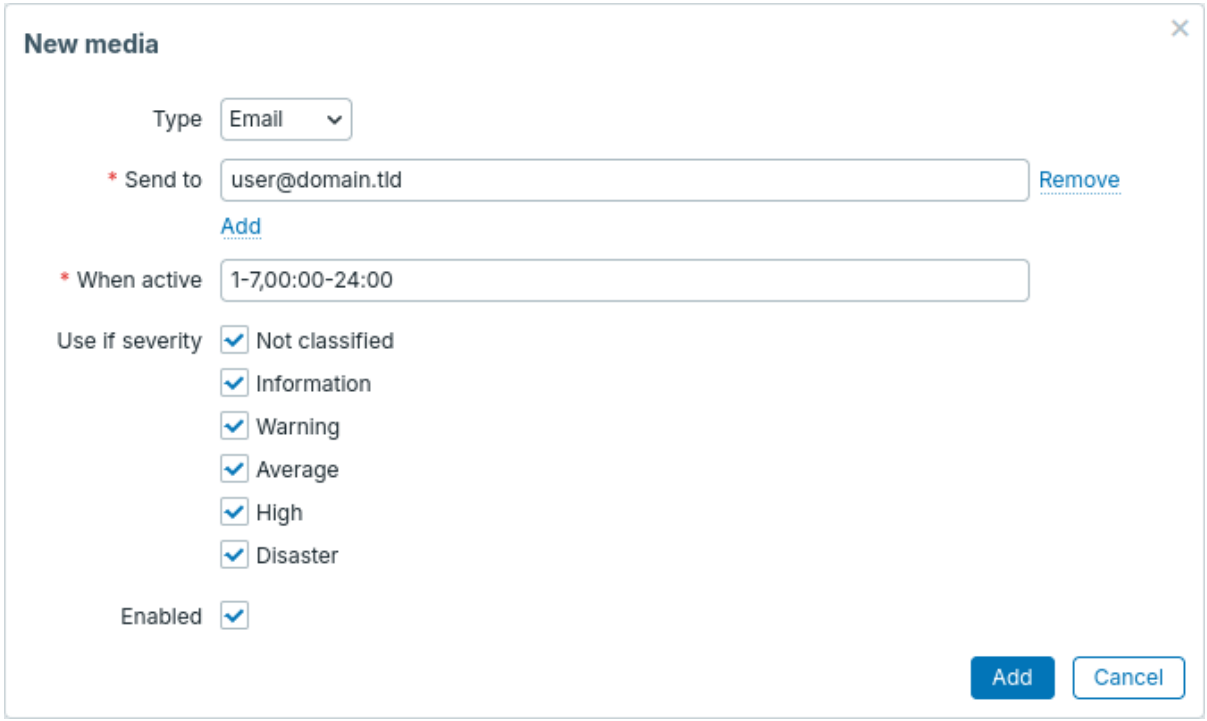
To view collected metrics, open the *Monitoring* > *Hosts* menu section and click the *Latest data* next to the host.



Set up problem alerts

Zabbix can notify you about a problem with your infrastructure using a variety of methods. This guide provides configuration steps for sending email alerts.

1. Go to the *User settings* > *Profile*, switch to the tab *Media* and add your email.



2. Follow the guide for Receiving problem notification.

Next time, when Zabbix detects a problem you should receive an alert via email.

Note:
On Windows, you can use CpuStres utility to simulate high CPU load and as a result receive a problem alert.

See also:

- Creating an item - how to start monitoring additional metrics (custom monitoring without templates).
- Zabbix agent items, Zabbix agent items for Windows - full list of metrics you can monitor using Zabbix agent on Windows.
- Problem escalations - how to create multi-step alert scenarios (e.g., first send message to the system administrator, then, if a problem is not resolved in 45 minutes, send message to the data center manager).

Monitor Apache via HTTP

Introduction

This page shows a quick and simple way to start monitoring an Apache web server without installing any additional software.

Who this guide is for

This guide is designed for new Zabbix users and contains the minimum set of steps required to enable basic monitoring of your Apache installation. If you are looking for deep customization options or require more advanced configuration, see Configuration section of Zabbix manual.

Prerequisites

Before proceeding with this installation guide, you must download and install Zabbix server and Zabbix frontend according to instructions for your OS.

Prepare Apache

1. Check, which Apache version you are using:

On RHEL-based system, run:

```
httpd -v
```

On Debian/Ubuntu, run:

```
apache2 -v
```

2. Make sure that the Status module is enabled in your Apache instance.

On RHEL-based system, run:

```
httpd -M | grep status
status_module (shared)
```

On Debian/Ubuntu, run:

```
apache2ctl -M | grep status
status_module (shared)
```

If you don't see status_module in the list, enable the module by running:

On RHEL-based system, run:

```
LoadModule status_module /usr/lib/apache2/modules/mod_status.so
```

On Debian/Ubuntu, run:

```
sudo /usr/sbin/a2enmod status
```

3. Edit Apache configuration file to allow access to status reports from Zabbix server IP.

On an RHEL-based system: `/etc/httpd/conf.modules.d/status.conf`:

```
sudo vi /etc/httpd/conf.modules.d/status.conf
```

On Debian/Ubuntu: `/etc/apache2/mods-enabled/status.conf`:

```
sudo vi /etc/apache2/mods-enabled/status.conf
```

Add the following lines to the file (**replace 198.51.100.255** with your Zabbix server IP address):

- For Apache 2.2:
  <Location /server-status> SetHandler server-status

  Order Deny,Allow Deny from all Allow from 198.51.100.255 </Location>
- For Apache 2.4:
  <Location "/server-status"> SetHandler server-status Require ip 198.51.100.255 </Location>

4. Restart Apache

On an RHEL-based system, run:

```
sudo systemctl restart httpd
```

On Debian/Ubuntu, run:

```
sudo systemctl restart apache2
```

5. To check, if everything is configured correctly, run (**replace 198.51.100.255** with your Zabbix server IP address):

```
curl 198.51.100.255/server-status
```

The response should contain Apache web server statistics.

Configure Zabbix for monitoring

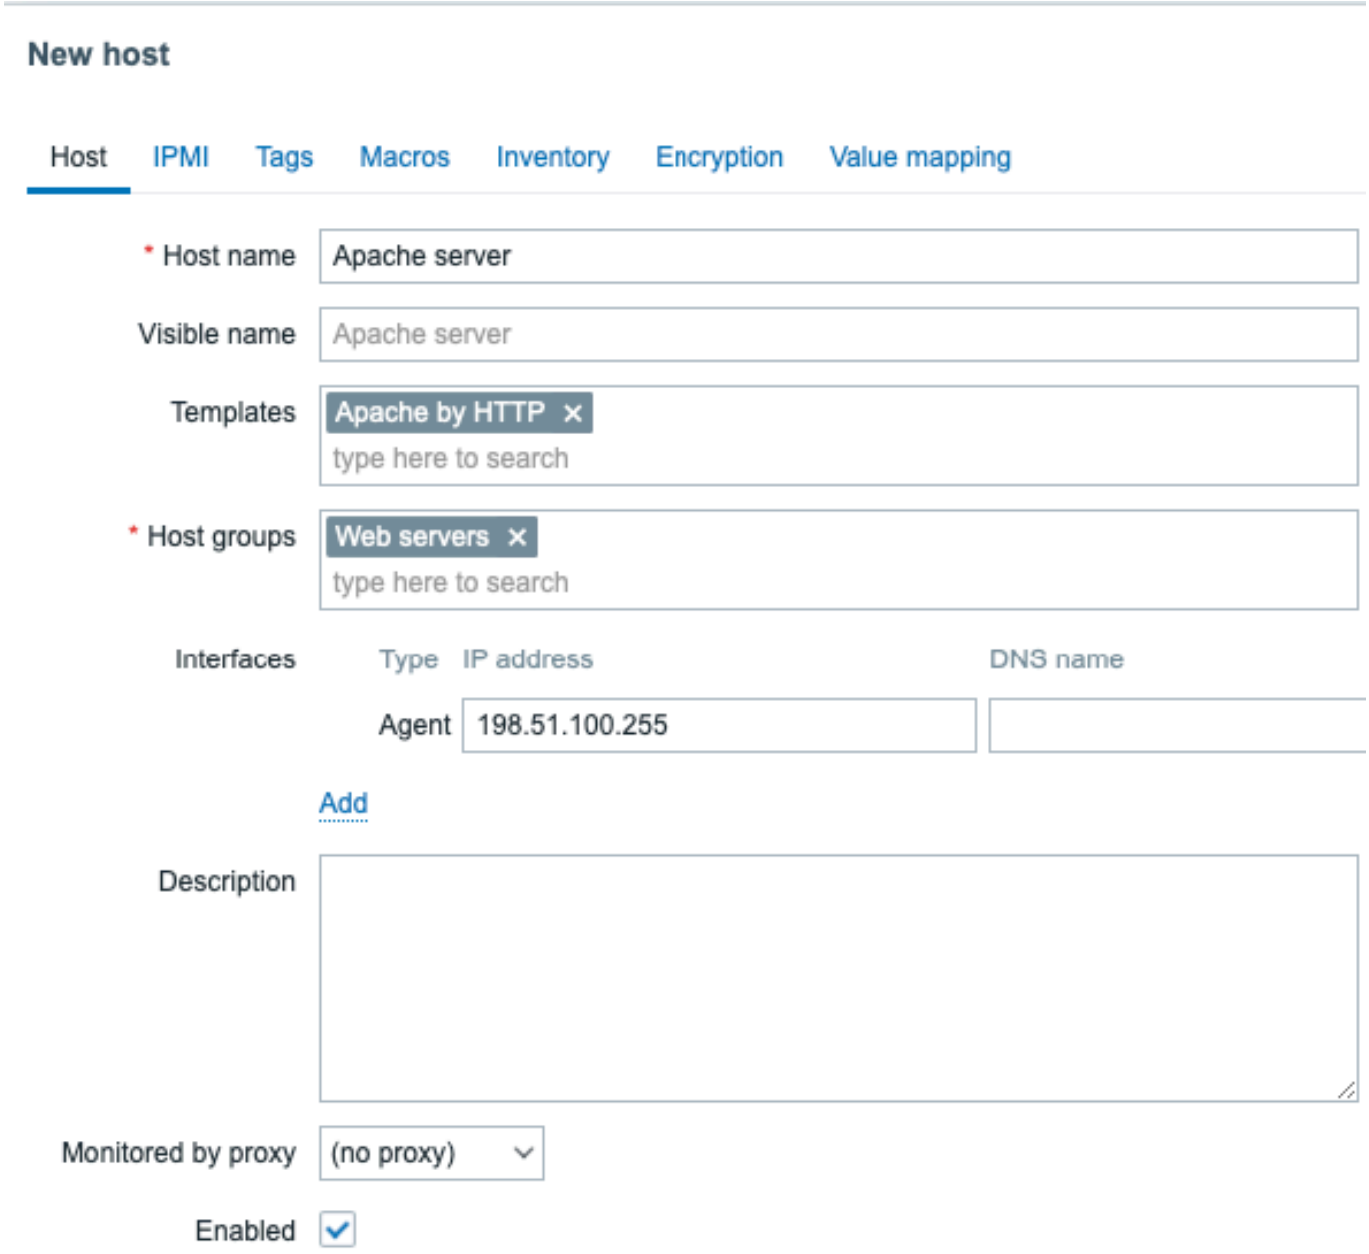
1. Log into Zabbix frontend.

2. Create a host in Zabbix web interface.

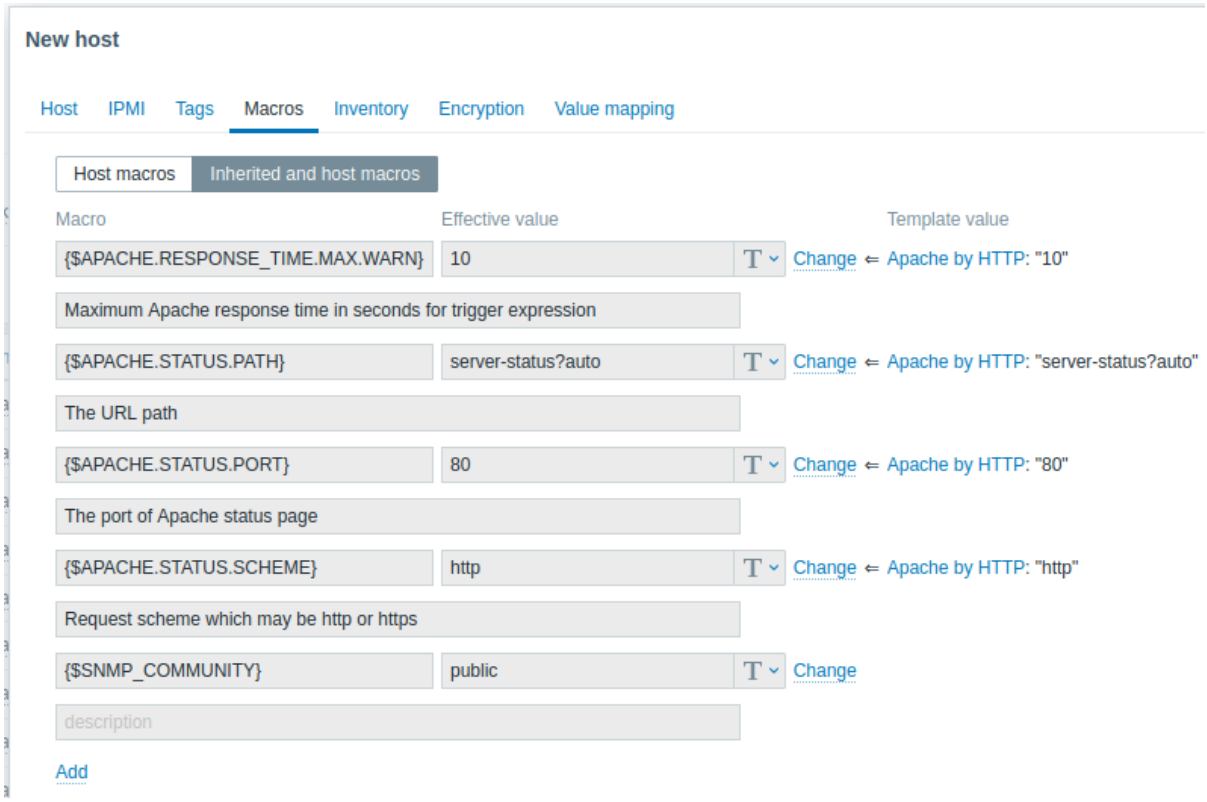
This host will represent your Apache server.

3. In the *Interfaces* parameter, add *Agent* interface and specify your Apache instance IP address. **You don't need to install Zabbix agent on the machine**, the interface will only be used for resolving {HOST.CONN} macro. This macro is used in template items to locate Apache instance.

4. In the *Templates* parameter, type or select *Apache by HTTP*.



5. Switch to the **Macros** tab and select *Inherited and host macros* mode. Check that values of the macros {$APACHE.STATUS.PORT} and {$APACHE.STATUS.SCHEME} suit your installation settings. By default, the port is 80 and the scheme is http. Change macro values if you use different port and/or scheme.



View collected metrics

Congratulations! At this point, Zabbix is already monitoring your Apache web server.

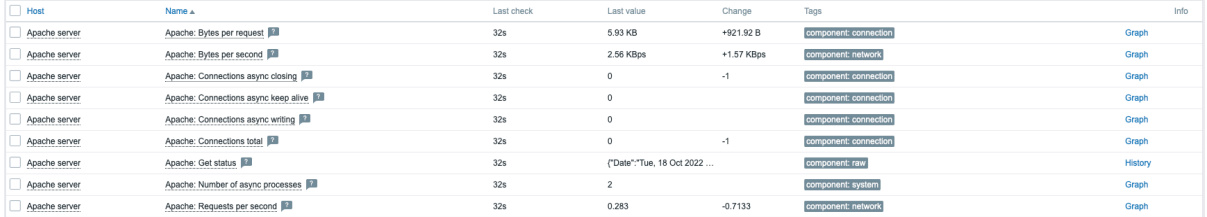
To view collected metrics, open the *Monitoring* > *Hosts* menu section and click the *Dashboards* next to the host.



This action will take you to the host dashboard with most important metrics collected from Apache /server-status page.



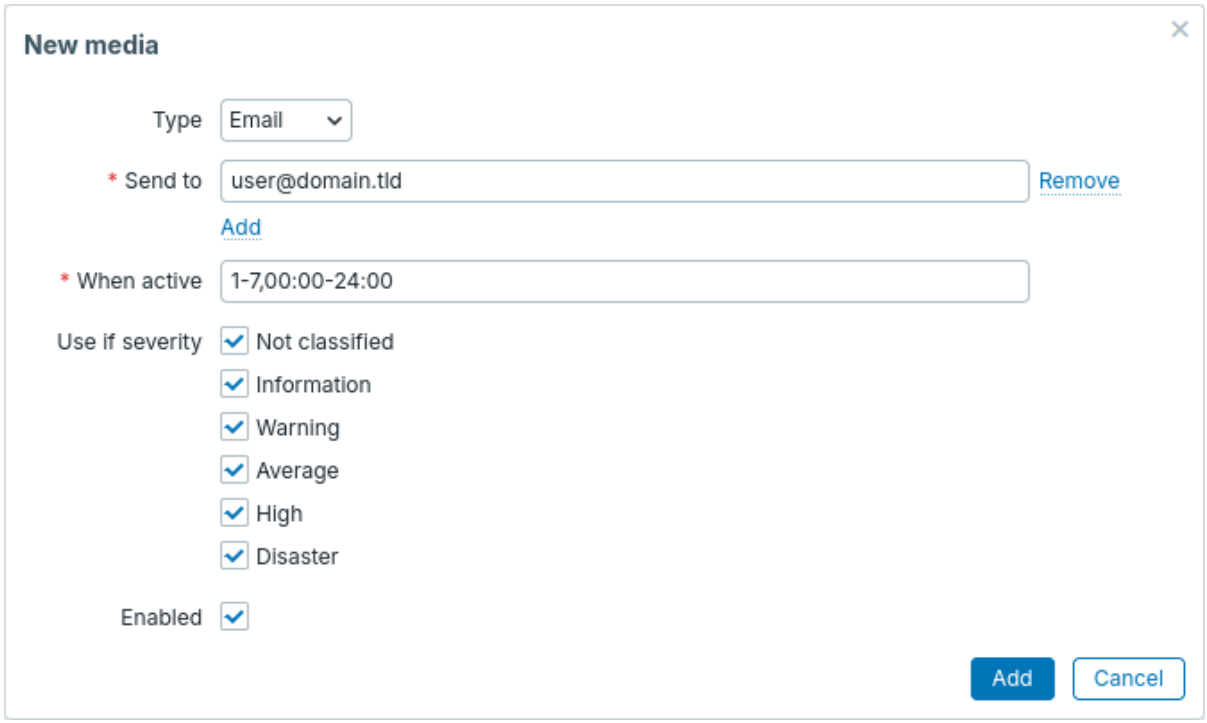
Alternatively, from the *Monitoring* > *Hosts*, you can click the *Latest data* to view all the latest collected metrics in a list.



Set up problem alerts

Zabbix can notify you about a problem with your infrastructure using a variety of methods. This guide provides configuration steps for sending email alerts.

1. Go to the *User settings* > *Profile*, switch to the tab *Media* and add your email.



2. Follow the guide for Receiving problem notification.

Next time, when Zabbix detects a problem you should receive an alert via email.

Test your configuration

To simulate real problem and receive a test problem alert:

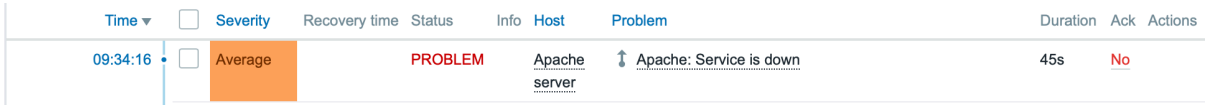
1. Open the *Apache server* host configuration in Zabbix.

2. Switch to the Macros tab and select *Inherited and host macros*.

3. Click *Change* next to {$APACHE.STATUS.PORT} macro and set a different port.

4. Click *Update* to save host configuration.

5. In a few minutes, Zabbix will detect the problem *Apache service is down*, because now it cannot connect to the instance. It will appear in the *Monitoring* > *Problems* section.



If the alerts are configured, you will also receive the problem notification.

6. Change the macro value back to resolve the problem and continue monitoring Apache.

See also:

- Web server hardening - recommended settings for greater web server security.
- Creating an item - how to start monitoring additional metrics.
- HTTP items - how to monitor custom metrics using HTTP agent.
- Problem escalations - how to create multi-step alert scenarios (e.g., first send message to the system administrator, then, if a problem is not resolved in 45 minutes, send message to the data center manager).

Monitor MySQL with Zabbix agent 2

Introduction

This page walks you through the steps required to start basic monitoring of a MySQL server.

To monitor a MySQL server, there are several approaches: Zabbix agent, Zabbix agent 2, or the Open Database Connectivity (ODBC) standard. The primary focus of this guide is on monitoring a MySQL server with Zabbix agent 2, which is the **recommended** approach due to its seamless configuration across various setups. However, this page also offers instructions for the other approaches, so feel free to choose the one that best suits your requirements.

Who this guide is for

This guide is designed for new Zabbix users and contains the minimum set of steps required to enable basic monitoring of a MySQL server. If you are looking for deep customization options or require more advanced configuration, see the Configuration section of Zabbix manual.

Prerequisites

Before proceeding with this guide, you need to download and install Zabbix server, Zabbix frontend and Zabbix agent 2 according to the instructions for your OS.

Depending on your environment, some steps in this guide may slightly differ. This guide is based on an environment running Ubuntu.

Create MySQL user

To monitor a MySQL server, Zabbix requires access to it and its processes. Your MySQL installation already has a user with the required level of access (the user "zabbix" that was created when installing Zabbix), however, this user has more privileges than necessary for simple monitoring (privileges to DROP databases, DELETE entries from tables, etc.). Therefore, a MySQL user for the purpose of *only* monitoring the MySQL server needs to be created.

1. Connect to the MySQL client, create a "zbx_monitor" user (replace <*password*> for the "zbx_monitor" user with a password of your choice), and GRANT the necessary privileges to the user:

```
mysql -u root -p
# Enter password:

mysql> CREATE USER 'zbx_monitor'@'%' IDENTIFIED BY '<password>';
mysql> GRANT REPLICATION CLIENT,PROCESS,SHOW DATABASES,SHOW VIEW ON *.* TO 'zbx_monitor'@'%';
mysql> quit;
```

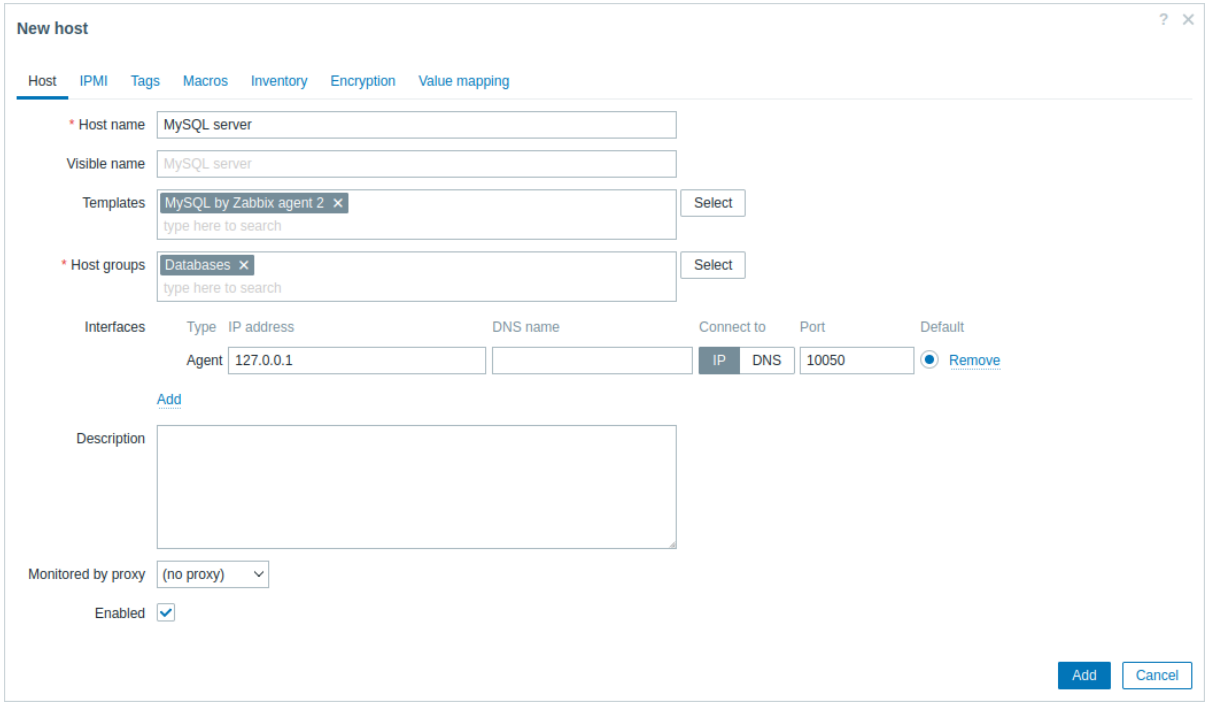
Once the user is created, you can move on to the next step.

Configure Zabbix frontend

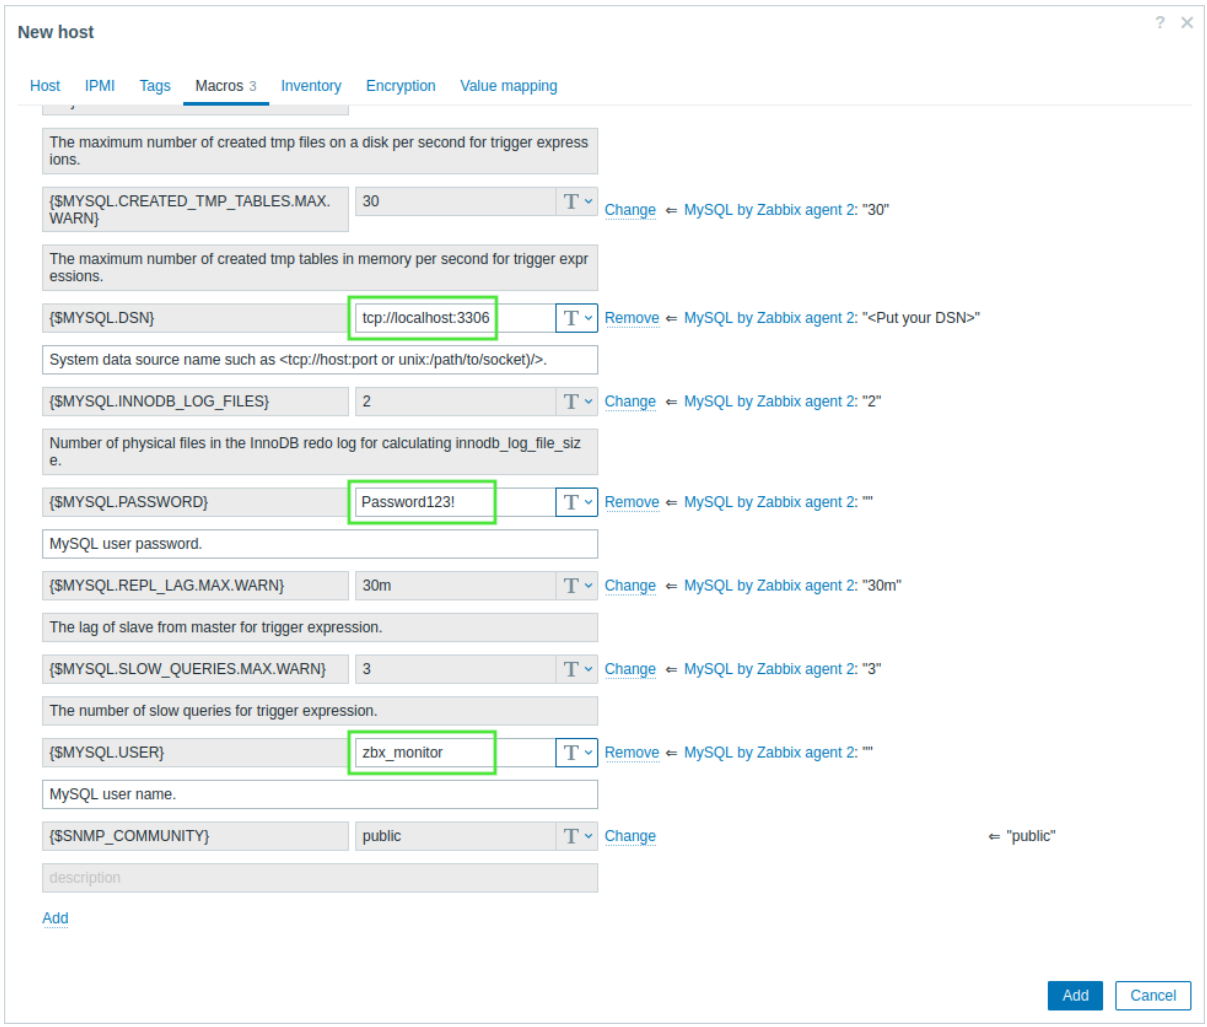
1. Log into Zabbix frontend.

2. Create a host in Zabbix web interface:

- In the *Host name* field, enter a host name (e.g., "MySQL server").
- In the *Templates* field, type or select the template "MySQL by Zabbix agent 2" that will be linked to the host.
- In the *Host groups* field, type or select a host group (e.g., "Databases").
- In the *Interfaces* field, add an interface of type "Agent" and specify your MySQL server IP address. This guide uses "127.0.0.1" (localhost) for monitoring a MySQL server that is installed on the same machine as Zabbix server and Zabbix agent 2.



- In the *Macros* tab, switch to *Inherited and host macros*, look for the following macros and click *Change* next to the macro value to update it:
  - {$MYSQL.DSN} - set the data source of the MySQL server (the connection string of a named session from the MySQL Zabbix agent 2 plugin configuration file). This guide uses the default data source "tcp://localhost:3306" for monitoring a MySQL server that is installed on the same machine as Zabbix server and Zabbix agent 2.
  - {$MYSQL.PASSWORD} - set the password of the previously created MySQL user "zbx_monitor".
  - {$MYSQL.USER} - set the name of the previously created MySQL user "zbx_monitor".

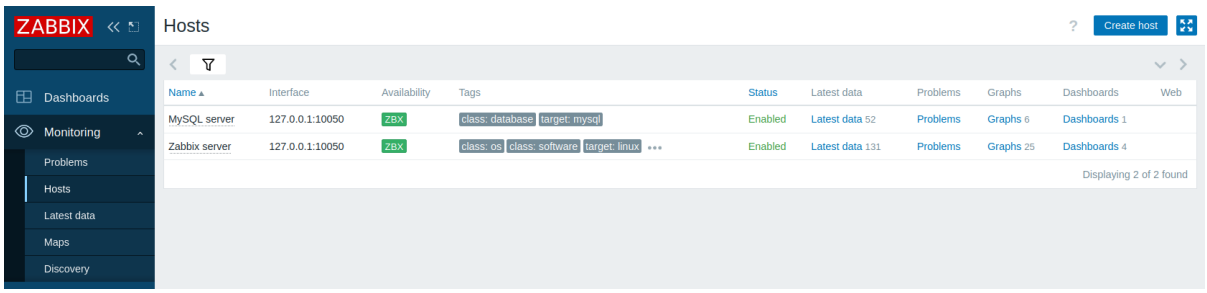


3. Click *Add* to add the host. This host will represent your MySQL server.

View collected metrics

Congratulations! At this point, Zabbix is already monitoring your MySQL server.

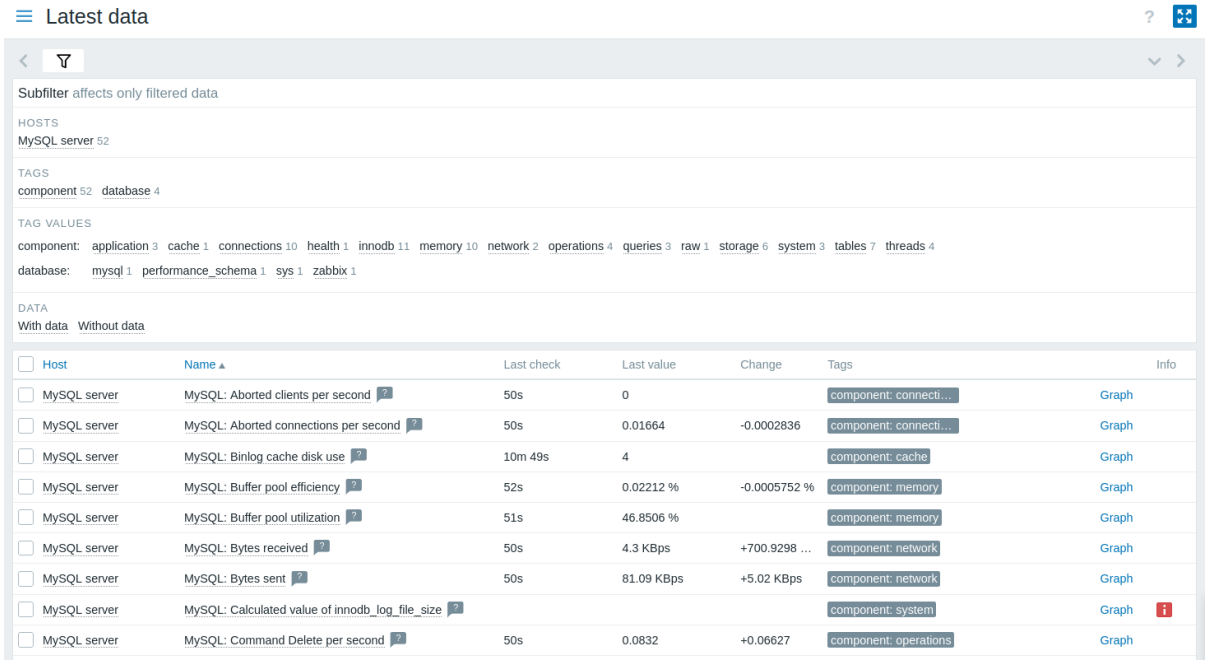
To view collected metrics, navigate to the *Monitoring > Hosts* menu section and click *Dashboards* next to the host.



This action will take you to the host dashboard (configured on the template level) with the most important metrics collected from the MySQL server.



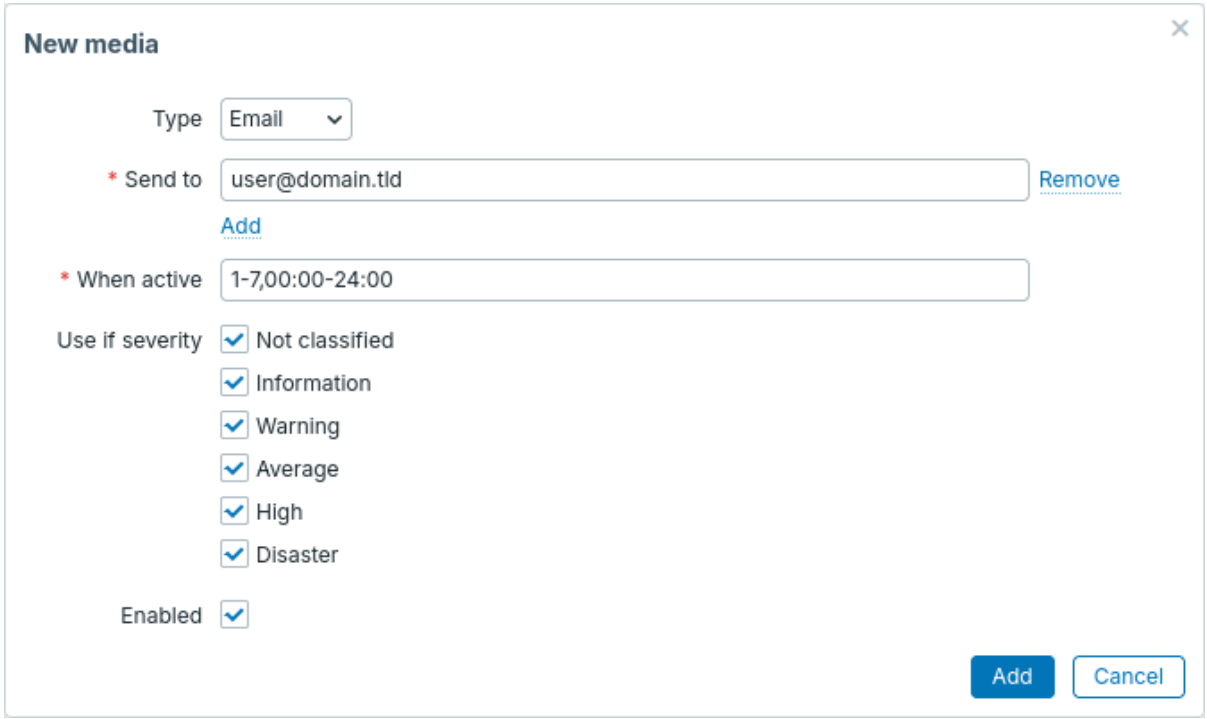
Alternatively, from the *Monitoring > Hosts* menu section, you can click *Latest data* to view all the latest collected metrics in a list. Note that the item *MySQL: Calculated value of innodb_log_file_size* is expected to have no data, as the value will be calculated from data in the last hour.



Set up problem alerts

Zabbix can notify you about a problem with your infrastructure using a variety of methods. This guide provides basic configuration steps for sending email alerts.

1. Navigate to *User settings > Profile*, switch to the *Media* tab and add your email.



2. Follow the guide for Receiving a problem notification.

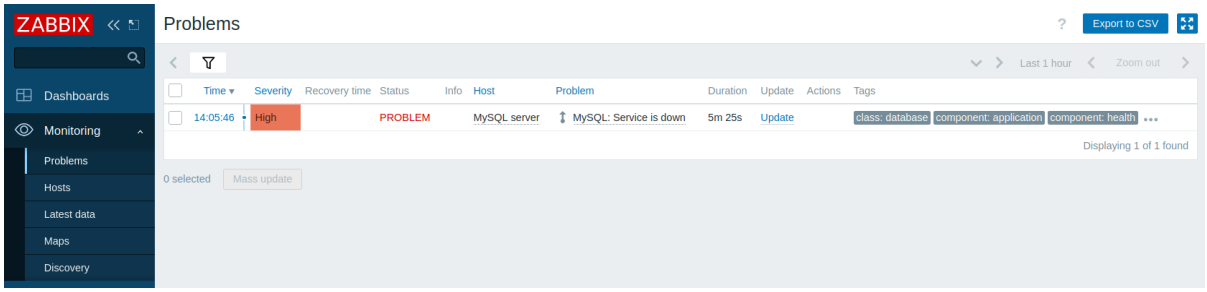
Next time, when Zabbix detects a problem, you should receive an alert via email.

Test your configuration

To test your configuration, we can simulate a real problem by updating the host configuration in Zabbix frontend.

1. Open your MySQL server host configuration in Zabbix.
2. Switch to the *Macros* tab and select *Inherited and host macros*.
3. Click *Change* next to, for example, the previously configured {$MYSQL.USER} macro value and set a different MySQL user name.
4. Click *Update* to update the host configuration.

5. In a few moments, Zabbix will detect the problem "MySQL: Service is down", because it will not be able to connect to the MySQL server. The problem will appear in *Monitoring > Problems*.



If alerts are configured, you will also receive the problem notification.

6. Change the macro value back to its previous value to resolve the problem and continue monitoring the MySQL server.

Other approaches to monitor MySQL

Instead of monitoring a MySQL server with Zabbix agent 2, you could also use Zabbix agent or the Open Database Connectivity (ODBC) standard. While using Zabbix agent 2 is recommended, there might be some setups that do not support Zabbix agent 2 or require a custom approach.

The key difference between Zabbix agent and ODBC lies in the data collection method - Zabbix agent is installed directly on the MySQL server and collects data using its built-in functionality, while ODBC relies on an ODBC driver to establish a connection to the MySQL server and retrieve data using SQL queries.

Although many of the configuration steps are similar to monitoring a MySQL server with Zabbix agent 2, there are some significant differences - you need to configure Zabbix agent or ODBC to be able to monitor a MySQL server. The following instructions walk you through these **differences**.

Monitor MySQL with Zabbix agent

To monitor a MySQL server with Zabbix agent, you need to download and install Zabbix server, Zabbix frontend and Zabbix agent according to the instructions for your OS.

Once you have successfully installed the required Zabbix components, you need to create a MySQL user as described in the *Create MySQL user* section.

After you have created the MySQL user, you need to configure Zabbix agent to be able to establish a connection with the MySQL server and monitor it. This includes configuring multiple user parameters for executing custom agent checks, as well as providing Zabbix agent with the necessary credentials for connecting to the MySQL server as the previously created "zbx_monitor" user.

Configure Zabbix agent

1. Navigate to the Zabbix agent additional configurations directory.

```
cd /etc/zabbix/zabbix_agentd.d
```

> **Attention:**
> The Zabbix agent additional configurations directory should be located in the same directory as your Zabbix agent configuration file; the default path depends on the OS and Zabbix installation method.

Instead of defining all of the necessary user parameters for monitoring the MySQL server in the Zabbix agent configuration file, these parameters will be defined in a separate file in the additional configurations directory.

2. Create a `template_db_mysql.conf` file in the Zabbix agent additional configurations directory.

```
vi template_db_mysql.conf
```

3. Copy the contents from the `template_db_mysql.conf` file (located in the Zabbix repository) to the `template_db_mysql.conf` file you created, and save.

4. Restart Zabbix agent to update its configuration.

```
systemctl restart zabbix-agent
```

Once you have configured Zabbix agent user parameters, you can move on to configure the credentials that will allow Zabbix agent to access the MySQL server.

5. Navigate to the Zabbix agent home directory (if it does not exist on your system, you need to create it; default: */var/lib/zabbix*).

```
cd /var/lib/zabbix
```

6. Create a *.my.cnf* file in the Zabbix agent home directory.

```
vi .my.cnf
```

7. Copy the following contents to the *.my.cnf* file (replace *<password>* with the password of the "zbx_monitor" user).

```
[client]
user='zbx_monitor'
password='<password>'
```

Configure Zabbix frontend and test your configuration

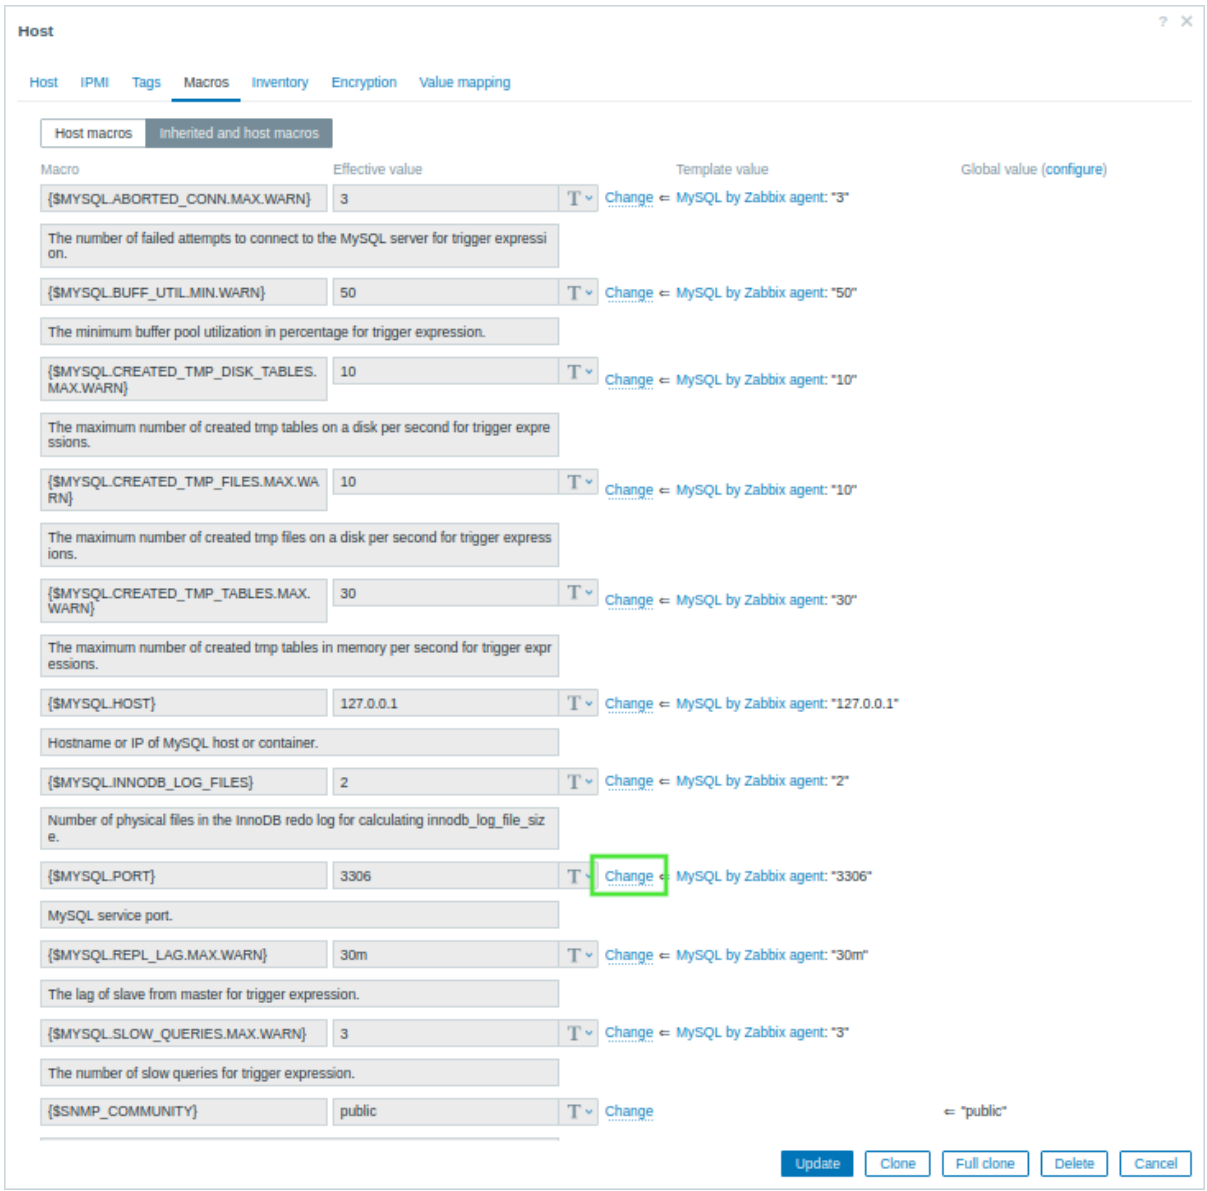
To configure Zabbix frontend, follow the instructions in the *Configure Zabbix frontend* section with the following adjustments:

- In the *Templates* field, type or select the template "MySQL by Zabbix agent" that will be linked to the host.
- Configuring *Macros* is not required.

Once you have configured Zabbix frontend, you can view collected metrics and set up problem alerts.

To test your configuration, follow the instructions in the *Test your configuration* section with the following adjustments:

- In the *Inherited and host macros* section of the MySQL server host configuration, click *Change* next to the {$MYSQL.PORT} macro value and set a different port (e.g., "6033").



Monitor MySQL with ODBC

To monitor a MySQL server with ODBC, you need to download and install Zabbix server and Zabbix frontend.

Once you have successfully installed the required Zabbix components, you need to create a MySQL user as described in the *Create MySQL user* section.

After you have created the MySQL user, you need to setup ODBC. This includes installing one of the most commonly used open source ODBC API implementations - unixODBC - and a unixODBC driver, as well as editing the ODBC driver configuration file.

Configure ODBC

1. Install unixODBC. The recommended way of installing unixODBC is to use the Linux operating system default package repositories.

```
apt install unixodbc
```

2. Install the MariaDB unixODBC database driver. Although you have a MySQL database, the MariaDB unixODBC driver is used for compatibility issues.

```
apt install odbc-mariadb
```

3. Check the location of the ODBC configuration files *odbcinst.ini* and *odbc.ini*.

```
odbcinst -j
```

The result of executing this command should be similar to the following.

```
unixODBC 2.3.9
DRIVERS............: /etc/odbcinst.ini
SYSTEM DATA SOURCES: /etc/odbc.ini
FILE DATA SOURCES..: /etc/ODBCDataSources
...
```

4. To configure the ODBC driver for monitoring a MySQL database, you need the driver name, which is located in the *odbcinst.ini* file. In the following *odbcinst.ini* file example, the driver name is "MariaDB Unicode".

```
[MariaDB Unicode]
Driver=libmaodbc.so
Description=MariaDB Connector/ODBC(Unicode)
Threading=0
UsageCount=1
```

5. Copy the following contents to the *odbc.ini* file (replace <*password*> with the password of the "zbx_monitor" user). This guide uses "127.0.0.1" (localhost) as the MySQL server address for monitoring a MySQL server that is installed on the same machine as the ODBC driver. Note the data source name (DSN) "test", which will be required when configure Zabbix frontend.

```
[test]
Driver=MariaDB Unicode
Server=127.0.0.1
User=zbx_monitor
Password=<password>
Port=3306
Database=zabbix
```

Configure Zabbix frontend and test your configuration

To configure Zabbix frontend, follow the instructions in the *Configure Zabbix frontend* section with the following adjustments:

- In the *Templates* field, type or select the template "MySQL by ODBC" that will be linked to the host.
- Configuring *Interfaces* is not required.
- The {$MYSQL.DSN} macro value In the *Inherited and host macros* section of the MySQL server host configuration should be set to the DSN name from the *odbc.ini* file.

Once you have configured Zabbix frontend, you can view collected metrics, set up problem alerts and test your configuration.

See also

- Creating an item - how to start monitoring additional metrics.
- Problem escalations - how to create multi-step alert scenarios (e.g., first send message to the system administrator, then, if a problem is not resolved in 45 minutes, send message to the data center manager).
- ODBC monitoring - how to set up ODBC on other Linux distributions, and how to start monitoring additional database-related metrics with ODBC.
- Template *MySQL by Zabbix agent* - additional information about the *MySQL by Zabbix agent* template.
- Template *MySQL by Zabbix agent 2* - additional information about the *MySQL by Zabbix agent 2* template.
- Template *MySQL by ODBC* - additional information about the *MySQL by ODBC* template.

Monitor VMware with Zabbix

Introduction

This page walks you through the steps required to start basic monitoring of VMware.

Who this guide is for

This guide is designed for new Zabbix users and contains the minimum set of steps required to enable basic monitoring of VMware. If you are looking for deep customization options or require more advanced configuration, see the Virtual machine monitoring section or the Configuration section of Zabbix manual.

Prerequisites

Before proceeding with this guide, you need to download and install Zabbix server and Zabbix frontend according to the instructions for your OS.

Depending on your environment, some steps in this guide may slightly differ. This guide is based on an environment running Ubuntu.

It is assumed that VMware is already configured. This guide does not cover the configuration of VMware.

Configure Zabbix server

To monitor VMware, the *vmware collector* Zabbix processes need to be enabled. For more information on how VMware monitoring is performed, see Virtual machine monitoring.

1. Open the Zabbix server configuration file.

```
vi /etc/zabbix/zabbix_server.conf
```

2. Locate and set the `StartVMwareCollectors` parameter in Zabbix server configuration file to 2 or more (the default value is 0).

```
###### Option: StartVMwareCollectors
#####       Number of pre-forked vmware collector instances.
#####
##### Mandatory: no
##### Range: 0-250
##### Default:
##### StartVMwareCollectors=0

StartVMwareCollectors=2
```

3. Restart Zabbix server.

```
systemctl restart zabbix-server
```

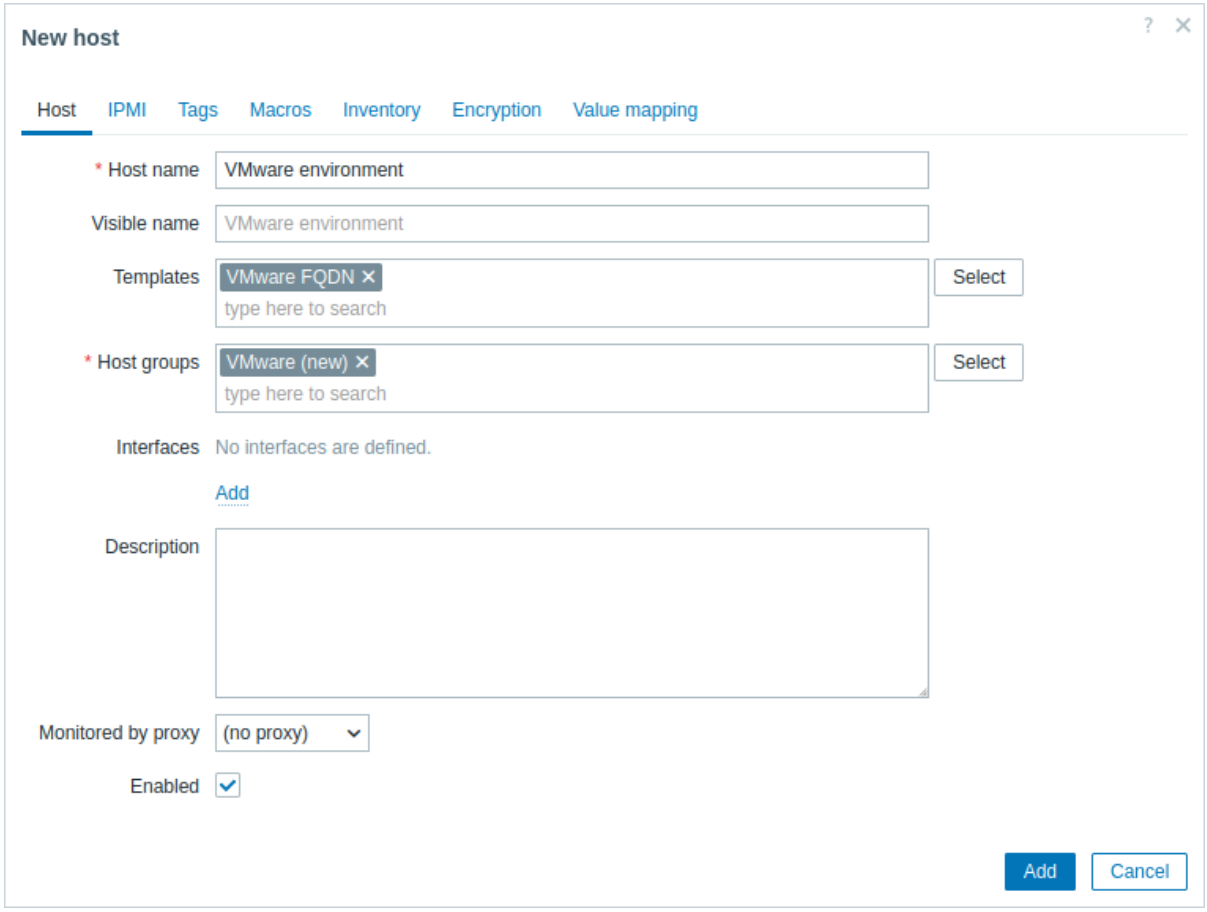
Once the *vmware collector* processes have been started, move on to the next step.

Configure Zabbix frontend

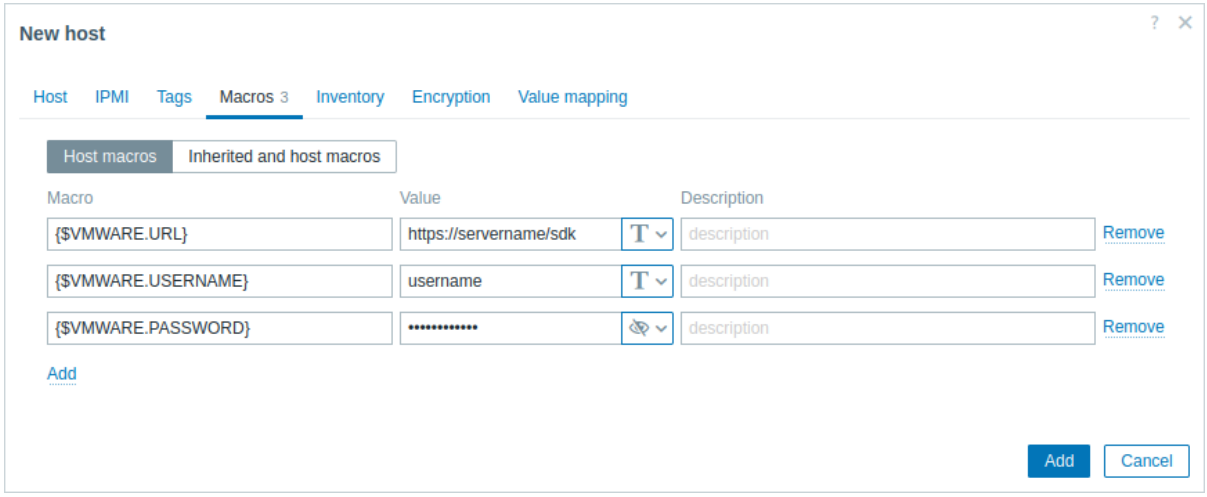
1. Log into Zabbix frontend.

2. Create a host in Zabbix web interface:

- In the *Host name* field, enter a host name (for example, ”VMware environment”).
- In the *Templates* field, type or select the ”VMware FQDN” (or ”VMware”) template. For more information on these templates, see Virtual machine monitoring.
- In the *Host groups* field, type or select a host group (for example, a new host group ”VMware”).



- In the *Macros* tab, set the following host macros:
  - {$VMWARE.URL} - VMware service (vCenter or ESXi hypervisor) SDK URL (https://servername/sdk)
  - {$VMWARE.USERNAME} - VMware service user name
  - {$VMWARE.PASSWORD} - VMware service {$VMWARE.USERNAME} user password



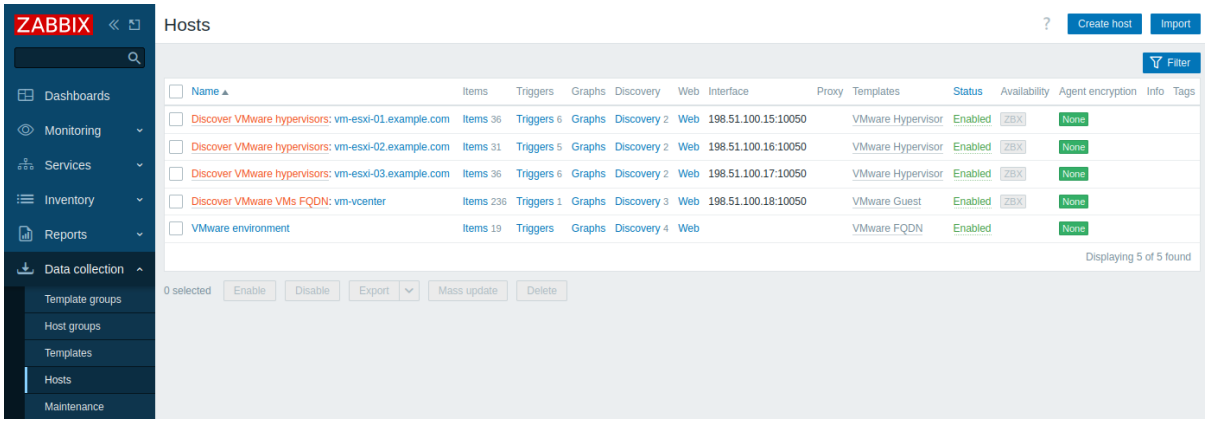
3. Click the *Add* button to create the host. This host will represent your VMware environment.

View collected metrics

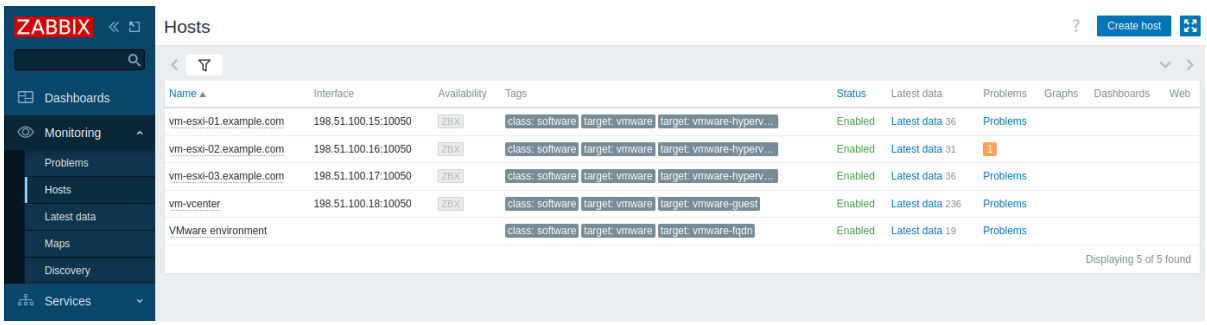
Congratulations! At this point, Zabbix is already monitoring your VMware environment.

Depending on the configuration of your VMware environment, Zabbix may discover and then create hosts for the discovered entities. Note that the discovery and creation of hosts can also be executed manually, if necessary.

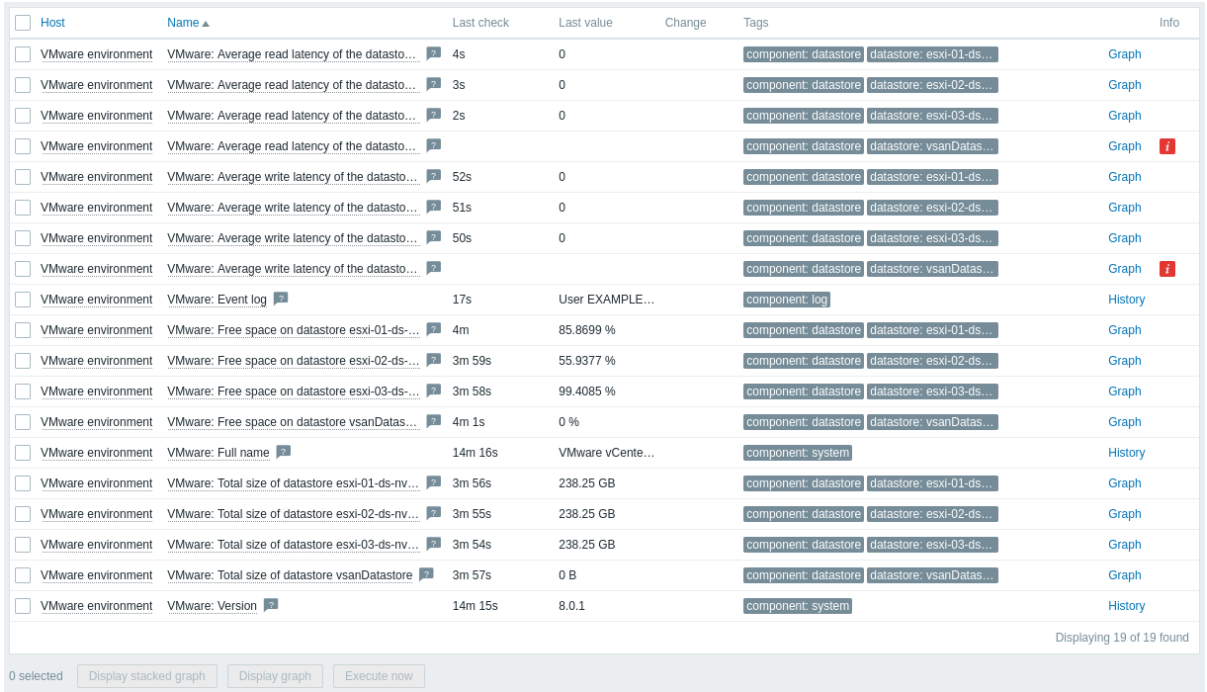
To view created hosts, navigate to the *Data collection > Hosts* menu section.



To view collected metrics, navigate to the *Monitoring > Hosts* menu section and click *Latest data* next to the created "VMware environment" host or one of the hosts that were created for the discovered entities.



This action will open a list of all the latest metrics collected from the selected host.

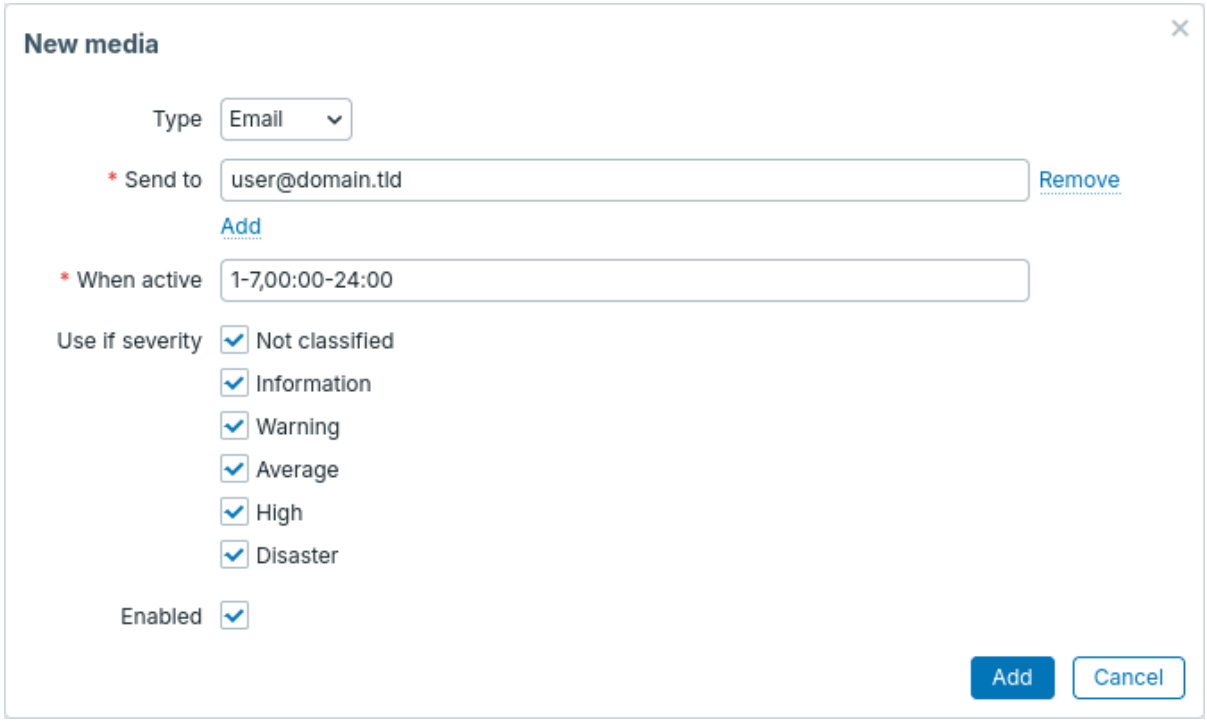


Note that some items have no data and the *Not supported* state. This is because Zabbix cannot find valid performance counters on the specific datastore, as it is not enabled in the VMware environment being monitored.

Set up problem alerts

Zabbix can notify you about a problem with your infrastructure using a variety of methods. This guide provides basic configuration steps for sending email alerts.

1. Navigate to *User settings > Profile*, switch to the *Media* tab and add your email.



2. Follow the guide for Receiving a problem notification.

Next time, when Zabbix detects a problem, you should receive an alert via email.

See also

- Creating an item - how to start monitoring additional metrics.
- Problem escalations - how to create multi-step alert scenarios (e.g., first send message to the system administrator, then, if a problem is not resolved in 45 minutes, send message to the data center manager).
- Virtual machine monitoring - additional information about VMware monitoring (data collection process, server configuration options, troubleshooting guidance, etc.).
- VMware monitoring item keys - a full list of VMware metrics that can be monitored using Zabbix.
- Template *VMware* - additional information about the *VMware* template.
- Template *VMware FQDN* - additional information about the *VMware FQDN* template.

Monitor network traffic with Zabbix

Introduction

This page walks you through the steps required to start basic monitoring of your network traffic with Zabbix.

Who this guide is for

This guide is designed for new Zabbix users and contains the minimum set of steps required to enable basic monitoring of your network traffic. If you are looking for deep customization options or require more advanced configuration, see the Configuration section of Zabbix manual.

Prerequisites

Before proceeding with this guide, you need to download and install Zabbix server, Zabbix frontend and Zabbix agent according to the instructions for your OS. Note that you should install Zabbix agent on the machine that requires traffic monitoring. This can be either the same host where Zabbix server is installed or a different host.

This guide will provide the instructions on configuring the network traffic monitoring of *eth0* interface on a separate machine named *Remote host*.

Configure Zabbix for monitoring

Zabbix agent can collect metrics in active or passive mode (simultaneously). See Passive and active agent checks for more info. In this guide, monitoring by means of passive checks will be described.

Configuring Zabbix agent

1. Open the agent configuration file on the machine where the agent is installed:

```
sudo vi /etc/zabbix/zabbix_agentd.conf
```

2. Add the IP address or DNS name of your Zabbix server to the *Server* parameter. For example:

```
Server=192.0.2.22
```

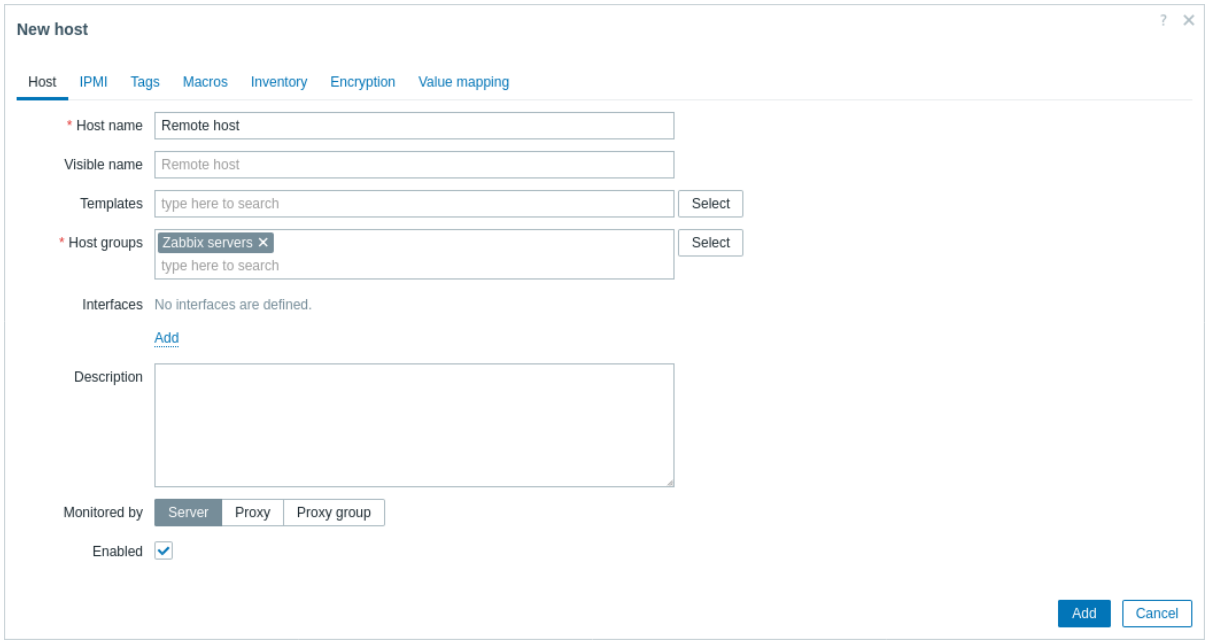
3. Restart Zabbix agent:

```
systemctl restart zabbix-agent
```

Zabbix frontend

1. Log into Zabbix frontend.

2. Create a host in Zabbix web interface, specifying the IP address or DNS name of the machine on which the agent is installed.

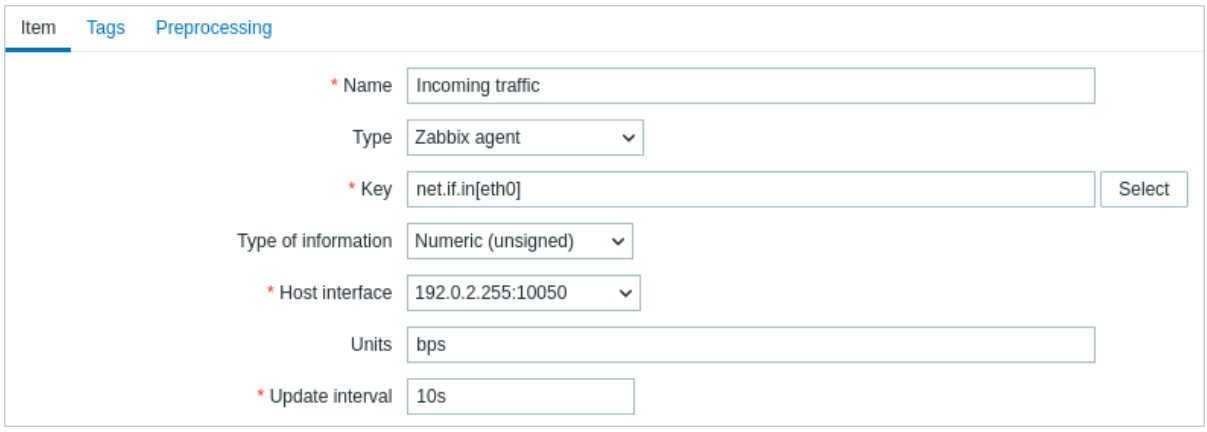


Create items

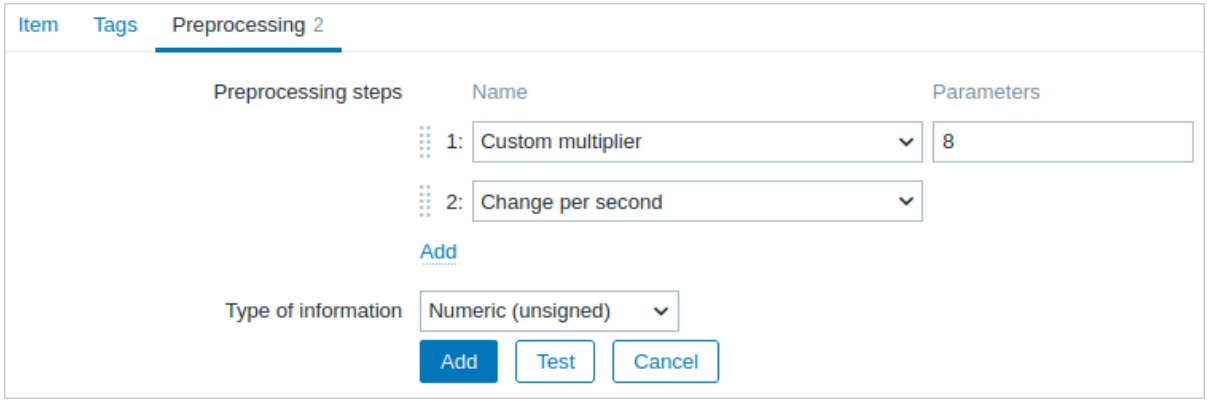
Follow the instructions on creating an item to add the items for traffic monitoring, namely:

- Incoming traffic
- Outgoing traffic
- Total traffic

A simple setup for the incoming traffic monitoring item would look as follows:



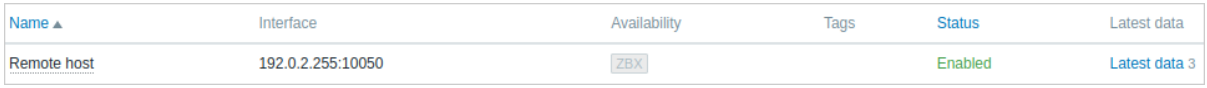
To make the collected data suitable for practical use, you can set some preprocessing steps when creating the items. In the given case, those can be multiplication by 8 (to convert bytes to bits) and presentation as change per second.



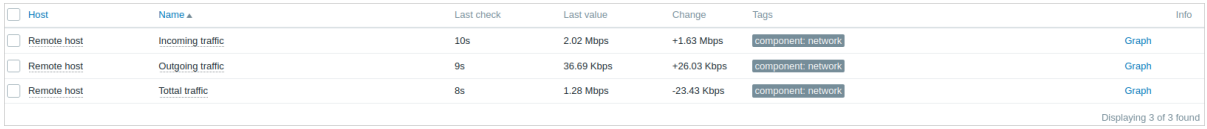
View collected data

Congratulations! At this point, Zabbix is already monitoring your network traffic.

To view collected metrics, open the *Monitoring > Hosts* menu section and click *Latest data* in the row of the host.

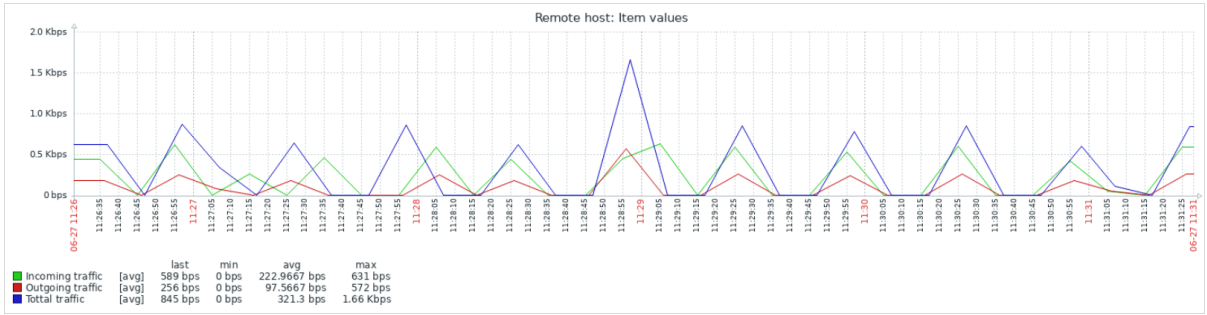


You will see the traffic data.



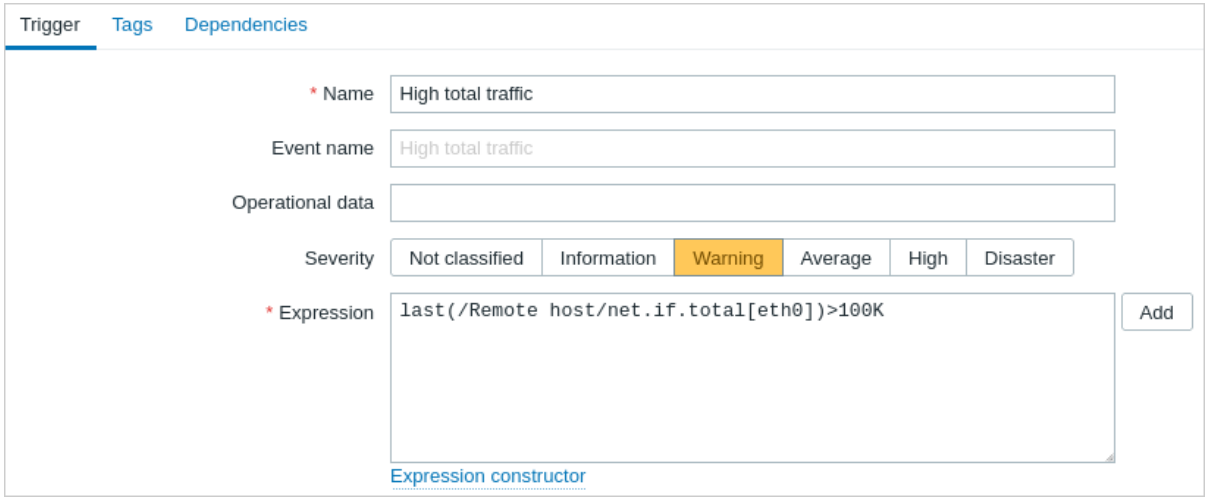
View graphs

The data collected can be displayed as graphs. To view those, in the *Latest data* section click *Graph* in the row of the item or select required items and click *Display graph* below.

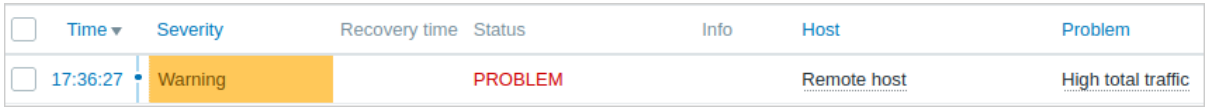


Configure triggers

You can set triggers to detect abnormal network traffic. See the instructions on configuring a trigger and add the trigger signalling that the total traffic is too high, for example:



Now make the traffic exceed the threshold you have set in the trigger expression and navigate to the *Monitoring > Problems* to check that the problem is listed there.



Set up problem alerts

There are several ways of getting notifications about the problem. Email being the most popular one, follow the instructions on setting up a problem notification via email. You can also choose other media types to be used for the notification delivery.

See also:

- Problem escalations - how to create multi-step alert scenarios (e.g., first send message to the system administrator, then, if a problem is not resolved in 45 minutes, send message to the data center manager).
- Problem acknowledgment - how to indicate that the problem is known, make comments on resolving it, suppress or close the problem.
- Monitor Linux with Zabbix agent - how to start basic monitoring of the most important items by linking a pre-configured template.

Monitor network traffic using active checks

Introduction

This page walks you through the steps required to start basic monitoring of your network traffic with Zabbix using active checks.

Who this guide is for

This guide is designed for new Zabbix users and contains the minimum set of steps required to enable basic monitoring of your network traffic using active checks. If you are looking for deep customization options or require more advanced configuration, see the Configuration section of Zabbix manual.

Prerequisites

Before proceeding with this guide, you need to download and install Zabbix server, Zabbix frontend and Zabbix agent according to the instructions for your OS. Note that you should install Zabbix agent on the machine that requires traffic monitoring. This can be either the same host where Zabbix server is installed or a different host.

This guide will provide the instructions on configuring the network traffic monitoring of *eth0* interface on a separate machine named *Remote host*.

Configure Zabbix for monitoring

Zabbix agent can collect metrics in active or passive mode (simultaneously). See Passive and active agent checks for more info. In this guide, monitoring by means of **active checks** will be described.

Configuring Zabbix agent

1. Open the agent configuration file on the machine where the agent is installed.

If you're using Zabbix agent:

```
sudo vi /etc/zabbix/zabbix_agentd.conf
```

If you're using Zabbix agent 2:

```
sudo vi /etc/zabbix/zabbix_agent2.conf
```

2. Add the IP address or DNS name (and optional port) of your Zabbix server to the ServerActive parameter. For example:

```
ServerActive=192.0.2.1:10051
```

Zabbix agent will use this address to connect to Zabbix server's trapper port (default: 10051) and request the configuration data for active checks.

3. Define the Hostname parameter, which must match the host name that will be defined in Zabbix frontend. In our example, it is:

```
Hostname=Remote host
```

The Hostname value must match because, for active checks, Zabbix agent uses it to retrieve the correct host configuration from the server. Specifically, the agent initiates a connection to the server and identifies itself using the Hostname value. The server then provides the monitoring configuration for that host. If these values differ, the agent will not receive the appropriate configuration, resulting in missing metrics or monitoring issues.

4. Restart Zabbix agent.

If you're using Zabbix agent:

```
systemctl restart zabbix-agent
```

If you're using Zabbix agent 2:

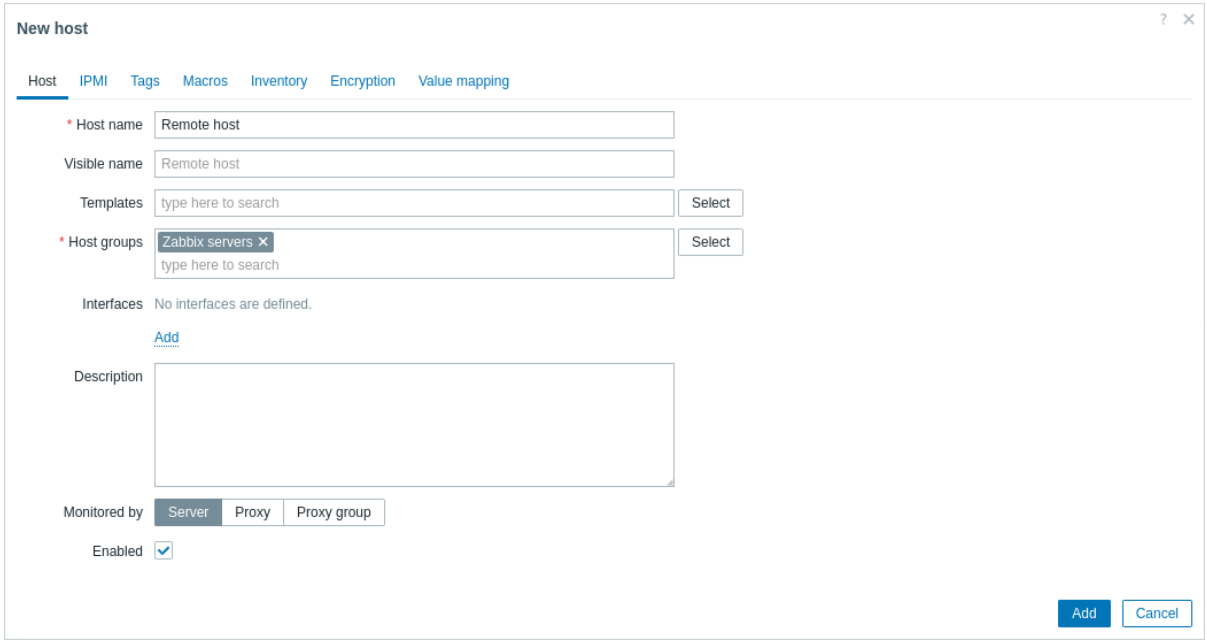
```
systemctl restart zabbix-agent2
```

Zabbix frontend

1. Log into Zabbix frontend.

2. Create a host in Zabbix web interface.

- In the *Host name* field, enter a host name (e.g., "Remote host") that matches the Hostname parameter value defined earlier in the agent configuration file.
- In the *Host groups* field, type or select a host group (e.g., "Zabbix servers").
- Leave *Interfaces* undefined; an interface is not necessary for active checks because the agent initiates the connection to the server instead of listening for a connection from the server.



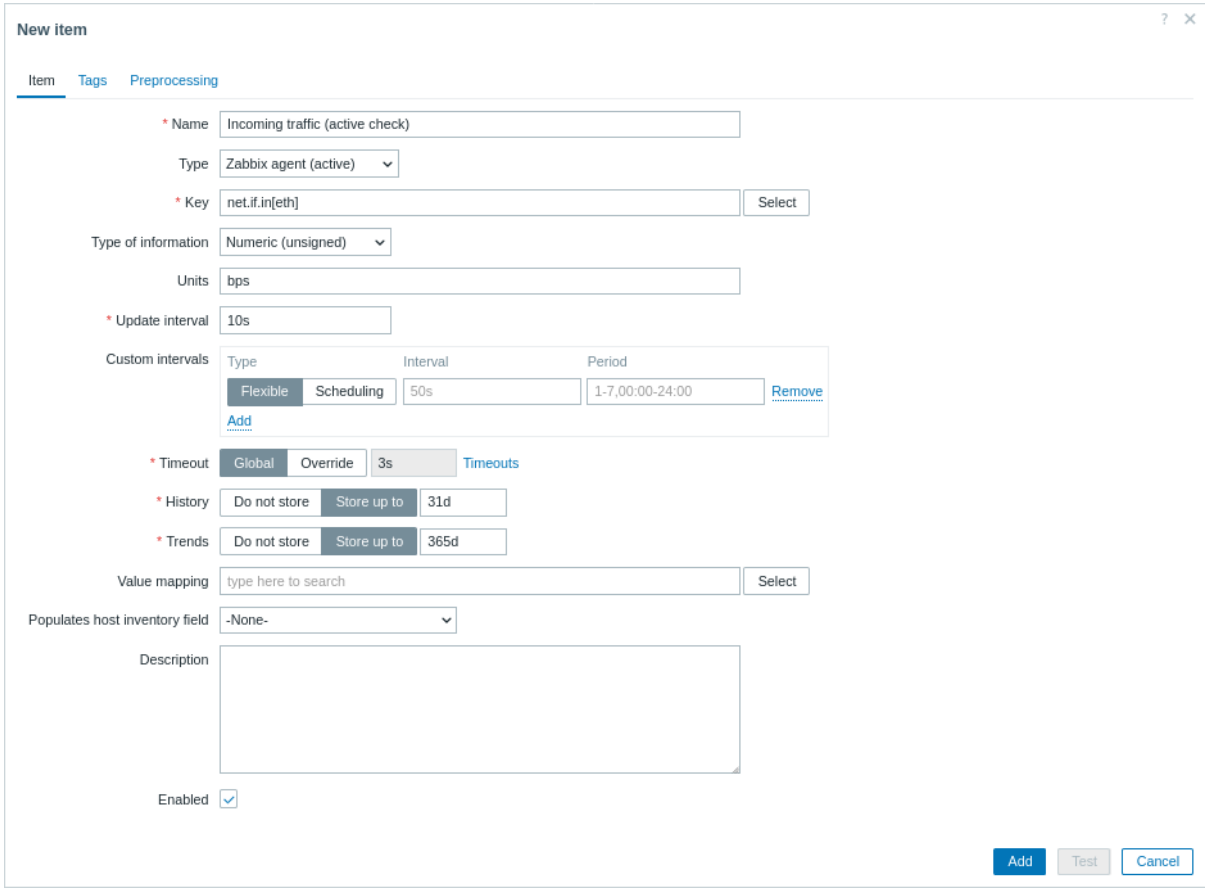
3. Click *Add* to add the host. This host will represent the monitored Linux machine.

Create items

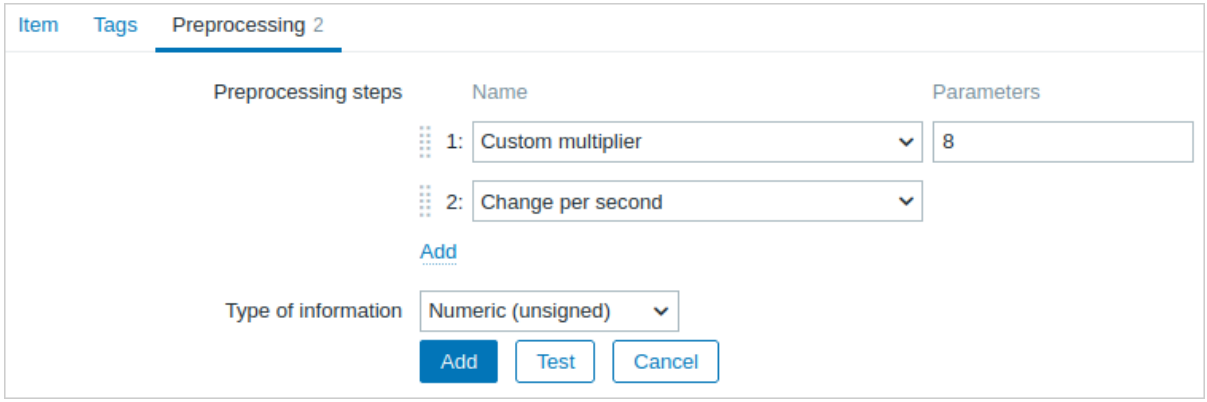
Follow the instructions on creating an item to add the items for traffic monitoring, namely:

- Incoming traffic
- Outgoing traffic
- Total traffic

A simple setup for the incoming traffic monitoring item with an active check would look as follows:



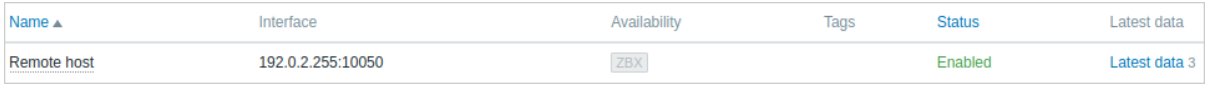
To make the collected data suitable for practical use, you can set some preprocessing steps when creating the items. In the given case, those can be multiplication by 8 (to convert bytes to bits) and presentation as change per second.



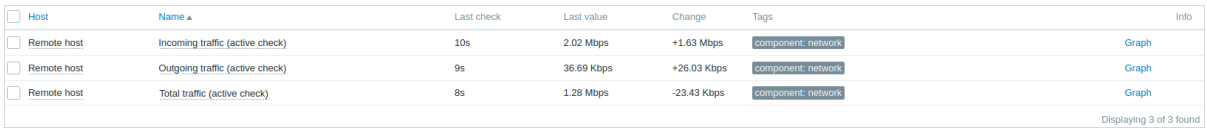
View collected data

Congratulations! At this point, Zabbix is already monitoring your network traffic.

To view collected metrics, open the *Monitoring > Hosts* menu section and click the *Latest data* in the row of the host.



You will see the traffic data.



See also:

- View graphs - how to display the collected data as graphs.
- Configure triggers - how to set up triggers to detect abnormal network traffic.
- Set up problem alerts - how to set up notifications about problem situations.

- Problem escalations - how to create multi-step alert scenarios (e.g., first send message to the system administrator, then, if a problem is not resolved in 45 minutes, send message to the data center manager).
- Problem acknowledgment - how to indicate that the problem is known, make comments on resolving it, suppress or close the problem.
- Monitor Linux with Zabbix agent - how to start basic monitoring of the most important items by linking a pre-configured template.
- Installation from packages - how to install Zabbix components using official RPM and DEB packages for various Linux distributions, ensuring access to the latest features and bug fixes.

Monitor websites with Browser items

Introduction

This page walks you through the steps required to start basic monitoring of websites with Browser items.

Who this guide is for

This guide is designed for new Zabbix users and contains the minimum set of steps required to enable basic monitoring of websites with Browser items. If you are looking for deep customization options or require more advanced configuration, see the Browser items page or the Configuration section of Zabbix manual.

Prerequisites

Before proceeding with this guide, you need to download and install Zabbix server and Zabbix frontend according to the instructions for your OS.

Depending on your environment, some steps in this guide may slightly differ. This guide is based on an environment running Ubuntu.

Configure WebDriver

Browser items require an automation framework (either Selenium Server or a plain WebDriver, for example, ChromeDriver) as a web testing endpoint that controls and interacts with a browser, executing test commands such as clicking buttons or entering text. As an example, this guide will use Selenium Server with Chrome in a Docker container.

It is assumed that Docker is already configured. This guide does not cover the configuration of Docker. For installation instructions, see Install Docker Engine on Ubuntu.

1. Launch Selenium Server with Chrome in a Docker container with the following options:

- **docker run --name browser** - runs a new Docker container named "browser";
- **-p 4444:4444** - maps port 4444 on your host machine to port 4444 on the container (this is the port used by Selenium Server to accept commands);
- **-p 7900:7900** - maps port 7900 on your host machine to port 7900 on the container (this is the port used by the Virtual Network Computing (VNC) server, allowing you to view the browser GUI remotely; requires a VNC client);
- **--shm-size="2g"** - allocates 2GB of shared memory to the container (this is important for Chrome to run properly, as it can require a significant amount of shared memory to avoid crashes);
- **-d** - runs the container in detached mode, meaning it will run in the background;
- **selenium/standalone-chrome:latest** - specifies the Docker image to use; in this case, the latest version of Selenium Server with Chrome.

```
docker run --name browser \
-p 4444:4444 \
-p 7900:7900 \
--shm-size="2g" \
-d selenium/standalone-chrome:latest
```

2. Ensure that the `browser` Docker container is running and accessible.

- Retrieve the IP address of the container (in this example, 192.0.2.1):

```
ip addr

##### 1: lo: <LOOPBACK,UP,LOWER_UP>
#####     ...
##### 3: docker0: <NO-CARRIER,BROADCAST,MULTICAST,UP> ...
#####     inet 192.0.2.1/16 brd 192.0.255.255 scope global docker0
#####     ...
```

- Test the connection to the container with Ncat:

```
nc -zv 192.0.2.1 4444

##### Connection to 192.0.2.1 4444 port [tcp/*] succeeded!
```

- Retrieve the webpage content from Selenium Server with curl:

```
curl -L 192.0.2.1:4444

##### <!DOCTYPE html>
##### <html lang="en">
#####
##### <head>
#####   <meta charset="utf-8"/>
#####   <link href="favicon.svg" rel="icon" type="image/svg">
#####   <meta content="width=device-width, initial-scale=1" name="viewport"/>
#####   <link href="logo192.png" rel="apple-touch-icon"/>
#####   <link href="manifest.json" rel="manifest"/>
#####   <title>Selenium Grid</title>
##### </head>
#####
##### <body>
##### ...
```

Note:
For troubleshooting, refer to Docker documentation.

Configure Zabbix server

Browser items are executed and processed by *browser poller* Zabbix processes that need to be enabled by adjusting the StartBrowserPollers server configuration parameter. Additionally, the WebDriverURL parameter should specify the previously configured web testing endpoint.

By default, the StartBrowserPollers parameter is set to 1, therefore you only need to specify the web testing endpoint.

1. Open the Zabbix server configuration file:

```
vi /etc/zabbix/zabbix_server.conf
```

2. Locate and set the WebDriverURL parameter in Zabbix server configuration file:

```
###### Option: WebDriverURL
#####WebDriver interface HTTP(S) URL. For example, http://localhost:4444 used with Selenium WebDriver star
#####
##### Mandatory: no
##### Default:
##### WebDriverURL=

WebDriverURL=192.0.2.1:4444
```

3. Restart Zabbix server.

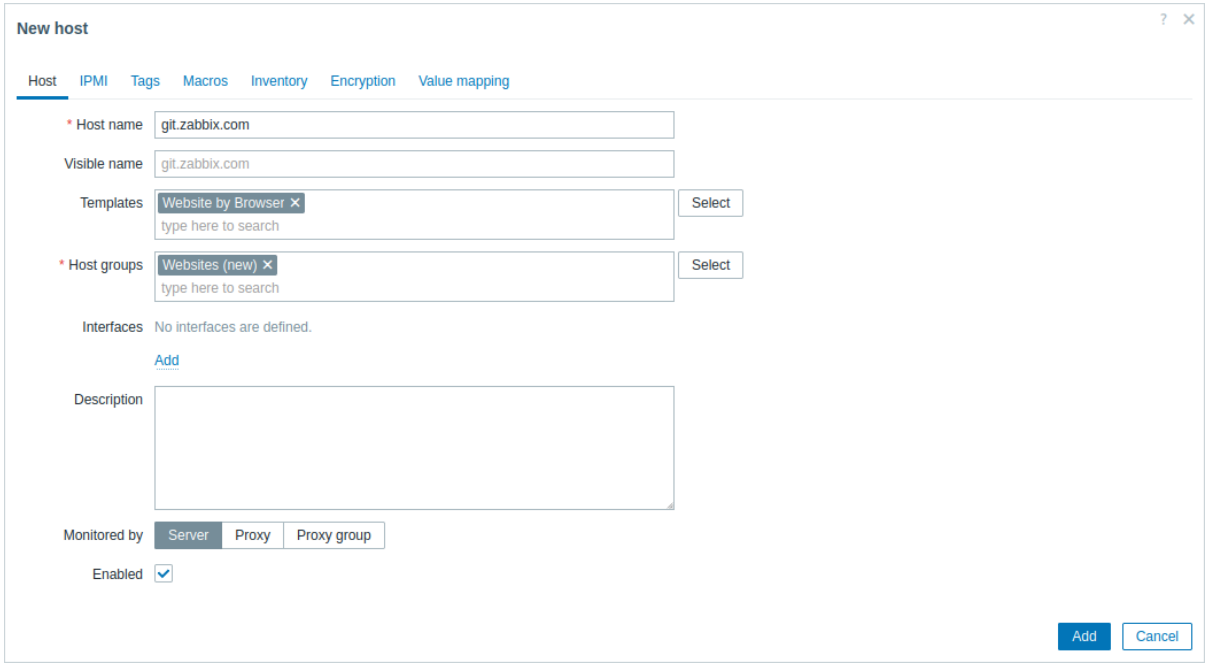
```
systemctl restart zabbix-server
```

Configure Zabbix frontend

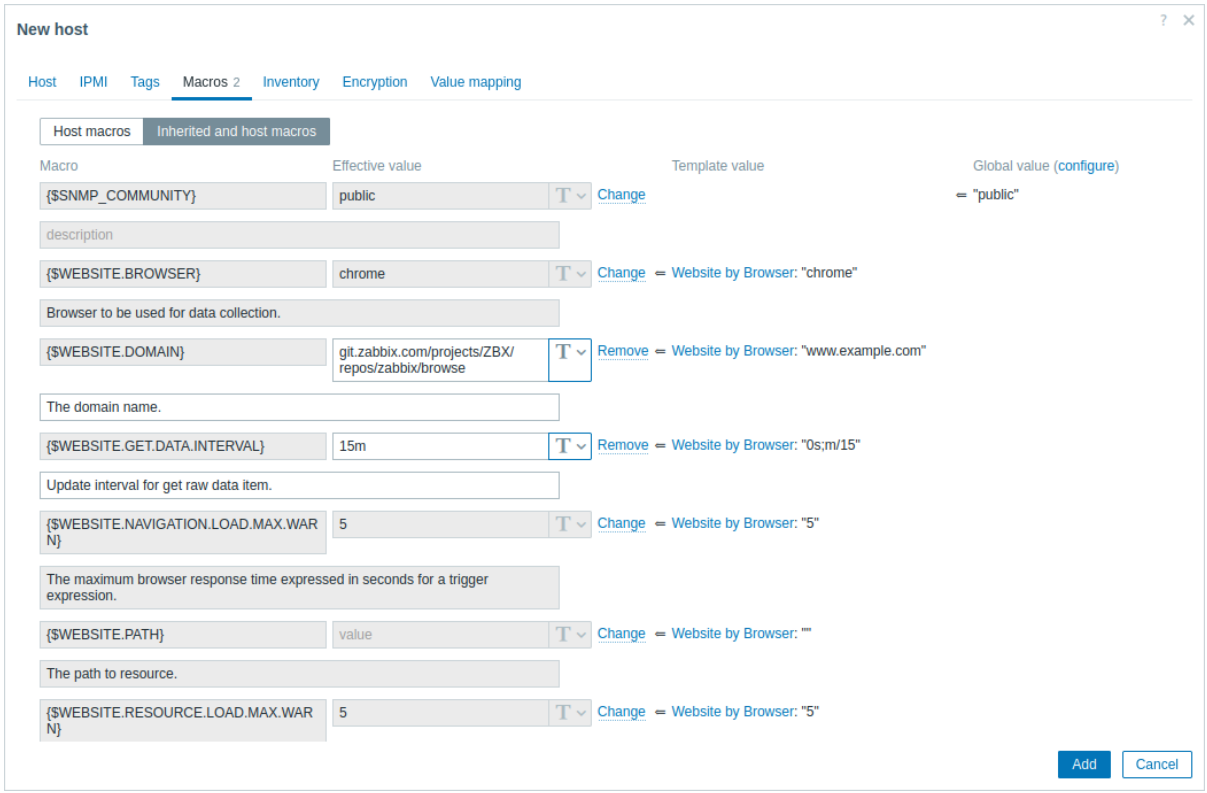
1. Log into Zabbix frontend.

2. Create a host in Zabbix web interface:

- In the *Host name* field, enter a host name (for example, ”git.zabbix.com”).
- In the *Templates* field, type or select the ”Website by Browser” template. For more information on this template, see Website by Browser.
- In the *Host groups* field, type or select a host group (for example, a new host group ”Websites”).



- In the *Macros* tab, switch to *Inherited and host macros*, look for the following macros and click *Change* next to the macro value to update it:
  - {$WEBSITE.DOMAIN} - Domain name (for example, git.zabbix.com/projects/ZBX/repos/zabbix/browse)
  - {$WEBSITE.GET.DATA.INTERVAL} - Item data update interval (for example, 15m)

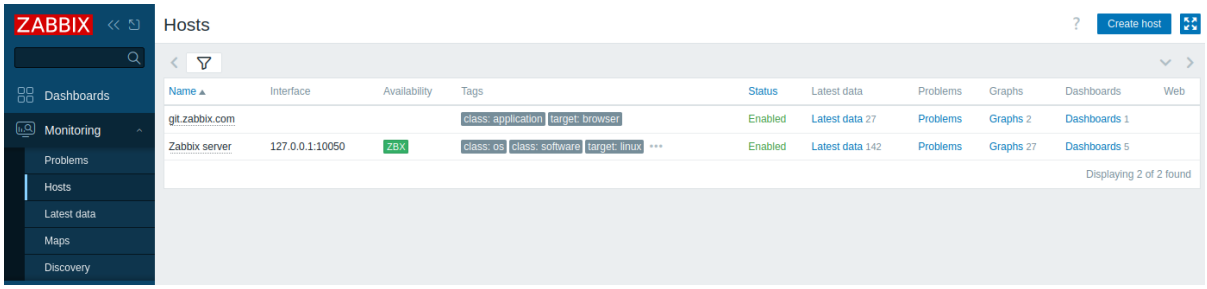


3. Click the *Add* button to create the host. This host will represent the website you want to monitor.

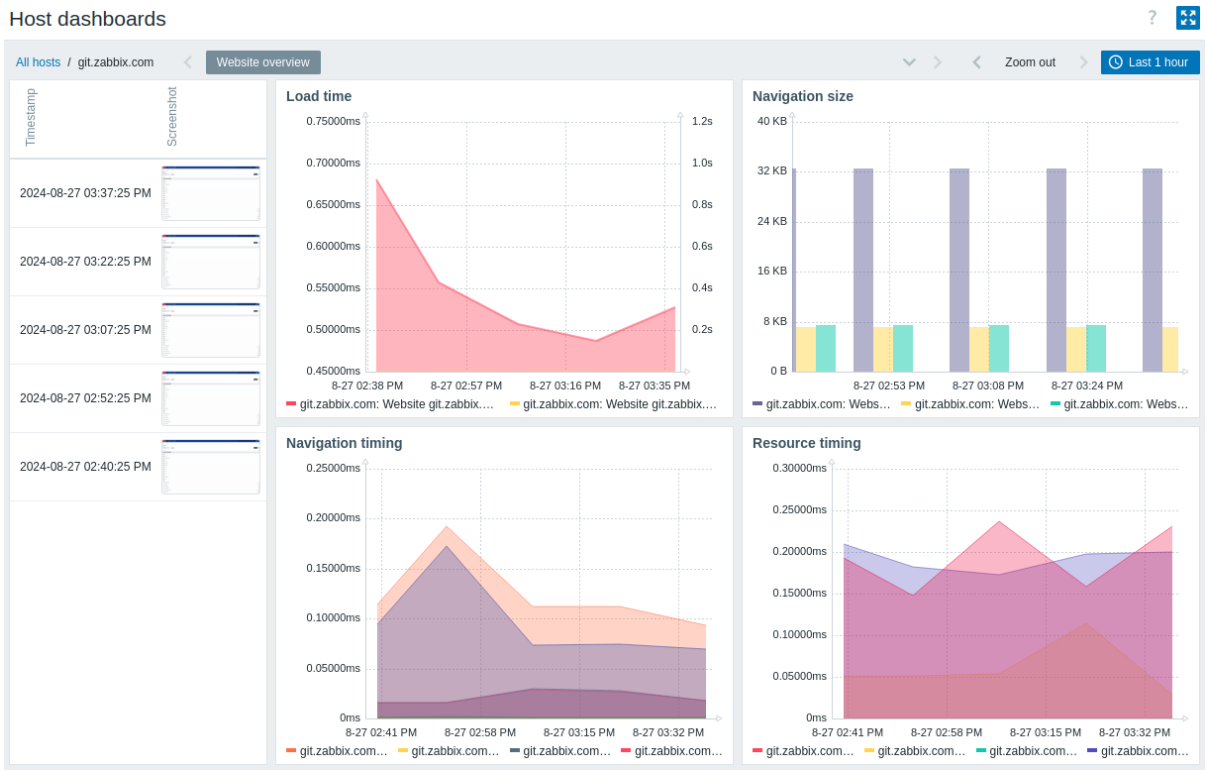
View collected metrics

Congratulations! At this point, Zabbix is already monitoring the website you specified.

To view collected metrics, navigate to the *Monitoring > Hosts* menu section and click *Dashboards* next to the host.



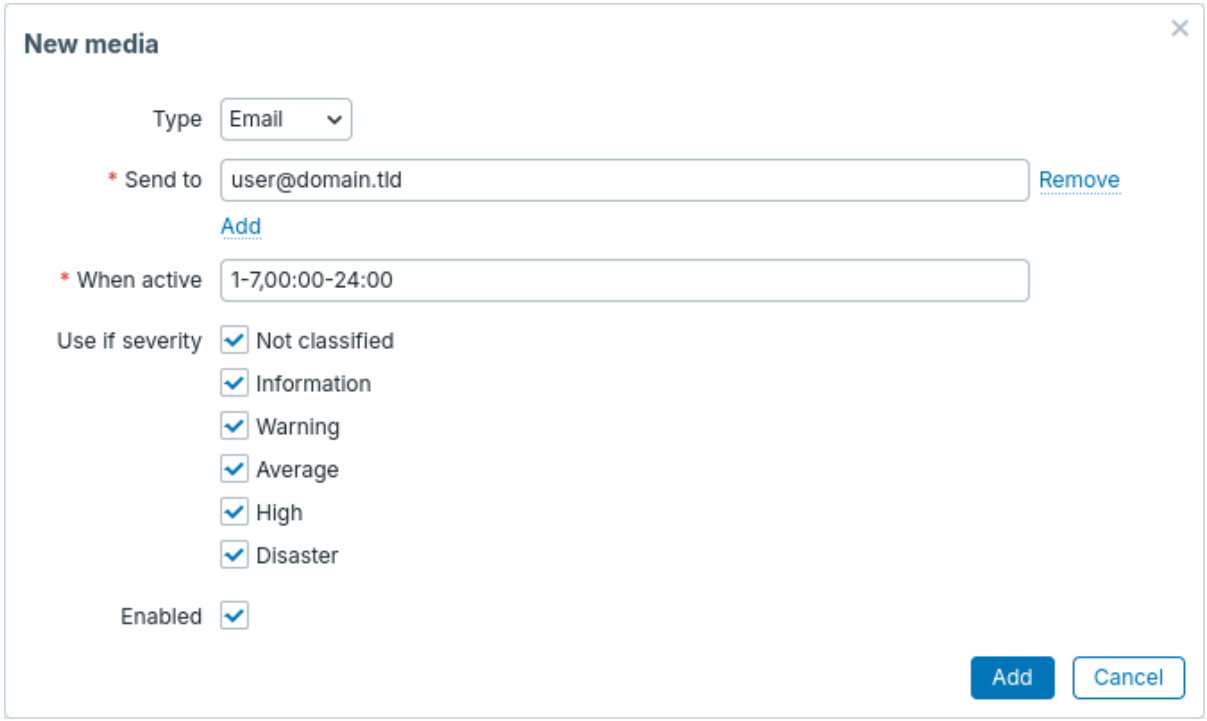
This action will take you to the host dashboard (configured on the template level) with the most important metrics collected from the website.



Set up problem alerts

Zabbix can notify you about a problem with your infrastructure using a variety of methods. This guide provides basic configuration steps for sending email alerts.

1. Navigate to *User settings > Profile*, switch to the *Media* tab and add your email.



2. Follow the guide for Receiving a problem notification.

Next time, when Zabbix detects a problem, you should receive an alert via email.

Test your configuration

To test your configuration, we can simulate a real problem by updating the host configuration in Zabbix frontend.

1. Open your website host configuration in Zabbix.

2. Switch to the *Macros* tab and select *Inherited and host macros*.

3. Click *Change* next to, for example, the previously configured {$WEBSITE.DOMAIN} macro value and set an incorrect domain name (for example, /git.zabbix.com/projects/ZBX/repos/zabbix/browse).

4. Click *Update* to update the host configuration.

5. In a few moments, Zabbix will detect the problem "Failed to get JSON of the requested website", because it will not be able to connect to the specified website. The problem will appear in *Monitoring > Problems*.



If alerts are configured, you will also receive the problem notification.

6. Change the macro value back to its previous value to resolve the problem and continue monitoring the website.

See also

- Creating an item - how to start monitoring additional metrics.
- Problem escalations - how to create multi-step alert scenarios (e.g., first send message to the system administrator, then, if a problem is not resolved in 45 minutes, send message to the data center manager).
- Browser items - how to configure Browser items.
- Template Website by Browser - additional information about the *Website by Browser* template.

Monitor website certificates with Zabbix agent 2 (passive)

Introduction

This guide provides a comprehensive overview of how to set up and monitor SSL/TLS certificates using the `web.certificate.get` key in Zabbix agent 2. It is designed to simplify certificate monitoring for single or multiple websites, ensuring administrators can quickly identify potential issues such as expired or invalid certificates.

Who this guide is for

This guide is designed for new Zabbix users and contains the minimum set of steps required to enable basic monitoring of website certificates. If you are looking for deep customization options or require more advanced configuration, see Configuration section of Zabbix manual.

Prerequisites

Before proceeding with this guide, you need to download and install Zabbix server, Zabbix frontend, and Zabbix agent 2 according to the instructions for your OS. This tutorial assumes that both the Zabbix server and agent are installed on the same machine; therefore, `127.0.0.1` is used in the configuration.

Configure Zabbix agent 2

1. Open the Zabbix agent configuration file:

```
sudo vi /etc/zabbix/zabbix_agent2.conf
```

2. Set the Server parameter to 127.0.0.1, since the agent and server are running on the same machine:

```
Server=127.0.0.1
```

3. Save the file and restart the Zabbix agent 2 service:

```
sudo systemctl restart zabbix-agent2
```

4. After setup and configuring Zabbix agent 2, test its availability with:

```
zabbix_get -s 127.0.0.1 -k web.certificate.get[<website_DNS_name>]
```

Zabbix agent 2 includes the WebCertificate plugin by default, so no separate installation or configuration is required.

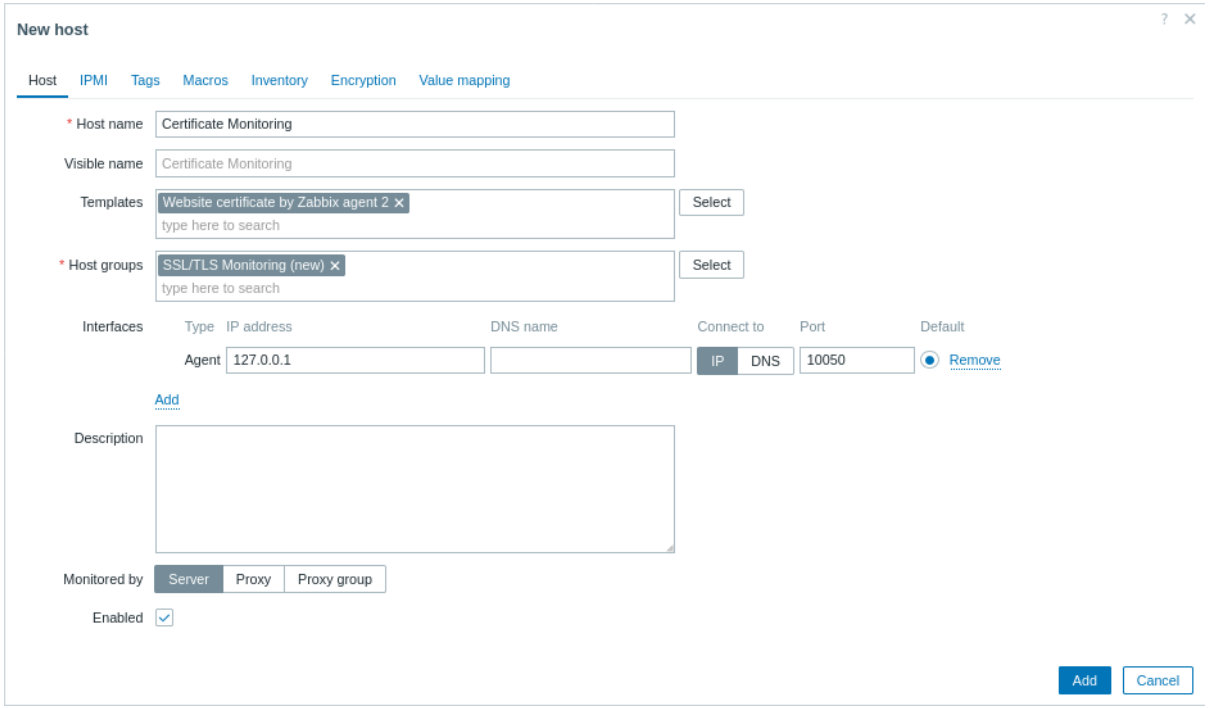
Configure Zabbix frontend

1. Log into Zabbix frontend.

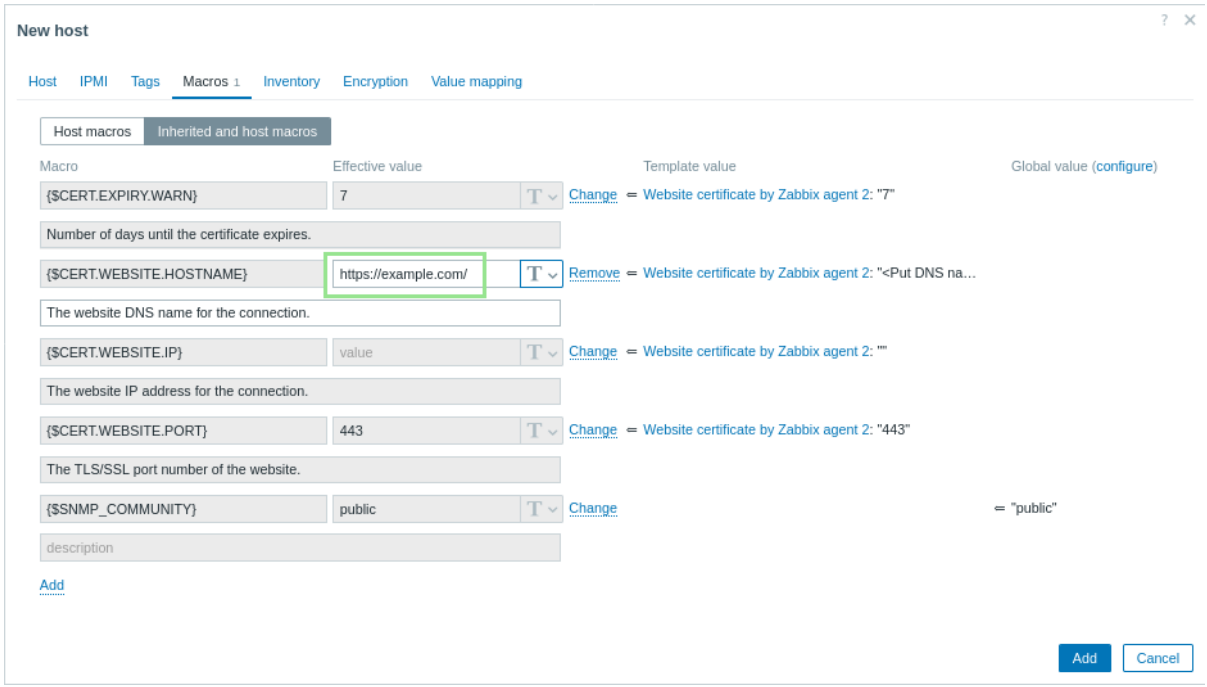
2. Navigate to *Monitoring* > *Hosts*.

3. Click an existing host where you want to monitor website certificates, or create a host if necessary:

- In the *Host name* field, enter a host name (e.g., ”Certificate Monitoring”).
- In the *Templates* field, type or select the template ”Website certificate by Zabbix agent 2” that will be linked to the host.
- In the *Host groups* field, type or select a host group (e.g., ”SSL/TLS Monitoring”).
- In the *Interfaces* field, add an interface of type ”Agent” and specify an IP address. For this example, ”127.0.0.1” is used.



- In the *Macros* tab, switch to *Inherited and host macros*, look for the following macros and click *Change* next to the macro value to update it:
  - `{$CERT.WEBSITE.HOSTNAME}` - enter the desired website DNS name as the value.



4. Click *Add* to add the host.

To monitor multiple websites, you can specify a comma-separated list of their hostnames in the `{$CERT.WEBSITE.HOSTNAME}` macro. Optionally, you can also provide a comma-separated list of ports in the `{$CERT.WEBSITE.PORT}` macro, where each port corresponds to the respective hostname in order. For example, to monitor example.com and example.org on port 8443:

- `{$CERT.WEBSITE.HOSTNAME}`: example.com,example.org
- `{$CERT.WEBSITE.PORT}`: 443,8443

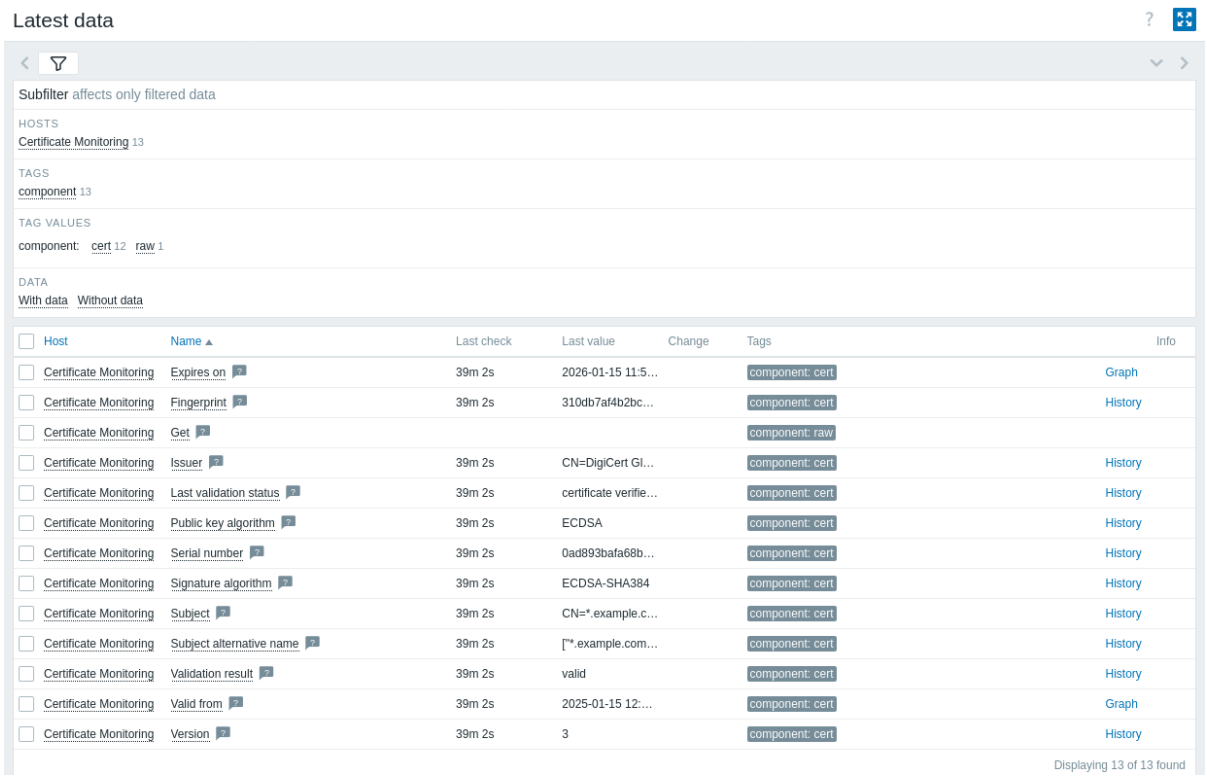
For each website specified in the `{$CERT.WEBSITE.HOSTNAME}` macro, Zabbix will create a corresponding set of items and triggers. This allows individual monitoring and alerting for each website's SSL certificate.

Host	Name ▲	Last check	Last value
Certificate Monitoring	Get data	11m 16s	{"error":{"message…
Certificate Monitoring	[example.com]: Expires on	11m 6s	2026-01-15 11:59:…
Certificate Monitoring	[example.com]: Fingerprint	11m 6s	310db7af4b2bc90…
Certificate Monitoring	[example.com]: Get		
Certificate Monitoring	[example.com]: Issuer	11m 6s	CN=DigiCert Glob…
Certificate Monitoring	[example.com]: Last validation status	11m 6s	certificate verified …
Certificate Monitoring	[example.com]: Public key algorithm	11m 6s	ECDSA
Certificate Monitoring	[example.com]: Serial number	11m 6s	0ad893bafa68b0b…
Certificate Monitoring	[example.com]: Signature algorithm	11m 6s	ECDSA-SHA384
Certificate Monitoring	[example.com]: Subject	11m 6s	CN=*.example.co…
Certificate Monitoring	[example.com]: Subject alternative name	11m 6s	["*.example.com","…
Certificate Monitoring	[example.com]: Validation result	11m 6s	valid
Certificate Monitoring	[example.com]: Valid from	11m 6s	2025-01-15 12:00:…
Certificate Monitoring	[example.com]: Version	11m 6s	3
Certificate Monitoring	[example.org]: Expires on	10m 52s	2026-01-15 11:59:…
Certificate Monitoring	[example.org]: Fingerprint	10m 52s	3b451cfce915637…
Certificate Monitoring	[example.org]: Get		
Certificate Monitoring	[example.org]: Issuer	10m 52s	CN=DigiCert Glob…
Certificate Monitoring	[example.org]: Last validation status	10m 52s	certificate verified …
Certificate Monitoring	[example.org]: Public key algorithm	10m 52s	ECDSA
Certificate Monitoring	[example.org]: Serial number	10m 52s	0722a749b558476…
Certificate Monitoring	[example.org]: Signature algorithm	10m 52s	ECDSA-SHA384
Certificate Monitoring	[example.org]: Subject	10m 52s	CN=*.example.org…
Certificate Monitoring	[example.org]: Subject alternative name	10m 52s	["*.example.org","e…
Certificate Monitoring	[example.org]: Validation result	10m 52s	valid
Certificate Monitoring	[example.org]: Valid from	10m 52s	2025-01-15 12:00:…

View collected metrics

Congratulations! At this point, Zabbix is already monitoring desired web certificate.

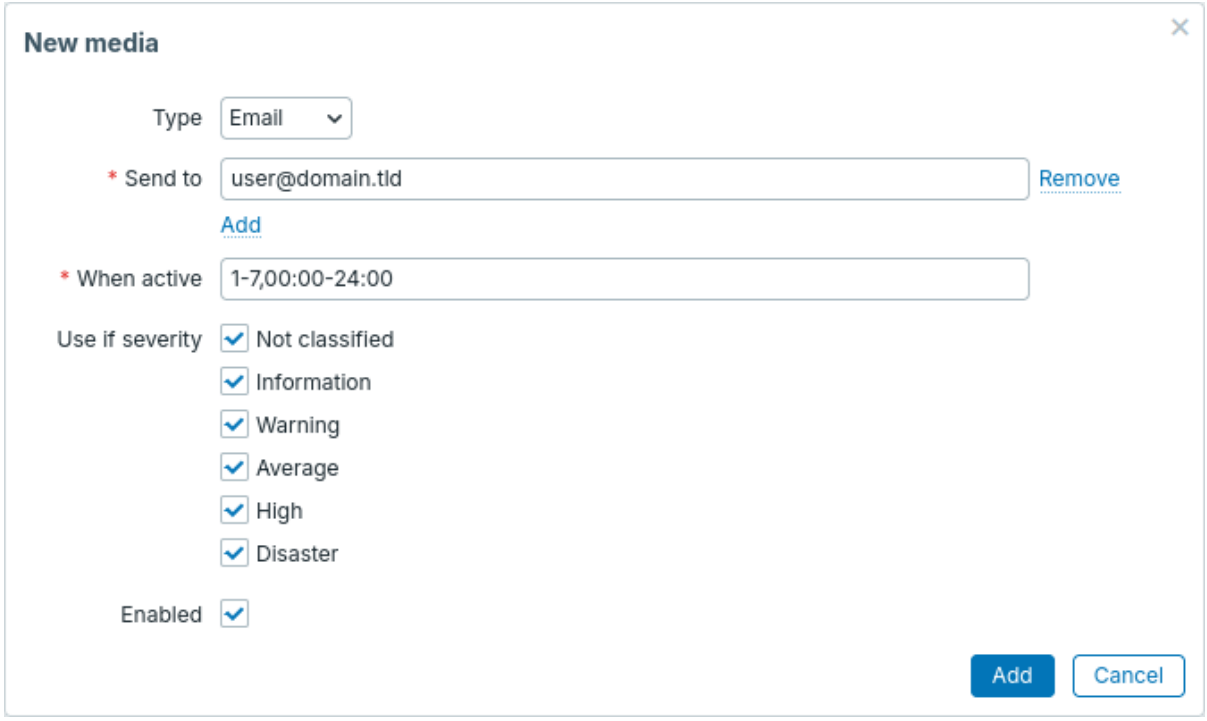
To view collected metrics, navigate to the *Monitoring > Hosts* menu section and click *Latest data* next to the host to view all the latest collected metrics in a list, like expiration date, issuer, and subject.



Set up problem alerts

Zabbix can notify you about infrastructure issues. This guide provides basic configuration steps for sending email alerts.

1. Navigate to *User settings > Profile*, switch to the *Media* tab and add your email.



2. Follow the guide for Receiving a problem notification.

Next time, when Zabbix detects a problem, you should receive an alert via email.

Test your configuration

To test your configuration, we can simulate a real problem by updating the host configuration in Zabbix frontend.

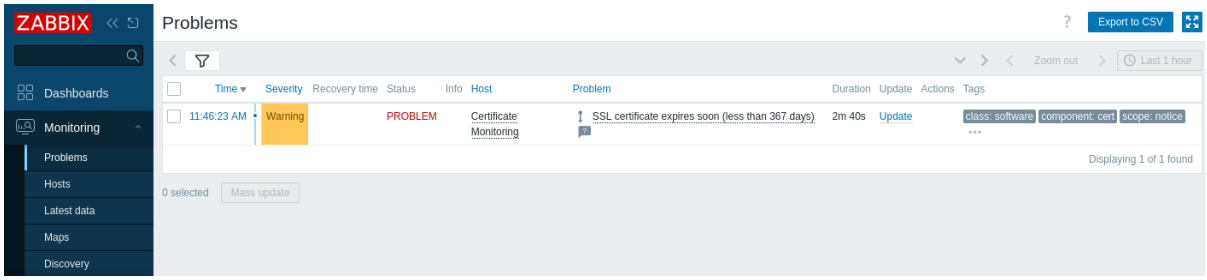
1. Open your ”Certificate Monitoring” host configuration in Zabbix.

2. Switch to the *Macros* tab and select *Inherited and host macros*.

3. Click *Change* next to the previously configured `{$CERT.EXPIRY.WARN}` macro value and set a very high number of days (greater than 365 days should be sufficient) to receive a warning before the certificate expires.

4. Click *Update* to update the host configuration.

5. In a few moments, Zabbix will detect the problem "SSL certificate expires soon", with the number of days until expiration. The problem will appear in *Monitoring > Problems*.



If alerts are configured, you will also receive the problem notification.

6. Change the macro value back to its previous value to resolve the problem and continue monitoring the certificate values.

See also

- Zabbix agent 2 - lists item keys.
- Template *Website certificate by Zabbix agent 2* - additional information about the *Website certificate by Zabbix agent 2* template.
- Template *Website certificate by Zabbix agent 2 active* - additional information about the *Website certificate by Zabbix agent 2 active* template.
- Monitor websites with Browser items - how to start basic monitoring of websites with Browser items.

Monitor a network switch or router with Zabbix

Introduction

This guide walks you through the steps required to start basic monitoring of your network switch or router using Zabbix. A Cisco router is used as an example, but the procedure applies to any SNMP-enabled network device.

Who this guide is for

This guide is designed for new Zabbix users and network administrators who want to quickly enable basic monitoring for network devices. If you require deep customization or advanced configuration options, please refer to the SNMP agent page or the Configuration section of Zabbix manual.

Prerequisites

Before proceeding with this guide, ensure that you have:

- Zabbix server and Zabbix frontend installed: install according to the instructions for your operating system (see Installation from packages and Web interface installation).
- Zabbix agent installed if monitoring local network metrics.
- SNMP-enabled device: a network switch or router (for example, a Cisco router) with SNMP enabled.
- MIB files installed: Installing MIB files enables Zabbix to translate numeric OIDs into human-readable names and descriptions. Without proper MIB support, you may see only numeric values, making it harder to configure items and troubleshoot issues.

To install MIB files on Ubuntu:

1. Install the MIB downloader package:

```
sudo apt-get update
sudo apt-get install snmp-mibs-downloader
```

If you need to add vendor-specific MIBs (e.g., from Cisco, Juniper), place them in the appropriate MIB directory:

- For Linux-based systems common locations include `/usr/share/snmp/mibs/` or `/usr/local/share/snmp/mibs/`.
- For Zabbix installations MIB files can be stored in `/var/lib/zabbix/mibs/`.

Ensure the `MIBDIRS` environment variable or the `snmp.conf` file includes the correct path.

To verify that your system recognizes the new MIBs, use:

```
snmptranslate -IR -On <MIB-NAME>::<object>
```

For detailed instructions, refer to your SNMP library documentation:

- Cisco MIBs
- Juniper MIBs

2. Edit `/etc/snmp/snmp.conf` and comment out the line that starts with `mibs :` to allow the system to load all available MIBs.

3. Verify by running an `snmpwalk` (for example, `snmpwalk -v 2c -c <your_community_string> <device_IP>`) and check that OIDs are displayed with descriptive names.

Depending on your environment, some steps in this guide may slightly differ. This guide is based on an environment running Ubuntu and a Cisco Catalyst 3750V2-24FS network device being monitored.

It is assumed that your network device is already physically installed and connected.

Configure the network device (Cisco router example)

For monitoring via SNMP, you must configure your network device to allow SNMP queries. The example below is for SNMPv2 and does not take into account existing settings. Caution: applying these commands may override current SNMP configurations.

For a Cisco router, the configuration typically involves steps written below.

SNMPv2 Example

1. Enable SNMP and set community string.

Log in to your Cisco router's console and enter configuration mode:

```
configure terminal
```

Then, enable SNMP by specifying a read-only community string. For example:

```
snmp-server community <your_community_string> RO
```

Replace `<your_community_string>` with your secure community string. Note: the RO (Read-Only) option allows SNMP to retrieve data from the device but prevents any configuration changes.

It is recommended to restrict SNMP access to only the necessary devices for security reasons. For further guidance on configuring access control lists (ACLs), refer to Cisco's official documentation.

2. Save the configuration.

Save your changes to ensure SNMP settings persist after a reboot:

```
write memory
```

SNMPv3 Example

SNMPv3 provides enhanced security with authentication and encryption. Its configuration is more secure than SNMPv2 and should be verified against your device-specific documentation.

1. Create an SNMP group.

Configure an SNMPv3 group with privacy (encryption) enabled:

```
configure terminal
snmp-server group <your_group> v3 priv
```

2. Create an SNMP user.

Add an SNMPv3 user with authentication and privacy. Replace the placeholders with your desired values:

```
snmp-server user <your_user> <your_group> v3 auth md5 <auth_password> priv aes 128 <priv_password>
```

3. Save the configuration:

```
write memory
```

For further details or model-specific instructions, you may refer to external Cisco SNMP configuration tutorials. This guide, however, provides the basic steps for enabling SNMP monitoring.

Configure Zabbix frontend

Create a host in Zabbix frontend

1. Log into Zabbix frontend.

2. Add a new host.

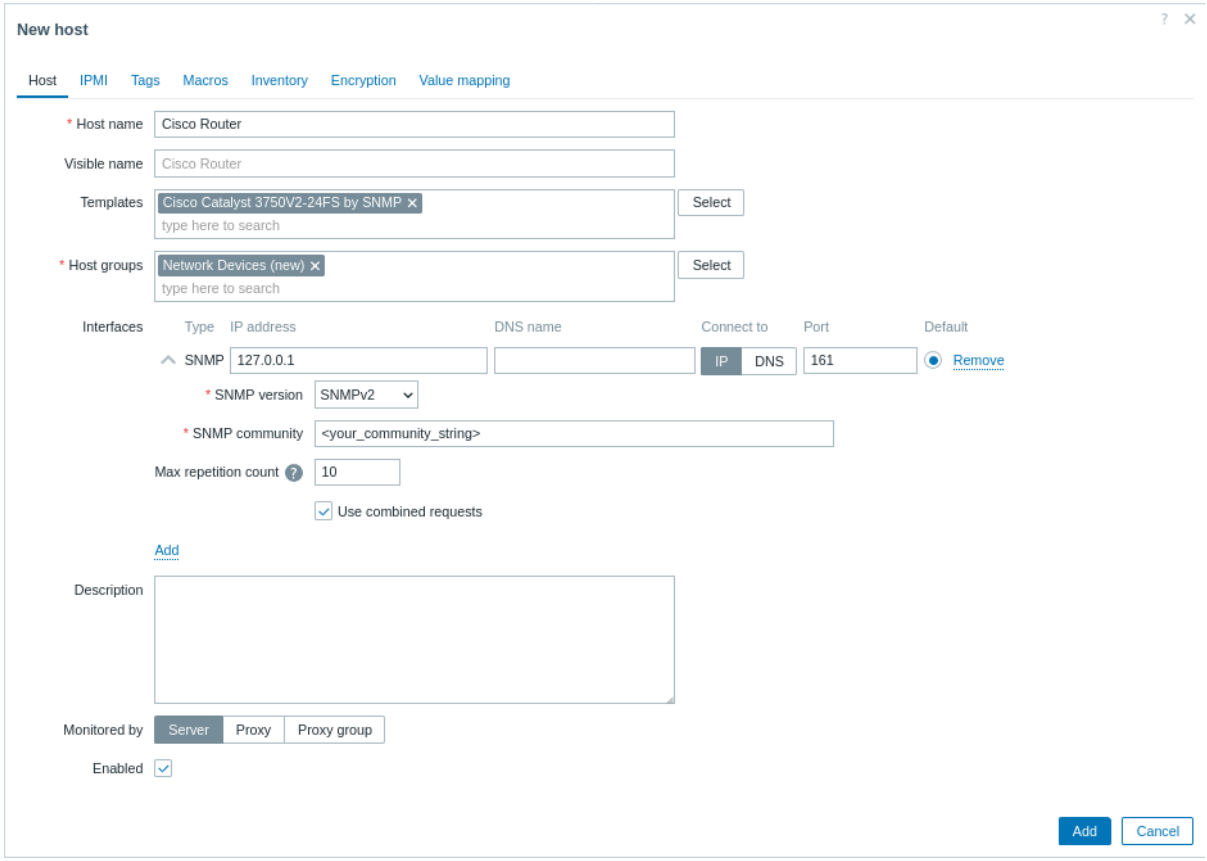
Navigate to *Data collection* > *Hosts* and click *Create host*.

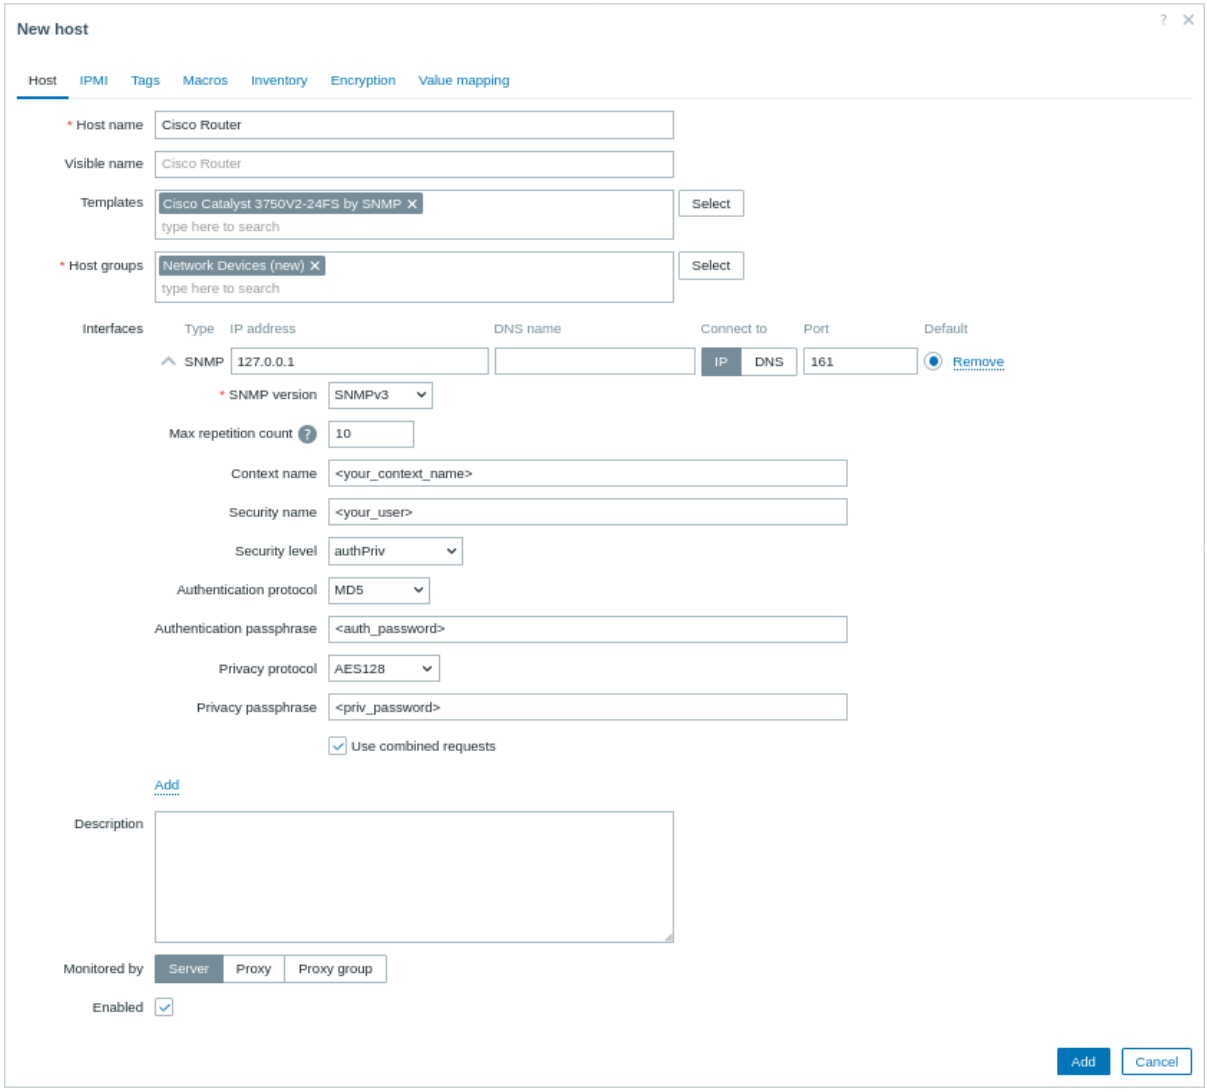
- *Host name*: enter a name for your device (e.g., ”Cisco Router”).
- Host groups: select an existing group or create a new group such as ”Network Devices”.
- Interfaces:
  - Click *Add* under Interfaces.
  - Choose *SNMP* as the interface type.
  - Enter the IP address or DNS name of your Cisco router.
  - Set the default SNMP port (usually 161).
  - Use the drop-down menu to select the appropriate SNMP version (e.g., SNMPv2).
  - For SNMPv1/v2, enter the community string in the *SNMP community* field. For SNMPv3, additional credentials (*Context name*, *Security name*, and *Security level*, etc.) will be prompted.

3. Link Templates

In the *Templates* field, select the SNMP template that best matches your device. Zabbix provides a range of pre-built SNMP templates for many device families. For example, if you are monitoring a Cisco device, choose the template that corresponds to your device’s OS or model (such as Cisco IOS SNMP or Cisco Catalyst 3750<device model> SNMP).

4. Click *Add* to save the host.



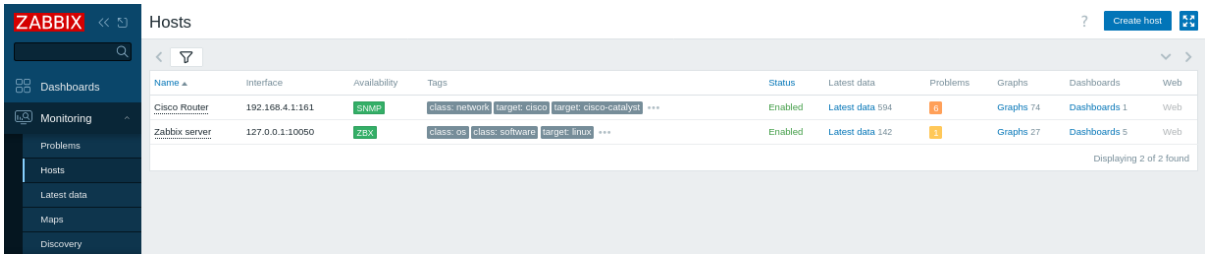


View collected metrics

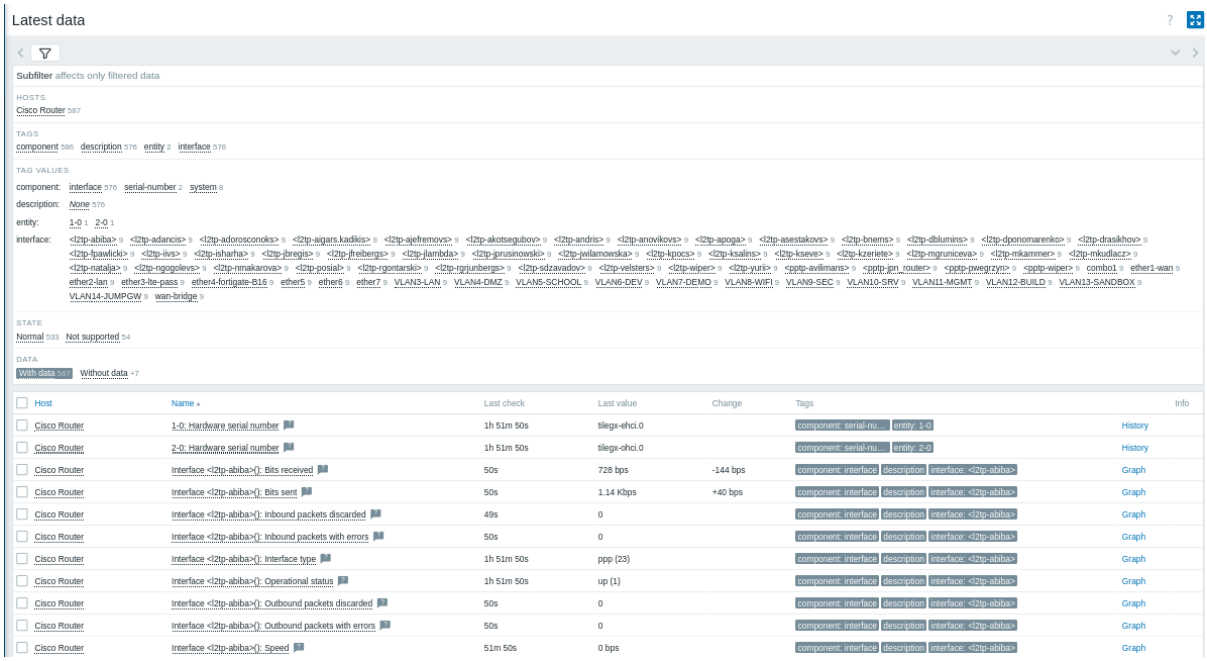
Congratulations! Zabbix is now set up to monitor your network device.

Latest Data:

- Navigate to Monitoring > Latest data in the Zabbix frontend.



- Select your ”Cisco Router” host (or discovered hosts) to view metrics such as hardware and network uptime, ICMP loss, ping, and response time, etc.



- Graphs and screens:

To visualize the performance data, click *Graphs* next to the SNMP items to see detailed metrics.

As a next step, you can:

- Add custom SNMP items to monitor additional metrics.
- Set up problem alerts to receive notifications about potential issues.

Create SNMP items

Once the host is set up, you can create items to monitor specific metrics. Note: this step is optional if you're using a template, as templates already contain default sets of items.

1. Identify the SNMP OID:

Use the `snmpwalk` command to list available OIDs on your device. For example:

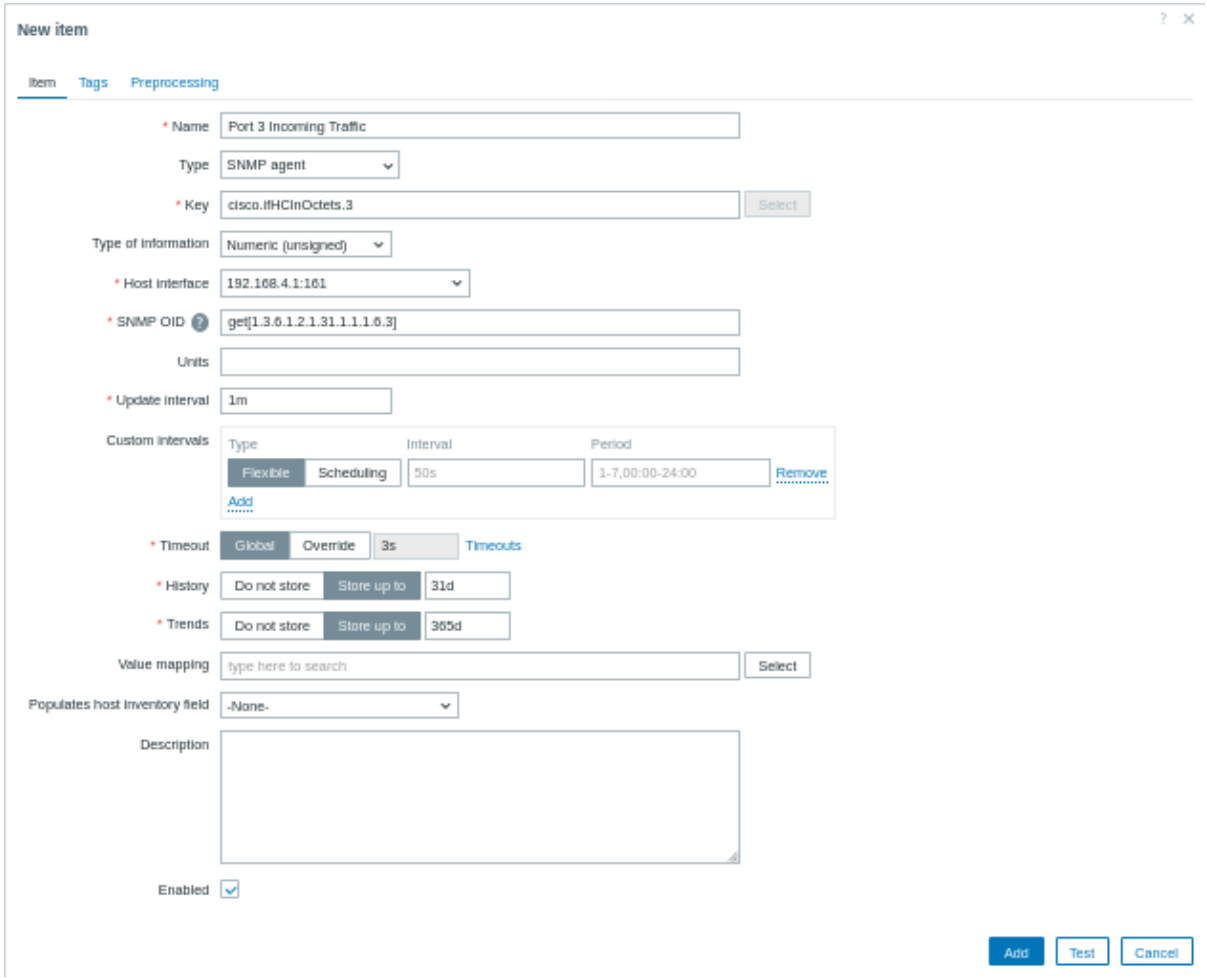
```
snmpwalk -v 2c -c <your_community_string> <device_IP> .
```

Find the OID for the metric you wish to monitor (for instance, IF-MIB::ifHCInOctets.3 for incoming traffic on port 3). To get the numeric OID, you can use:

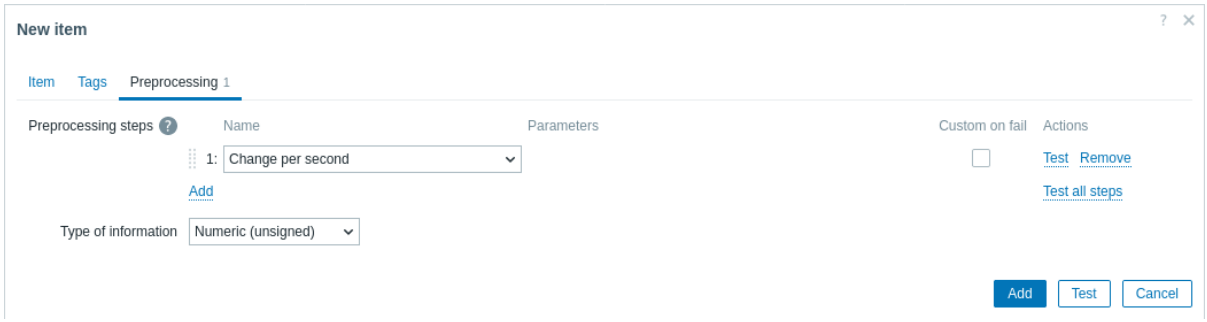
```
snmpget -v 2c -c <your_community_string> -On <device_IP> IF-MIB::ifHCInOctets.3
```

2. Create an SNMP item:

- Navigate to *Data collection > Hosts* and click the *Items* tab for your SNMP host and click *Create item*.
- *Name*: enter a descriptive name (e.g., "Port 3 Incoming Traffic").
- *Type*: select *SNMP agent*.
- *Key*: provide a meaningful key (e.g., `cisco.ifHCInOctets.3`).
- *Host interface*: ensure the SNMP interface is selected.
- *SNMP OID*: enter the OID using one of the supported formats, for example:
  - `get[1.3.6.1.2.1.31.1.1.1.6.3]` for a single value;
  - `walk[1.3.6.1.2.1.31.1.1.1.6.3]` to retrieve a subtree of values asynchronously.



- *Preprocessing* (if needed): if the item returns a cumulative counter (such as interface traffic), navigate to the *Preprocessing* tab, and add a preprocessing step like "Change per second" to calculate the rate.



To retrieve multiple values in one SNMP transaction, you can specify several OIDs using the syntax `walk[OID1,OID2,...]`.

Translating OIDs between Numeric and MIB Names

When working with SNMP, you might need to convert between numeric OIDs and their corresponding MIB names. This translation helps in identifying and troubleshooting metrics more easily.

- Translating a MIB name to a numeric OID: use the `snmptranslate` command with the `-On` option. For example, to translate the MIB name `IF-MIB::ifHCInOctets.3` to its numeric OID, run:

```
snmptranslate -On IF-MIB::ifHCInOctets.3
```

This command might output:

```
.1.3.6.1.2.1.31.1.1.1.6.3
```

- Translating a numeric OID to its MIB name: use the `snmptranslate` command with the `-IR` (or `-m ALL`) option to reverse the translation. For example, to translate the numeric OID `.1.3.6.1.2.1.31.1.1.1.6.3` back to its MIB name, run:

```
snmptranslate -IR -On .1.3.6.1.2.1.31.1.1.1.6.3
```

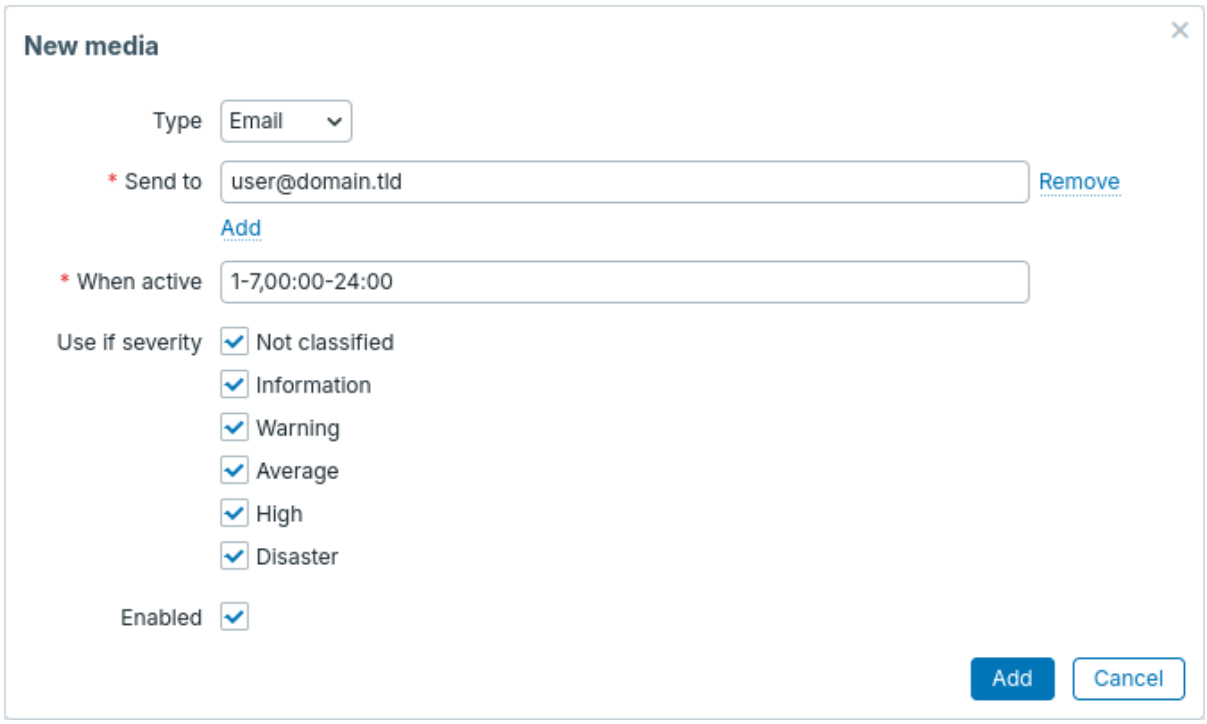
This command might output:

```
IF-MIB::ifHCInOctets.3
```

Set up problem alerts

This guide provides basic configuration steps for sending email alerts.

1. Navigate to *User settings* > *Profile*, switch to the *Media* tab and add your email.



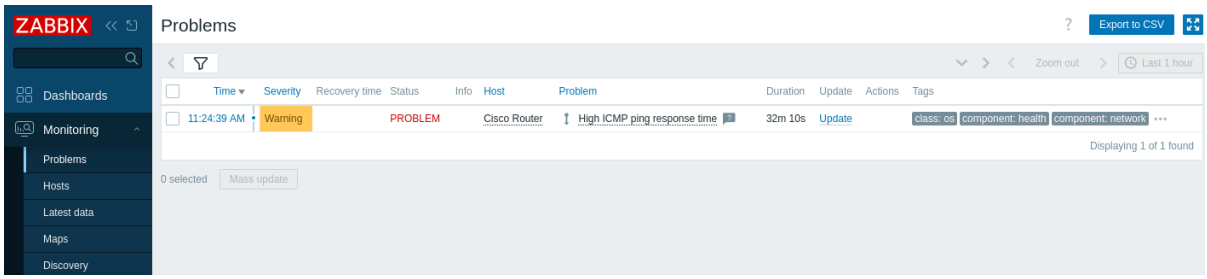
2. Follow the guide for Receiving a problem notification.

Next time, when Zabbix detects a problem, you should receive an alert via email.

Test your configuration

To ensure that Zabbix correctly detects network performance issues, simulate a real problem by increasing the ICMP ping response time threshold.

1. Open your "Cisco Router" host configuration in Zabbix.
2. Navigate to the *Macros* tab and select *Inherited and host macros*.
3. Locate the `{$ICMP_RESPONSE_TIME_WARN}` macro (or a similar response time threshold macro).
4. Set a very low value (e.g., 0.001) to trigger an alert when the ping response exceeds this value.
5. Click *Update* to apply the changes.
6. Wait a few moments for Zabbix to detect the simulated issue.
7. Navigate to *Monitoring* > *Problems* to verify that an alert appears (e.g., "High ICMP ping response time").



If alerts are configured, you should also receive a problem alert.

8. Revert the macro value to its original setting and click *Update* to save the changes.
9. Confirm that the problem is resolved and disappears from the *Problems* section.

Troubleshooting SNMP Monitoring

If you notice that the SNMP icon in the Zabbix frontend appears RED or no data is collected, try the following steps:

1. Check SNMP connectivity.

For SNMPv2 run the following command from your Zabbix server:

```
snmpwalk -v 2c -c <community_string> <device_IP> .
```

This command verifies that the device responds to SNMP queries.

For SNMPv3, include the appropriate SNMPv3 credentials:

```
snmpwalk -v3 -u <your_user> -l authPriv -a MD5 -A <auth_password> -x AES -X <priv_password> <device_IP> .
```

This verifies that SNMPv3 credentials are correct and the device is responding securely.

2. Ensure that MIB files are installed and enabled as described in the prerequisites. To ensure that, following command must not give error when you query a network device:

```
snmpwalk -v 2c -c <your_community_string> <device_IP> ifInOctets
```

This should return translated OIDs without errors.

3. Confirm that the SNMP version and credentials configured in Zabbix match those set on your device. For instance, review the SNMP settings in the Zabbix host configuration and verify them against your device's configuration. On a Cisco device, you might check the SNMP settings by running:

```
show running-config | include snmp
```

This ensures that the community string (for SNMPv2) or SNMPv3 user details are correct.

4. Verify that SNMP is correctly enabled on your network device. On a Cisco router, log in to the console and run:

```
show running-config | include snmp
```

This command displays the active SNMP configuration and helps confirm that SNMP is properly configured.

5. Ensure that no firewalls or network issues are blocking SNMP traffic (typically on port 161) between the Zabbix server and the device. You can test connectivity using:

nc -zv <device_IP> 161

`nc -zv` checks if port 161 is open and listening on the device.

Additionally, if you are using UFW on Ubuntu, check the firewall status:

```
sudo ufw status
```

Or, for iptables:

```
sudo iptables -L -n
```

6. Review the Zabbix server log files for any SNMP-related errors to help pinpoint the issue:

```
tail -f /tmp/zabbix_server.log
```

`tail -f` allows you to monitor log updates in real time.

See also:

- Creating an item - learn how to add additional metrics.
- SNMP agent - additional information on SNMP monitoring with Zabbix.
- Standardized templates for network devices - information on available SNMP templates.
- Discovery of SNMP OIDs - additional information on SNMP discovery on a switch.
- Configuring a network discovery rule - additional information on how to configure a network discovery rule used by Zabbix to discover hosts and services.

Monitor Windows event log using active checks

Introduction

This guide explains how to monitor Windows event logs with Zabbix using active checks. With Zabbix Windows-specific item keys, you can collect and analyze critical events (such as failed logon attempts, system errors, etc.) in real time.

Who this guide is for

This guide is designed for new Zabbix users and network administrators who want to monitor Windows event logs. For advanced configuration options, refer to the Windows-specific item keys documentation.

Prerequisites

Before proceeding with this guide, you need to download and install Zabbix server and Zabbix frontend according to the instructions for your OS. You also need Zabbix agent downloaded and installed on the Windows machine you want to monitor.

Configure Zabbix agent for Windows event log monitoring

1. Open the Zabbix agent configuration file on your Windows host and ensure that the `ServerActive` parameter is set to the IP address of your Zabbix server, and the `Hostname` parameter matches the host name that will be defined in Zabbix frontend. This enables the agent to request active checks for its host and from the specified Zabbix server. For example:

```
ServerActive=192.0.2.1
Hostname=MyWindowsHost
```

2. Restart Zabbix agent service to apply the changes:

```
net stop "Zabbix Agent" && net start "Zabbix Agent"
```

3. Check that the Windows host is running:

- Ensure Zabbix agent service is running on the Windows host.
- Check that the Windows host can connect to Zabbix server on port 10051. To test connectivity from the Windows host, open PowerShell and run the following command:

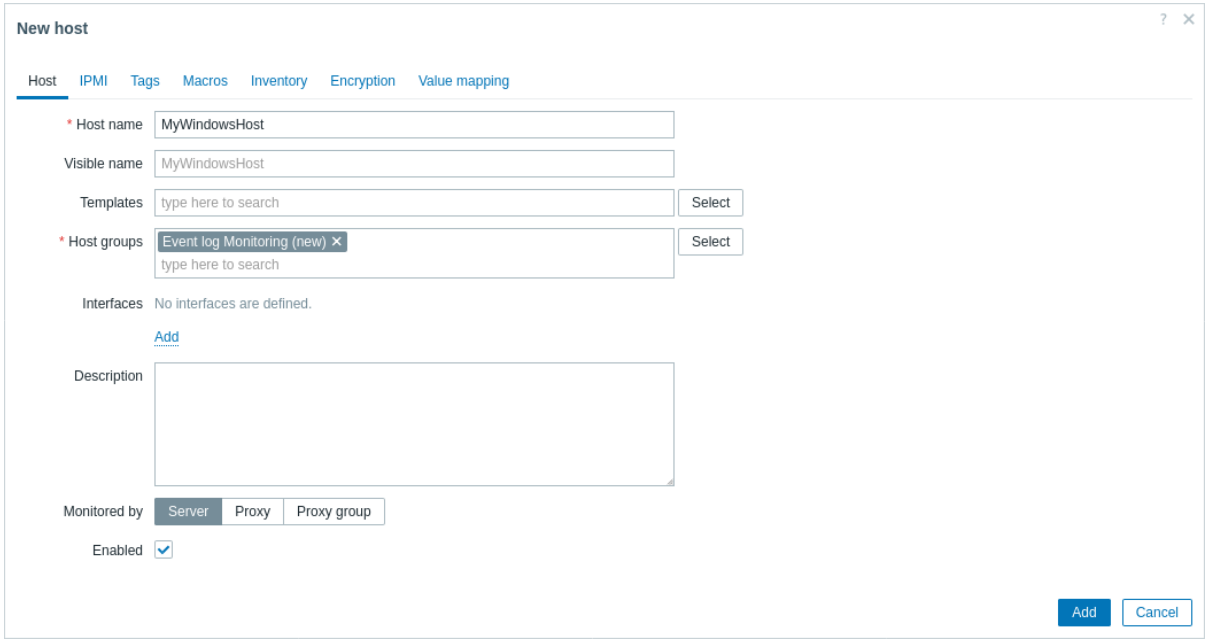
```
Test-NetConnection -ComputerName <Zabbix-server-IP> -Port 10051
```

Configure Zabbix frontend

1. Navigate to *Data collection* > *Hosts* and create a host:

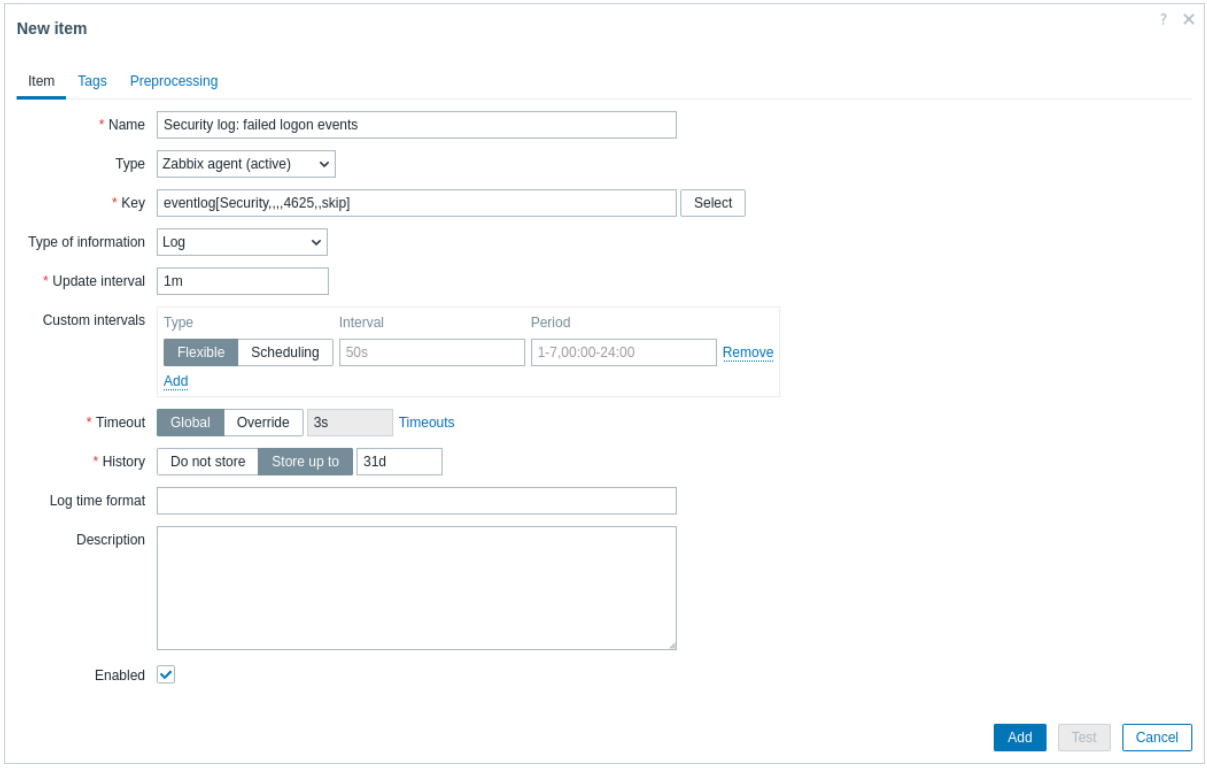
- In the *Host name* field, enter a host name (e.g., "MyWindowsHost").
- In the *Host groups* field, type or select a host group (e.g., "Event log Monitoring").
- Click *Add* to save the configured host.

Note:
In the *Templates* field you may add the "Windows by Zabbix agent active" template to help you troubleshoot by observing whether other active items on the same host are updating.



2. Create a new item with the following parameters:

- In the *Name* field, enter a descriptive item name (e.g., "Security log: failed logon events").
- In the *Type* drop-down, select "Zabbix agent (active)" (required for Event log monitoring).
- In the *Key* field, use the eventlog item key. For example, to monitor failed logon attempts (Event ID: 4625) in the Security log and to ignore entries older than the item's last check (using the `skip` parameter), enter the following item key: `eventlog[Security,,,,4625,,skip]`
- In the *Type of information* drop-down, select "Log".



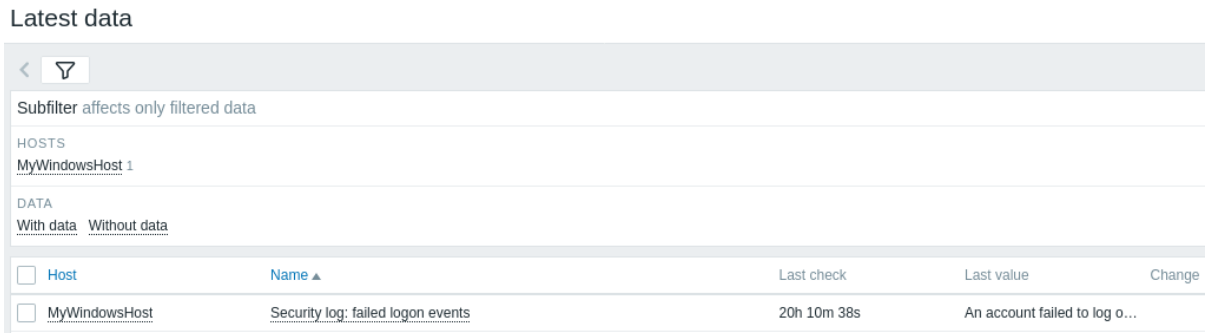
3. Click *Add* to save the item.

Test and view collected metrics

Congratulations! Zabbix is now set up to collect your Windows event logs. To verify that event logs are being collected, you can test the ”Security log: failed logon events” item by logging out of your Windows account and attempting to log in using incorrect credentials.

Then, view the collected logs in Zabbix frontend:

1. Navigate to *Monitoring* > *Latest data* in Zabbix frontend.



2. Filter by your ”MyWindowsHost” host in the *Name* field.

3. Click *History* to view recorded log values.

MyWindowsHost: Security log: failed logon events

Timestamp	Local time	Source	Severity	Event ID	Value
2025-04-07 11:13:11 PM	2025-04-07 11:12:19 PM	Microsoft-Windows-Security-Auditing	Failure Audit	4625	An account failed to log on. Subject: Security ID: NT AUTHORITY\SYSTEM Account Name: DESKTOP-NNEM26U$ Account Domain: WORKGROUP Logon ID: 0x3E7 Logon Type: 2 Account For Which Logon Failed: Security ID: NULL SID Account Name: Admin Account Domain: DESKTOP-NNEM26U Failure Information: Failure Reason: Unknown user name or bad password. Status: 0xC000006D Sub Status: 0xC000006A

4. If log values are absent, proceed to the Troubleshooting section of the guide.

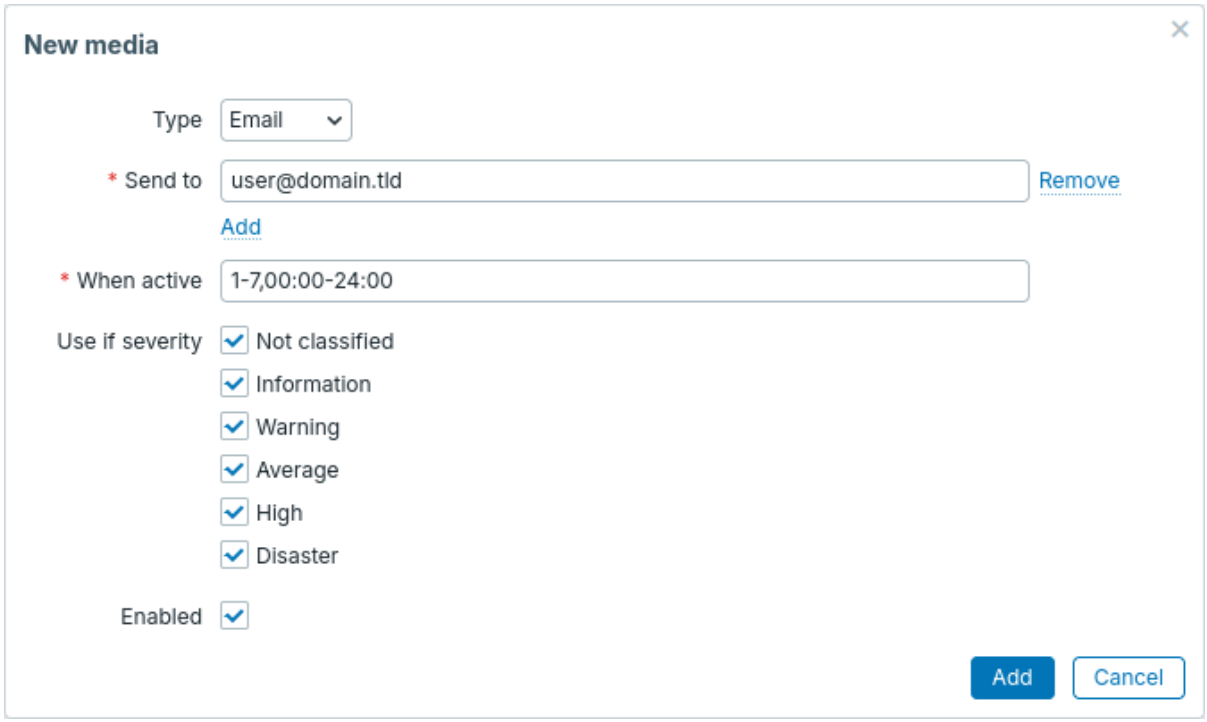
Set up problem alerts

This guide provides basic configuration steps for sending email alerts.

1. Navigate to Data collection > Hosts to define a trigger that fires when your event log item records the pattern you care about. For example, to catch failed logon attempts in the Security log, use the find() function:

find(/MyWindowsHost/eventlog[Security,,,4625,skip],10m,"like","Logon failed")

2. Navigate to *User settings > Profile*, switch to the *Media* tab and add your email.



3. Follow the guide for Receiving a problem notification.

Next time, when Zabbix detects a problem, you should receive an alert via email.

Troubleshooting

If you encounter issues with collecting or viewing Windows event logs, use the tips below to identify and resolve common problems:

1. On Zabbix server (Linux) list your iptables rules with the following command:

```
sudo iptables -L -n
```

and verify there's an ACCEPT rule for TCP port 10051.

2. Make sure your `eventlog[...]` key uses the exact log name (case-sensitive), event ID, mode (e.g., skip) and other parameters exactly as shown in the Windows-specific item keys.

See also:

- Creating an item - learn how to add additional metrics.
- Zabbix agent on Microsoft Windows - detailed installation instructions.
- Monitor Windows with Zabbix agent - a comprehensive guide on setting up basic monitoring for Windows machines using Zabbix agent.
- Windows-specific item keys - detailed information on Windows-specific item keys supported by Zabbix agents, including those for event log monitoring.
- Log file monitoring - instructions on configuring Zabbix for centralized monitoring and analysis of log files, applicable to Windows event logs.

Installation: additional guides

If you are looking for the easiest way of installing Zabbix, please see Installation and first steps.

The guides below cover **additional** or system-specific ways of installing Zabbix or its components:

- Installation from containers
- Installation from sources
- Notes on installation from packages
- Zabbix agent installation from Windows MSI
- Building Zabbix agent 2 on Windows
- Building Zabbix agent on Windows
- macOS agent installation from PKG
- Zabbix agent on IBM AIX
- Building Zabbix agent on macOS
- Zabbix appliance
- AWS deployment guide for Zabbix server
- AWS deployment guide for Zabbix proxy
- Azure deployment guide for Zabbix server
- Azure deployment guide for Zabbix proxy
- Google Cloud deployment guide

Installation from containers

Overview

This page describes how to deploy Zabbix using Docker or Docker Compose.

Docker Compose is the quickest way to deploy Zabbix. It reads a configuration file and automatically starts all the containers that make up a complete Zabbix setup, in the correct order.

Docker (manual deployment) achieves the same result by deploying each component step by step.

Prerequisites

Before you begin, make sure Docker (1.12.0 or later) is installed on your system. If it is not, follow the Docker installation guide. For Docker Compose deployment, Docker Compose (2.24.0 or later) is also required.

Some Zabbix components require specific ports to be open on the host running Docker (e.g., 10051/TCP for Zabbix server, 162/UDP for SNMP traps). See Requirements for a full list of ports used by Zabbix components. For Zabbix server and agent, the default port can be changed by setting the `ZBX_LISTENPORT` environment variable on the respective container.

Available Docker images

Zabbix provides a Docker image for each Zabbix component, all published on Docker Hub. Each image is used to create a container running that component.

Zabbix component	Docker image
Agent	zabbix/zabbix-agent
Agent 2	zabbix/zabbix-agent2
Server (MySQL)	zabbix/zabbix-server-mysql
Server (PostgreSQL)	zabbix/zabbix-server-pgsql
Web interface (Apache + MySQL)	zabbix/zabbix-web-apache-mysql
Web interface (Apache + PostgreSQL)	zabbix/zabbix-web-apache-pgsql
Web interface (Nginx + MySQL)	zabbix/zabbix-web-nginx-mysql
Web interface (Nginx + PostgreSQL)	zabbix/zabbix-web-nginx-pgsql
Proxy (SQLite3)	zabbix/zabbix-proxy-sqlite3
Proxy (MySQL)	zabbix/zabbix-proxy-mysql
Java gateway	zabbix/zabbix-java-gateway
Web service	zabbix/zabbix-web-service
SNMP traps	zabbix/zabbix-snmptraps
MCP server	zabbix/zabbix-mcp-server

Note:
To use SNMP traps, the SNMP traps container must share a volume with the Zabbix server or proxy container (see how to use this image on Docker Hub and the example below).

All Zabbix component images on Docker Hub are based on the latest major versions of the supported operating systems. These images are automatically rebuilt when the underlying OS images are updated.

Image tags

Each image supports tags to select the base operating system and Zabbix version:

```
zabbix/<image>:<os>-<version>
```

Supported `<os>` values:

- `alpine` - Alpine Linux
- `ubuntu` - Ubuntu
- `centos` - CentOS Stream
- `ol` - Oracle Linux
- `ltsc2022` - Windows 11 LTSC 2022 (Zabbix agent only)

Supported `<version>` values:

- `latest` - Latest stable Zabbix version on Alpine Linux
- `<os>-latest` - Latest stable Zabbix version on the selected OS
- `<os>-trunk` - Latest development (nightly) build on the selected OS
- `<os>-X.X-latest` - Latest Zabbix minor release of a specific Zabbix major version on the selected OS
- `<os>-X.X.*` - Specific Zabbix minor release on the selected OS

Examples:

```
#### Latest stable Zabbix proxy (SQLite3) on Alpine Linux:
docker pull zabbix/zabbix-proxy-sqlite3:latest

#### Latest stable Zabbix proxy (SQLite3) on Ubuntu:
docker pull zabbix/zabbix-proxy-sqlite3:ubuntu-latest

#### Latest development (nightly) build of Zabbix server (MySQL) on Ubuntu:
docker pull zabbix/zabbix-server-mysql:ubuntu-trunk

#### Latest 8.2 minor release of Zabbix server (MySQL) on Alpine Linux:
docker pull zabbix/zabbix-server-mysql:alpine-8.2-latest

#### Version 8.2.1 of Zabbix server (MySQL) on Alpine Linux:
docker pull zabbix/zabbix-server-mysql:alpine-8.2.1
```

Docker Compose

Docker Compose is the quickest way to deploy Zabbix. It reads a configuration file (Compose file) and handles the entire setup—pulling Docker images, creating an internal network between containers, setting up storage, initializing the database, and starting everything in the correct order.

The official Zabbix Dockerfiles repository provides ready-to-use Docker Compose files and an `.env`-based configuration system supporting different operating systems, database backends, and Zabbix component configuration.

1. Clone the repository, navigate to it, and switch to version 8.2:

```
git clone https://github.com/zabbix/zabbix-docker.git
cd zabbix-docker
git checkout 8.2
```

2. (Optional) Customize the deployment with environment variables. This step can be skipped for a default setup.

3. Deploy the default setup—Zabbix server, web interface (running on NGINX), and a database backend (MySQL or PostgreSQL), each running in its own container on Alpine Linux.

```
#### With MySQL as the database:
docker compose -f ./compose.yaml up -d

#### With PostgreSQL as the database:
docker compose -f ./compose_pgsql.yaml up -d
```

After the containers are up and running (typically within 1–3 minutes), the Zabbix server starts and the web interface becomes available for login.

Note:
Use `docker compose ps` to check container status. All containers (except the `zabbix-docker-server-db-init-1` container) should have the `Up` status. If any show `Exited`, check their logs for errors with `docker logs -f <container-name>`.

Environment variables

Docker Compose behavior and Zabbix component configuration can both be customized using environment variables.

Compose-level variables (defined in the `.env` file) control which Docker images, ports, and network IP ranges are used. You can use these variables inline before the `docker compose` command, or edit the `.env` file.

For example, the following command deploys a complete multi-container Zabbix setup using an Ubuntu-based image (instead of the default Alpine Linux) and exposing the web interface via NGINX on custom HTTP (8282) and HTTPS (8443) ports:

```
OS=ubuntu \
ZABBIX_WEB_NGINX_HTTP_PORT=8282 \
ZABBIX_WEB_NGINX_HTTPS_PORT=8443 \
docker compose -f ./compose.yaml up -d
```

Component-level variables (defined in the `env_vars/.env_<component>` files) control Zabbix component configuration. Edit the respective `.env_<component>` file before running the `docker compose` command.

For example, you can adjust Zabbix server passive poller count and cache settings, as well as configure the web interface timezone by editing the following variables:

```
#### env_vars/.env_srv
ZBX_STARTPOLLERS=20
ZBX_CACHESIZE=64M
ZBX_HISTORYCACHESIZE=32M

#### env_vars/.env_web
PHP_TZ=Europe/Paris
```

Note:
Component-level environment variables correspond to Zabbix component configuration parameters using a different naming style (e.g., `ZBX_STARTPOLLERS` corresponds to `StartPollers`). Some variables are Docker-specific, and some configuration parameters cannot be changed (e.g., `PIDFile` and `LogType`). When using environment variables, please refer to the *Environment variables* section in each component's Docker image description on Docker Hub.

Volumes

Docker Compose stores persistent data in the `zbx_env/` directory created alongside the Compose file. This directory preserves component data across container restarts and updates.

The contents of `zbx_env/` are pre-defined for each component's image. For example:

- `zbx_env/etc/` - used by the Zabbix web interface container
- `zbx_env/usr/` and `zbx_env/var/` - used by the Zabbix server container

For more information on each volume, see the *Allowed volumes* section in each component's Docker image description on Docker Hub.

Makefile

The Zabbix Dockerfiles repository also provides a Makefile as a shortcut for common Docker Compose tasks. Instead of using the full `docker compose` command, you can use shorter `make` commands (run `make help` to see the available options):

```
#### Default deployment (Zabbix server, web interface, MySQL, all on Alpine Linux)
make up

#### Custom deployment (Zabbix server, web interface with custom NGINX ports, PostgreSQL, all on Ubuntu)
make up \
  OS=ubuntu \
  DB=pgsql \
  ZABBIX_WEB_NGINX_HTTP_PORT=8282 \
  ZABBIX_WEB_NGINX_HTTPS_PORT=8443
```

Attention:
When stopping or removing containers, always specify the same database type that was deployed (e.g., `make down DB=pgsql`).

By default, `make up` starts only a minimal set of services (Zabbix server, web interface, and database). This keeps the setup lightweight and avoids starting components that may not be needed. To include additional components, you can use the following Compose profiles:

```
#### Default deployment + Zabbix agent, Java gateway, web service, and SNMP traps:
make up COMPOSE_PROFILES=full

#### COMPOSE_PROFILES=full + Zabbix proxy (MySQL and SQLite3):
make up COMPOSE_PROFILES=all
```

Docker (manual deployment)

Use manual deployment when you want to deploy Zabbix step by step, run individual components, integrate them with an existing environment, or use an alternative container runtime such as Podman.

For example, to deploy a Zabbix proxy, run the following command:

```
docker run --name zabbix-proxy-sqlite3 \
  -e ZBX_SERVER_HOST=192.0.2.1 \
  -e ZBX_PROXYMODE=0 \
  -e ZBX_HOSTNAME=zabbix-proxy-sqlite3 \
  -v zabbix-proxy-data:/var/lib/zabbix/db_data \
  --init \
  -d \
  zabbix/zabbix-proxy-sqlite3:alpine-8.2-latest
```

This command:

- Pulls the `zabbix/zabbix-proxy-sqlite3:alpine-8.2-latest` Docker image.
- Creates and starts a `zabbix-proxy-sqlite3` container based on the Docker image (with `--init` and `-d` flags).
- Configures the Zabbix proxy `Server` parameter via the `ZBX_SERVER_HOST` environment variable. This parameter specifies the IP address of the Zabbix server from which the proxy will retrieve configuration data and to which it will send collected data. Note that other variables required for active proxy operation use default values and can be omitted.

Note:
Environment variables correspond to Zabbix component configuration parameters using a different naming style (e.g., `ZBX_SERVER_HOST` corresponds to `Server`). Some variables are Docker-specific, and some configuration parameters cannot be changed (e.g., `PIDFile` and `LogType`). When using environment variables, please refer to the *Environment variables* section in each component's Docker image description on Docker Hub.

- Connects a storage location managed by Docker (e.g., `/var/lib/docker/volumes/zabbix-proxy-data/`) to the container's `/var/lib/zabbix` directory so that Zabbix proxy data is stored permanently, even if the container is removed.

Note:
For more information on each volume, see the *Allowed volumes* section in each component's Docker image description on Docker Hub.

After deploying the Zabbix proxy container, you can proceed to add the proxy in your Zabbix web interface and configure hosts to be monitored by this proxy.

The examples below cover three additional deployment scenarios:

- Zabbix server (MySQL) with Java gateway
- Zabbix server (PostgreSQL) with SNMP traps
- Zabbix server (MySQL) with Java gateway on RHEL 8–10

For more examples, please refer to each component's Docker image description on Docker Hub.

Deploy Zabbix server (MySQL) with Java gateway

This example shows how to deploy a Zabbix server with a MySQL backend, a Java gateway for JMX monitoring, and an NGINX-based web interface.

1. Create a dedicated Docker network, so that all Zabbix component containers can reach each other by container name:

```
docker network create --subnet 172.20.0.0/16 --ip-range 172.20.240.0/20 zabbix-net
```

2. Run an empty MySQL server container (replace `zabbix_pwd` and `root_pwd` with strong passwords, and use those values in the following steps):

```
docker run --name mysql-server -t \
  -e MYSQL_DATABASE="zabbix" \
  -e MYSQL_USER="zabbix" \
  -e MYSQL_PASSWORD="zabbix_pwd" \
  -e MYSQL_ROOT_PASSWORD="root_pwd" \
  --network=zabbix-net \
  --restart unless-stopped \
  -d mysql:8.4-oracle \
  --character-set-server=utf8mb4 --collation-server=utf8mb4_bin
```

3. Follow the MySQL logs and wait until MySQL is ready for connections (use `Ctrl+C` to exit the logs):

```
docker logs -f mysql-server

#### ... [Server] /usr/sbin/mysqld: ready for connections. ...
```

4. Run a Zabbix Java gateway container:

```
docker run --name zabbix-java-gateway -t \
  --network=zabbix-net \
  --restart unless-stopped \
  -d zabbix/zabbix-java-gateway:alpine-8.2-latest
```

5. Enable `log_bin_trust_function_creators` in the MySQL server container. This temporarily relaxes security restrictions on stored functions; the setting will be disabled in later steps. It allows the Zabbix database user to create stored functions without the `SUPER` privilege, which MySQL requires when binary logging is enabled (default since MySQL 8.2). Run the following command:

```
docker exec -it mysql-server mysql -u root -p'root_pwd' \
  -e "SET GLOBAL log_bin_trust_function_creators = 1;"
```

6. Run a Zabbix server container and link it to the Java Gateway and MySQL server containers:

```
docker run --name zabbix-server-mysql -t \
  -e DB_SERVER_HOST="mysql-server" \
  -e MYSQL_DATABASE="zabbix" \
  -e MYSQL_USER="zabbix" \
  -e MYSQL_PASSWORD="zabbix_pwd" \
  -e ZBX_JAVAGATEWAY="zabbix-java-gateway" \
  --network=zabbix-net \
  -p 10051:10051 \
  --restart unless-stopped \
  --init \
  -d zabbix/zabbix-server-mysql:alpine-8.2-latest
```

7. Follow the Zabbix server logs and wait until Zabbix server has finished initializing the database schema (use `Ctrl+C` to exit the logs):

```
docker logs -f zabbix-server-mysql

#### ... [info]: ** Creating 'zabbix' schema in MySQL
#### ... [info]: ** Database schema successfully created!
```

8. Disable `log_bin_trust_function_creators` in the MySQL server container. This re-enforces the security restriction that prevents non-SUPER users from creating stored functions. Run the following command:

```
docker exec -it mysql-server mysql -u root -p'root_pwd' \
  -e "SET GLOBAL log_bin_trust_function_creators = 0;"
```

9. Run a Zabbix web interface container and link it to the Zabbix server and MySQL server containers:

```
docker run --name zabbix-web-nginx-mysql -t \
  -e ZBX_SERVER_HOST="zabbix-server-mysql" \
  -e DB_SERVER_HOST="mysql-server" \
  -e MYSQL_DATABASE="zabbix" \
  -e MYSQL_USER="zabbix" \
  -e MYSQL_PASSWORD="zabbix_pwd" \
  -e PHP_TZ="Europe/Riga" \
  --network=zabbix-net \
  -p 80:8080 \
  --restart unless-stopped \
  -d zabbix/zabbix-web-nginx-mysql:alpine-8.2-latest
```

After the containers are up and running, the Zabbix server starts and the web interface becomes available for login.

Note:
Use `docker ps` to check container status. All containers should have the `Up` status. If any show `Exited`, check their logs for errors with `docker logs -f <container-name>`.

Deploy Zabbix server (PostgreSQL) with SNMP traps

This example shows how to deploy a Zabbix server with a PostgreSQL backend, SNMP traps, and an NGINX-based web interface.

1. Create a dedicated Docker network, so that all Zabbix component containers can reach each other by container name:

```
docker network create --subnet 172.20.0.0/16 --ip-range 172.20.240.0/20 zabbix-net
```

2. Run an empty PostgreSQL server container (replace `zabbix_pwd` with a strong password, and use that value in the following steps):

```
docker run --name postgres-server -t \
  -e POSTGRES_USER="zabbix" \
  -e POSTGRES_PASSWORD="zabbix_pwd" \
  -e POSTGRES_DB="zabbix" \
  --network=zabbix-net \
  --restart unless-stopped \
  -d postgres:latest
```

3. Run a Zabbix SNMP traps container:

```
docker run --name zabbix-snmptraps -t \
  -v /zbx_instance/snmptraps:/var/lib/zabbix/snmptraps:rw \
  -v /var/lib/zabbix/mibs:/usr/share/snmp/mibs:ro \
  --network=zabbix-net \
  -p 162:1162/udp \
  --restart unless-stopped \
  -d zabbix/zabbix-snmptraps:alpine-8.2-latest
```

4. Run a Zabbix server container and link it to the PostgreSQL server and SNMP traps containers:

```
docker run --name zabbix-server-pgsql -t \
  -e DB_SERVER_HOST="postgres-server" \
  -e POSTGRES_USER="zabbix" \
  -e POSTGRES_PASSWORD="zabbix_pwd" \
  -e POSTGRES_DB="zabbix" \
  -e ZBX_ENABLE_SNMP_TRAPS="true" \
  --network=zabbix-net \
  -p 10051:10051 \
  --volumes-from zabbix-snmptraps \
  --restart unless-stopped \
  --init \
  -d zabbix/zabbix-server-pgsql:alpine-8.2-latest
```

5. Follow the Zabbix server logs and wait until Zabbix server has finished initializing the database schema (use `Ctrl+C` to exit the logs):

```
docker logs -f zabbix-server-pgsql

#### ... [info]: ** Creating 'zabbix' schema in PostgreSQL
#### ... [info]: ** Database schema successfully created!
```

6. Run a Zabbix web interface container and link it to the Zabbix server and PostgreSQL server containers:

```
docker run --name zabbix-web-nginx-pgsql -t \
  -e ZBX_SERVER_HOST="zabbix-server-pgsql" \
  -e DB_SERVER_HOST="postgres-server" \
  -e POSTGRES_DB="zabbix" \
  -e POSTGRES_USER="zabbix" \
  -e POSTGRES_PASSWORD="zabbix_pwd" \
  -e PHP_TZ="Europe/Riga" \
  --network=zabbix-net \
  -p 443:8443 \
  -p 80:8080 \
  -v /etc/ssl/nginx:/etc/ssl/nginx:ro \
  --restart unless-stopped \
  -d zabbix/zabbix-web-nginx-pgsql:alpine-8.2-latest
```

After the containers are up and running, the Zabbix server starts and the web interface becomes available for login.

> **Note:**
> Use `docker ps` to check container status. All containers should have the `Up` status. If any show `Exited`, check their logs for errors with `docker logs -f <container-name>`.

Deploy Zabbix server (MySQL) with Java gateway on RHEL 8-10

This example shows how to deploy a Zabbix server with a MySQL backend, a Java gateway for JMX monitoring, and an NGINX-based web interface, all running on Red Hat Enterprise Linux 8, 9, or 10.

On Red Hat Enterprise Linux, the recommended container runtime is Podman instead of Docker. Podman works similarly to Docker but does not require a background service running as root, which makes it a better fit for Red Hat environments.

1. Create a new pod with the name `zabbix` and exposed ports for Zabbix web interface and Zabbix server trapper:

```
podman pod create --name zabbix -p 80:8080 -p 10051:10051
```

2. (Optional) Run a Zabbix agent container in the `zabbix` pod location:

```
podman run --name zabbix-agent \
  -e ZBX_SERVER_HOST="127.0.0.1,localhost" \
  --restart=always \
  --pod=zabbix \
  -d registry.connect.redhat.com/zabbix/zabbix-agent-80:latest
```

3. Create a `./mysql/` directory on the RHEL host:

```
mkdir -p ./mysql
```

4. Run an empty MySQL server container (replace `zabbix_pwd` and `root_pwd` with strong passwords, and use those values in the following steps):

```
podman run --name mysql-server -t \
  -e MYSQL_DATABASE="zabbix" \
  -e MYSQL_USER="zabbix" \
  -e MYSQL_PASSWORD="zabbix_pwd" \
  -e MYSQL_ROOT_PASSWORD="root_pwd" \
  -v ./mysql/:/var/lib/mysql/:Z \
  --restart=always \
  --pod=zabbix \
  -d mysql:8.4 \
  --character-set-server=utf8mb4 --collation-server=utf8mb4_bin
```

5. Follow the MySQL logs and wait until MySQL is ready for connections (use `Ctrl+C` to exit the logs):

```
podman logs -f mysql-server

#### ... [Server] /usr/sbin/mysqld: ready for connections. ...
```

6. Run a Zabbix Java gateway container:

```
podman run --name zabbix-java-gateway -t \
  --restart=always \
  --pod=zabbix \
  -d registry.connect.redhat.com/zabbix/zabbix-java-gateway-80
```

7. Enable `log_bin_trust_function_creators` in the MySQL server container. This temporarily relaxes security restrictions on stored functions; the setting will be disabled in later steps. It allows the Zabbix database user to create stored functions without the `SUPER` privilege, which MySQL requires when binary logging is enabled (default since MySQL 8.0). Run the following command:

```
podman exec -it mysql-server mysql -u root -p'root_pwd' \
  -e "SET GLOBAL log_bin_trust_function_creators = 1;"
```

8. Run a Zabbix server container and link it to the Java Gateway and MySQL server containers:

```
podman run --name zabbix-server-mysql -t \
  -e DB_SERVER_HOST="127.0.0.1" \
  -e MYSQL_DATABASE="zabbix" \
  -e MYSQL_USER="zabbix" \
  -e MYSQL_PASSWORD="zabbix_pwd" \
  -e ZBX_JAVAGATEWAY="127.0.0.1" \
  --restart=always \
  --init \
  --pod=zabbix \
  -d registry.connect.redhat.com/zabbix/zabbix-server-mysql-80
```

9. Follow the Zabbix server logs and wait until Zabbix server has finished initializing the database schema (use `Ctrl+C` to exit the logs):

```
podman logs -f zabbix-server-mysql

#### ... [info]: ** Creating 'zabbix' schema in MySQL
#### ... [info]: ** Database schema successfully created!
```

10. Disable `log_bin_trust_function_creators` in the MySQL server container. This re-enforces the security restriction that prevents non-SUPER users from creating stored functions. Run the following command:

```
podman exec -it mysql-server mysql -u root -p'root_pwd' \
  -e "SET GLOBAL log_bin_trust_function_creators = 0;"
```

11. Run a Zabbix web interface container and link it to the Zabbix server and MySQL server containers:

```
podman run --name zabbix-web-mysql -t \
  -e ZBX_SERVER_HOST="127.0.0.1" \
  -e DB_SERVER_HOST="127.0.0.1" \
  -e MYSQL_DATABASE="zabbix" \
  -e MYSQL_USER="zabbix" \
  -e MYSQL_PASSWORD="zabbix_pwd" \
  -e PHP_TZ="Europe/Riga" \
  --restart=always \
  --pod=zabbix \
  -d registry.connect.redhat.com/zabbix/zabbix-web-mysql-80
```

Note:
Pod `zabbix` exposes 80/TCP port (HTTP) to the host machine from 8080/TCP of `zabbix-web-mysql` container.

After the containers are up and running, the Zabbix server starts and the web interface becomes available for login.

Note:
Use `docker ps` to check container status. All containers should have the `Up` status. If any show `Exited`, check their logs for errors with `docker logs -f <container-name>`.

Installation from sources

You can get the very latest version of Zabbix by compiling it from the sources. See also Getting Zabbix source code.

A step-by-step tutorial for installing Zabbix from the sources is provided here.

Installing Zabbix daemons

1 Download the source archive

Go to the Zabbix download page and download the source archive. Once downloaded, extract the sources, by running:

```
tar -zxvf zabbix-8.2.0.tar.gz
```

Note:
Enter the correct Zabbix version in the command. It must match the name of the downloaded archive.

2 Create user account

All Zabbix daemon processes run under unprivileged system users. If a Zabbix daemon is started from an unprivileged user account, it will continue running as that user.

In the default configuration, if a daemon is started as `root`, it will switch to the `zabbix` user account, which must be present. To create a `zabbix` system user and group, run the commands listed below.

RedHat-based system:

```
groupadd --system zabbix
useradd --system -g zabbix -d /usr/lib/zabbix -s /sbin/nologin -c "Zabbix Monitoring System" zabbix
```

Debian-based system:

```
addgroup --system --quiet zabbix
adduser --quiet --system --disabled-login --ingroup zabbix --home /var/lib/zabbix --no-create-home zabbix
```

There is no need to create a separate user account for Zabbix frontend.

Security recommendation

If the Zabbix server and agent run on the same machine, it is recommended to run them under **separate user accounts**. Running both as the same user allows the agent to access the server's configuration file, which could expose sensitive information—such as the database password—to any Admin-level user in Zabbix.

Attention:
Running Zabbix as `root`, `bin`, or any other account with special rights is a security risk.

Home directory (optional)

Zabbix processes do not require a home directory, so creating one is generally not recommended. However, if you need functionality that requires a home directory (for example, storing MySQL credentials in $HOME/.my.cnf), you can create it using the commands listed below.

On RedHat-based systems, run:

```
mkdir -m u=rwx,g=rwx,o= -p /usr/lib/zabbix
chown zabbix:zabbix /usr/lib/zabbix
```

On Debian-based systems, run:

```
mkdir -m u=rwx,g=rwx,o= -p /var/lib/zabbix
chown zabbix:zabbix /var/lib/zabbix
```

3 Create Zabbix database

For Zabbix server and proxy daemons, as well as Zabbix frontend, a database is required. It is not needed to run Zabbix agent.

SQL scripts are provided for creating database schema and inserting the dataset. Zabbix proxy database needs only the schema while Zabbix server database requires also the dataset on top of the schema.

Having created a Zabbix database, proceed to the following steps of compiling Zabbix.

4 Configure the sources

C99 with GNU extensions is required for building Zabbix server, Zabbix proxy or Zabbix agent. This version can be explicitly specified by setting CFLAGS="-std=gnu99":

```
export CFLAGS="-std=gnu99"
```

Note:
If installing from Zabbix Git repository, it is required to run first:
`./bootstrap.sh`

When configuring the sources for a Zabbix server or proxy, you must specify the database type to be used. Only one database type can be compiled with a server or proxy process at a time.

To see all of the supported configuration options, inside the extracted Zabbix source directory run:

```
./configure --help
```

To configure the sources for a Zabbix server and agent, you may run something like:

```
./configure --enable-server --enable-agent --with-mysql --enable-ipv6 --with-net-snmp --with-libcurl --wit
```

To configure the sources for a Zabbix server (with PostgreSQL etc.), you may run:

```
./configure --enable-server --with-postgresql --with-net-snmp
```

To configure the sources for a Zabbix proxy (with SQLite etc.), you may run:

```
./configure --prefix=/usr --enable-proxy --with-net-snmp --with-sqlite3 --with-ssh2
```

To configure the sources for a Zabbix agent, you may run:

```
./configure --enable-agent
```

or, for Zabbix agent 2:

```
./configure --enable-agent2
```

Note:
A supported Go version must be installed to build Zabbix agent 2.

Notes on compilation options:

- `--enable-agent` - compiles Zabbix agent, as well as Zabbix get and Zabbix sender command-line utilities.
- `--with-libcurl` - required for virtual machine monitoring, SMTP authentication, and `web.page.*` Zabbix agent items, and telemetry query items. See also: Requirements (libcurl).

- `--with-libxml2` - required for virtual machine monitoring.
- `--with-libpcre2[=DIR]` - Zabbix always compiles with the PCRE2 library; this option only allows you to specify a custom PCRE2 installation path.
- `--with-mysql=/path/to/mysql_config` - specifies the path to a particular MySQL client library configuration. Useful when multiple versions of MySQL or MariaDB are installed.
- `--enable-static` - statically links libraries (not supported on Solaris). Use this if you plan to distribute compiled binaries to systems without the required libraries. Not recommended when building Zabbix server. To build the server statically, a static version of every external library is required. The configure script does not check for this automatically.
- `--with-stacksize=<value>` - sets the per-thread stack size in kilobytes (e.g., `--with-stacksize=512`). You can increase this value if Zabbix crashes or freezes due to stack overflows (e.g., during preprocessing on systems with low default thread stack limits).
- `--with-apm` (use only with `--enable-proxy`) - required for application performance monitoring (APM) on Zabbix proxy. See also: Requirements (g++, gRPC, protobuf, Abseil).

Attention:
If `./configure` fails due to missing libraries or other issues, please check the `config.log` file for detailed error information.
For example, if `libssl` is missing, the immediate error message may be misleading:
```
checking for main in -lmysqlclient... no
configure: error: Not found mysqlclient library
```
In this case, `config.log` reveals the actual cause:
```
/usr/bin/ld: cannot find -lssl
/usr/bin/ld: cannot find -lcrypto
```

See also:

- Compiling Zabbix with encryption support
- Known compilation issues

5 Make and install everything

Note:
If installing from Zabbix Git repository, it is required to run first:
```
$ make dbschema
```

```
make install
```

This step should be run as a user with sufficient permissions (commonly 'root', or by using `sudo`).

Running `make install` will by default install the daemon binaries (zabbix_server, zabbix_agentd, zabbix_proxy) in /usr/local/sbin and the client binaries (zabbix_get, zabbix_sender) in /usr/local/bin.

Note:
To specify a different location than `/usr/local`, use a `--prefix` key in the previous step of configuring sources, for example, `--prefix=/home/zabbix`. In this case daemon binaries will be installed under `<prefix>/sbin`, while utilities under `<prefix>/bin`. Man pages will be installed under `<prefix>/share`.

6 Review and edit configuration files

- edit the Zabbix agent configuration file (`zabbix_agentd.conf`).

You need to configure this file for every host with `zabbix_agentd` installed.

You must specify the Zabbix server **IP address** in the file. Connections from other hosts will be denied.

- edit the Zabbix server configuration file (`zabbix_server.conf`).

You must specify the database name, user and password (if using any).

The rest of the parameters will suit you with their defaults if you have a small installation (up to ten monitored hosts). You should change the default parameters if you want to maximize the performance of Zabbix server (or proxy) though.

- if you have installed a Zabbix proxy, edit the proxy configuration file (`zabbix_proxy.conf`).

You must specify the server IP address and proxy hostname (must be known to the server), as well as the database name, user and password (if using any).

Note:
With SQLite the full path to database file must be specified; DB user and password are not required.

7 Start up the daemons

Run zabbix_server on the server side.

```
zabbix_server
```

Note:
Make sure that your system allows allocation of 36MB (or a bit more) of shared memory; otherwise, the server may not start and you will see "Cannot allocate shared memory for <type of cache>." in the server log file. This may happen on FreeBSD, Solaris 8.

Run zabbix_agentd on all the monitored machines.

```
zabbix_agentd
```

Note:
Make sure that your system allows allocation of 2MB of shared memory; otherwise, the agent may not start and you will see "Cannot allocate shared memory for collector." in the agent log file. This may happen on Solaris 8.

If you have installed Zabbix proxy, run zabbix_proxy.

```
zabbix_proxy
```

Installing Zabbix web interface

Copying PHP files

Zabbix frontend is written in PHP, so to run it a PHP supported webserver is needed. Installation is done by simply copying the PHP files from the ui directory to the webserver HTML documents directory.

Common locations of HTML documents directories for Apache web servers include:

- /usr/local/apache2/htdocs (default directory when installing Apache from source)
- /srv/www/htdocs (OpenSUSE, SLES)
- /var/www/html (Debian, Ubuntu, Fedora, RHEL)

It is recommended to use a subdirectory instead of the HTML root. To create a subdirectory and copy Zabbix frontend files into it, execute the following commands, replacing the <htdocs> with actual directory:

```
mkdir <htdocs>/zabbix
cd ui
cp -a . <htdocs>/zabbix
```

If planning to use any other language than English, see Installation of additional frontend languages for instructions.

Installing frontend

Please see Web interface installation page for information about Zabbix frontend installation wizard.

Installing Zabbix agent 2 loadable plugins

Installing Zabbix agent 2 loadable plugins is only required if you want to monitor targets not covered by built-in plugins (e.g., MongoDB servers or clusters, PostgreSQL and its forks, etc.). See the full list of loadable plugins and built-in plugins.

Attention:
Before installing a plugin, please check its README file. It may contain specific requirements and installation instructions.

To install from sources, first download and extract the loadable plugin source archive.

To compile the plugin, navigate to the extracted plugin directory and run `make`:

```
make
```

Note:
A supported Go version must be installed to build Zabbix agent 2 loadable plugins.

The plugin executable may be placed anywhere as long as it is loadable by Zabbix agent 2. Specify the path to the plugin binary in the plugin configuration file, e.g., in `postgresql.conf` for the PostgreSQL plugin:

```
Plugins.PostgreSQL.System.Path=/path/to/executable/zabbix-agent2-plugin-postgresql
```

The path to the plugin configuration file must be specified in the `Include` parameter of the Zabbix agent 2 configuration file:

```
Include=/path/to/plugin/configuration/file/postgresql.conf
```

Proceed to setup for more details on configuring plugins.

Loadable plugins provided by Zabbix use simple makefiles with the following build targets:

- `make` - build the plugin
- `make clean` - delete all files that are created by building the plugin
- `make check` - run self-tests (requires a real monitoring target, e.g., a PostgreSQL database)
- `make style` - check Go code style with `golangci-lint`
- `make format` - format Go code with `go fmt`
- `make dist` - create a source archive including all dependencies

Installing Java gateway

It is required to install Java gateway only if you want to monitor JMX applications. Java gateway is lightweight and does not require a database.

To install from sources, first download and extract the source archive.

To compile Java gateway, run the `./configure` script with `--enable-java` option. It is recommended that you specify the `--prefix` option to request installation path other than the default /usr/local, because installing Java gateway will create a whole directory tree, not just a single executable.

```
./configure --enable-java --prefix=$PREFIX
```

To compile and package Java gateway into a JAR file, run `make`. Note that for this step you will need `javac` and `jar` executables in your path.

```
make
```

Now you have a zabbix-java-gateway-$VERSION.jar file in `src/zabbix_java/bin`. If you are comfortable with running Java gateway from `src/zabbix_java` in the distribution directory, then you can proceed to instructions for configuring and running Java gateway. Otherwise, make sure you have enough privileges and run `make install`.

```
make install
```

Proceed to setup for more details on configuring and running Java gateway.

Installing Zabbix web service

Installing Zabbix web service is only required if you want to use scheduled reports.

To install from sources, first download and extract the source archive.

To compile Zabbix web service, run the `./configure` script with `--enable-webservice` option.

> **Note:**
> A supported Go version must be installed to build Zabbix web service.

Run zabbix_web_service on the machine, where the web service is installed:

```
zabbix_web_service
```

Proceed to setup for more details on configuring Scheduled reports generation.

Installing Zabbix MCP server

To install from sources, first download and extract the source archive.

To compile Zabbix MCP server, run the following command in the MCP server source directory:

```
make build
```

> **Note:**
> A supported Go version must be installed to build Zabbix MCP server.

To report a specific version in the output of `zabbix_mcp_server -V` and in the `User-Agent` of Zabbix API requests, set the `VERSION` variable, for example, `make build VERSION=8.2.0`.

To install the binary to `/usr/local/sbin` and the example configuration file to `/etc/zabbix`, run:

```
sudo make install
```

Edit the configuration file `/etc/zabbix/zabbix_mcp_server.conf` and set at least the following parameters:

- `FrontendURL` - the Zabbix API URL, ending with `api_jsonrpc.php`;
- `AllowedIP` - the addresses of MCP clients that are allowed to connect.

To encrypt connections from MCP clients, set `TLSAccept=cert` and specify `TLSCertFile` and `TLSKeyFile`, or run the MCP server behind a trusted reverse proxy that handles TLS.

If the `zabbix` user does not exist (for example, if only the MCP server is installed on the machine), create it as described in Create user account.

Validate the configuration file and start zabbix_mcp_server:

```
zabbix_mcp_server -c /etc/zabbix/zabbix_mcp_server.conf -T
zabbix_mcp_server -c /etc/zabbix/zabbix_mcp_server.conf
```

Proceed to MCP server for more details on configuring the server and connecting MCP clients.

Getting Zabbix source code

There are several ways of getting Zabbix source code:

- You can download the released stable versions from the official Zabbix website
- You can download nightly builds from the official Zabbix website developer page
- You can get the latest development version from the Git source code repository system:
  - The primary location of the full repository is at https://git.zabbix.com/scm/zbx/zabbix.git
  - Master and supported releases are also mirrored to Github at https://github.com/zabbix/zabbix

A Git client must be installed to clone the repository. The official commandline Git client package is commonly called **git** in distributions. To install, for example, on Debian/Ubuntu, run:

```
sudo apt-get update
sudo apt-get install git
```

To grab all Zabbix source, change to the directory you want to place the code in and execute:

```
git clone https://git.zabbix.com/scm/zbx/zabbix.git
```

Compilation issues

These are the known issues regarding Zabbix compilation from sources. For all other cases, see the Known issues page.

Library in a non-standard location

Zabbix allows you to specify a library located in a non-standard location. In the example below, Zabbix will run `curl-config` from the specified non-standard location and use its output to determine the correct libcurl to use.

```
$ ./configure --enable-server --with-mysql --with-libcurl=/usr/local/bin/curl-config
```

This will work if it is the only libcurl installed in the system, but might not if there is another libcurl installed in a standard location (by the package manager, for example). Such is the case when you need a newer version of the library for Zabbix and the older one for other applications.

Therefore, specifying a component in a non-standard location will not always work when the same component also exists in a standard location.

For example, if you use a newer libcurl installed in `/usr/local` with the libcurl package still installed, Zabbix might pick up the wrong one and compilation will fail:

```
usr/bin/ld: ../../src/libs/zbxhttp/libzbxhttp.a(http.o): in function 'zbx_http_convert_to_utf8':
/tmp/zabbix-master/src/libs/zbxhttp/http.c:957: undefined reference to 'curl_easy_header'
collect2: error: ld returned 1 exit status
```

Here, the function `curl_easy_header()` is not available in the older `/usr/lib/x86_64-linux-gnu/libcurl.so`, but is available in the newer `/usr/local/lib/libcurl.so`.

The problem lies with the order of linker flags, and one solution is to specify the full path to the library in an LDFLAGS variable:

```
$ LDFLAGS="-Wl,--no-as-needed /usr/local/lib/libcurl.so" ./configure --enable-server --with-mysql --with-l
```

Note the `-Wl,--no-as-needed` option which might be needed on some systems (see also: default linking options on Debian-based systems).

Stack size too small on some systems

If Zabbix crashes or freezes due to stack overflows, you can increase the per-thread stack size using the `--with-stacksize` option when configuring the sources. This issue may occur on systems with low default thread stack limits, especially during preprocessing, where multiple threads are created.

The following example sets the stack size to 512 KB per thread:

```
./configure --enable-server --with-mysql --with-stacksize=512
```

You can check the system thread stack limits at runtime using the `ulimit -s` command on Linux-based systems.

Notes on installation from packages

Overview

Official **pre-release** packages of Zabbix 8.2 are available on the Zabbix website, where you can select your operating system and Zabbix component to generate installation instructions and commands.

Along with the package installation notes below, please also consult the Installation and setup page.

Packages are available for the following Linux distributions:

- Red Hat Enterprise Linux and its derivatives: AlmaLinux, Amazon Linux 2023, CentOS Stream, CentOS 7, Oracle Linux, Rocky Linux
- Debian, Ubuntu, Raspberry Pi OS, Raspbian
- SUSE Linux Enterprise Server, openSUSE Leap

> **Attention:**
> Some OS distributions (in particular, Debian-based distributions) provide their own Zabbix packages. These packages are **not** supported by Zabbix and may be outdated or missing the latest features and bug fixes. It is recommended to use only official packages from the Zabbix Official Repository. If you previously installed Zabbix from your operating system's repository, see the steps for upgrading Zabbix packages from OS repositories.

Packages support MySQL/PostgreSQL database and Apache/Nginx web server. Note that Zabbix server and proxy cannot share the same database; use different database names if both are installed on the same host.

If necessary, separate packages for Zabbix agent/agent 2, Zabbix get, and Zabbix sender are available at the Zabbix Official Repository.

On Debian- or Ubuntu-based systems, installing an additional Zabbix package with `apt` may also select upgrades for already installed Zabbix packages from the same repository. To install a package without upgrading an installed Zabbix server or proxy package, temporarily hold that package before installation and unhold it afterwards:

```
apt-mark hold zabbix-server-mysql
apt install zabbix-sender
apt-mark unhold zabbix-server-mysql
```

Zabbix also provides Zabbix agent pre-compiled binaries for non-Linux operating systems; see:

- Windows agent installation from MSI
- macOS agent installation from PKG
- Zabbix agent on IBM AIX
- Legacy binaries (for older/less common systems, such as HP-UX, NetBSD, Tru64, and older versions of SLES)

Package installation notes

The following notes apply to all systems:

- If using PostgreSQL, `DBHost=localhost` (or an IP address) in Zabbix server/proxy configuration makes PostgreSQL use a network socket instead of a local UNIX socket; see SELinux configuration for related setup instructions.
- If using TimescaleDB, see additional TimescaleDB setup.
- If installing Zabbix Java gateway (for monitoring JMX applications), see additional setup instructions for RHEL-based systems and Debian-based systems.
- If installing Zabbix MCP server, install the `zabbix-mcp-server` package and update the configuration file (at least `FrontendURL` and `AllowedIP`) before starting the service.
- For running Zabbix agent as root, see Running agent as root.

- Application performance monitoring (APM) on Zabbix proxy is supported by these packages:
  - RHEL 10 and its derivatives - install the Zabbix proxy APM package (e.g., `zabbix-proxy-mysql-apm` instead of `zabbix-proxy-mysql`); it also requires you to install the EPEL repository and enable the CRB dependency package repositories.
  - Debian 12 or later, Raspberry Pi OS (Raspbian) 12 or later, Ubuntu 24.04 or later - APM is included in Zabbix proxy packages.
  - On other systems, you can install Zabbix proxy from sources with the `--with-apm` option.

The following notes apply to RHEL and its derivatives:

- If you've enabled the EPEL repository (e.g., for the Zabbix proxy APM package), which also provides Zabbix packages, it must be excluded from package resolution before installing official Zabbix packages; see Accidental installation of EPEL Zabbix packages.
- For installing Zabbix packages on Red Hat UBI environments, see Zabbix packages for RHEL on Red Hat UBI environments.
- For using ICMP ping items, packages for `fping` are also available at the Zabbix Official Repository.

SELinux configuration

Zabbix uses socket-based inter-process communication. On systems where Security-Enhanced Linux (SELinux) is enabled, you may need to add SELinux rules to allow Zabbix create/use UNIX domain sockets in the SocketDir directory. Socket files are used by Zabbix server (alerter, preprocessing, IPMI) and Zabbix proxy (IPMI), and they're present while the process is running.

Having SELinux enabled in enforcing mode, execute the following commands to enable communication between Zabbix frontend and server:

For RHEL 7 (and later), AlmaLinux, CentOS Stream, Oracle Linux, Rocky Linux 8 (and later):

```
setsebool -P httpd_can_connect_zabbix on
```

If the database is accessed over the network (including `localhost` for PostgreSQL), also allow Zabbix frontend to connect to the database:

```
setsebool -P httpd_can_network_connect_db on
```

For RHEL prior to 7:

```
setsebool -P httpd_can_network_connect on
setsebool -P zabbix_can_network on
```

After applying SELinux settings, restart Apache:

```
systemctl restart httpd
```

Optionally, you can install a pre-defined `zabbix-selinux-policy` package from the Zabbix Official Repository. This package is provided for all supported OS versions to simplify Zabbix deployment and prevent users from turning off SELinux because of the configuration complexity.
 For version 10 of AlmaLinux, CentOS Stream, RHEL, Rocky Linux, and Oracle Linux, the SELinux policy is provided by the `zabbix-selinux-policy-main` package; the `zabbix-selinux-policy` package has a dependency on either `zabbix-selinux-policy-main` or `selinux-policy-targeted-extra`. The `zabbix-selinux-policy-main` package conflicts with `selinux-policy-targeted-extra`; therefore, both packages cannot be installed at the same time.

> **Attention:**
> For maximum security, it is recommended to set custom SELinux settings.

The `zabbix-selinux-policy` (`zabbix-selinux-policy-main` for version 10 of AlmaLinux, CentOS Stream, RHEL, Rocky Linux, and Oracle Linux) package contains a basic SELinux policy that allows Zabbix to create and use sockets and enables HTTPd connection to PostgreSQL (used by frontend).

The source `zabbix_policy.te` file contains the following rules:

```
module zabbix_policy 1.2;

require {
    type zabbix_t;
    type zabbix_port_t;
    type zabbix_var_run_t;
    type postgresql_port_t;
    type httpd_t;
    class tcp_socket name_connect;
    class sock_file { create unlink };
    class unix_stream_socket connectto;
```

```
}

####============= zabbix_t ==============
allow zabbix_t self:unix_stream_socket connectto;
allow zabbix_t zabbix_port_t:tcp_socket name_connect;
allow zabbix_t zabbix_var_run_t:sock_file create;
allow zabbix_t zabbix_var_run_t:sock_file unlink;
allow httpd_t zabbix_port_t:tcp_socket name_connect;

####============= httpd_t ==============
allow httpd_t postgresql_port_t:tcp_socket name_connect;
```

Debuginfo packages

Debuginfo packages contain debugging symbols for Zabbix binaries. They are not required for normal installation or operation but are useful for advanced troubleshooting.

Please note that Debuginfo packages are available for the following versions: Red Hat Enterprise Linux 7 and earlier (including derivatives), and SUSE Linux Enterprise Server 15 and earlier (including derivatives).

To enable the zabbix-debuginfo repository:

- On RHEL 7, edit /etc/yum.repos.d/zabbix.repo and set enabled=1 for the zabbix-debuginfo section:

```
[zabbix-debuginfo]
name=Zabbix Official Repository debuginfo - $basearch
baseurl=http://repo.zabbix.com/zabbix/8.2/stable/rhel/7/$basearch/debuginfo/
enabled=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-ZABBIX-A14FE591
gpgcheck=1
```

- On SUSE, edit /etc/zypp/repos.d/zabbix.repo and set enabled=1 for the zabbix-debuginfo section:

```
[zabbix-debuginfo]
name=Zabbix Official Repository debuginfo
type=rpm-md
baseurl=https://repo.zabbix.com/zabbix/8.2/unstable/sles/15/$basearch/debuginfo/
gpgcheck=1
gpgkey=https://repo.zabbix.com/zabbix/8.2/unstable/sles/15/$basearch/debuginfo/repodata/repomd.xml.key
enabled=0
update=1
```

Once enabled, install the packages:

- On RHEL, install a single package with debug information for all Zabbix components:

```
dnf install zabbix-debuginfo
```

- On SUSE, install component-specific debuginfo packages:

```
zypper install zabbix-<component>-debuginfo
```

Enabling unstable release repositories

The instructions below are for enabling unstable Zabbix release repositories (disabled by default) used for minor Zabbix version release candidates.

First, install or update to the latest zabbix-release package. To enable rc packages on your system do the following:

Red Hat Enterprise Linux

Open the /etc/yum.repos.d/zabbix.repo file and set enabled=1 for the zabbix-unstable repo.

```
[zabbix-unstable]
name=Zabbix Official Repository (unstable) - $basearch
baseurl=https://repo.zabbix.com/zabbix/8.2/unstable/rhel/8/$basearch/
enabled=1
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-ZABBIX-A14FE591
```

Debian/Ubuntu

Open the /etc/apt/sources.list.d/zabbix.list and uncomment "Zabbix unstable repository".

```
#### Zabbix unstable repository
deb https://repo.zabbix.com/zabbix/8.2/unstable/debian bullseye main
deb-src https://repo.zabbix.com/zabbix/8.2/unstable/debian bullseye main
```

SUSE

Open the /etc/zypp/repos.d/zabbix.repo file and set enable=1 for the zabbix-unstable repo.

```
[zabbix-unstable]
name=Zabbix Official Repository
type=rpm-md
baseurl=https://repo.zabbix.com/zabbix/8.2/unstable/sles/15/x86_64/
gpgcheck=1
gpgkey=https://repo.zabbix.com/zabbix/8.2/unstable/sles/15/x86_64/repodata/repomd.xml.key
enabled=1
update=1
```

Windows agent installation from MSI

Overview

Zabbix agent can be installed on Windows using 32-bit or 64-bit MSI installer packages, available for download.

Minimum OS requirements for MSI installation are:

- **For Zabbix agent:** Windows XP (64-bit) or Windows Server 2003
- **For Zabbix agent 2:** Windows 10 (32-bit) or Windows Server 2016

32-bit packages cannot be installed on 64-bit systems.

Packages include:

- TLS support (TLS configuration is optional)
- Zabbix get and Zabbix sender utilities (can be installed alongside Zabbix agent/agent 2 or separately)

> **Attention:**
> Zabbix agent 2 packages do not include loadable plugins (MongoDB, PostgreSQL, MSSQL), which need to be downloaded and installed separately.

Installation can be done using the Setup Wizard or the command line.

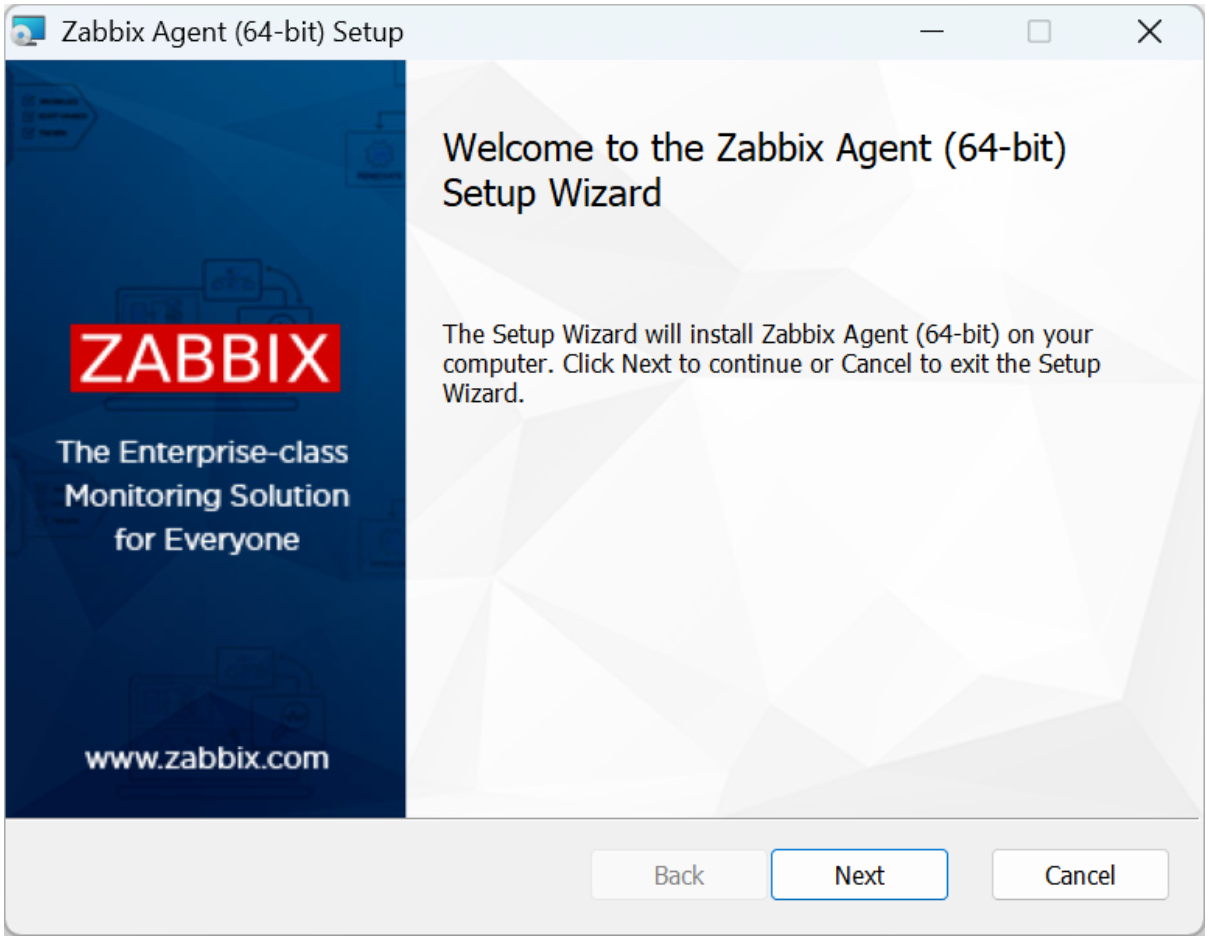
Although installation using MSI packages is fully supported, installing at least Microsoft .NET Framework 2 is recommended for proper error handling.

> **Attention:**
> It is recommended to use the default installation location provided by the installer. Using a custom location without the necessary permissions may compromise the security of the installation.

Installation from Setup Wizard

The following installation steps apply to both Zabbix agent and Zabbix agent 2.

1. Double-click the downloaded MSI file to start the installation:

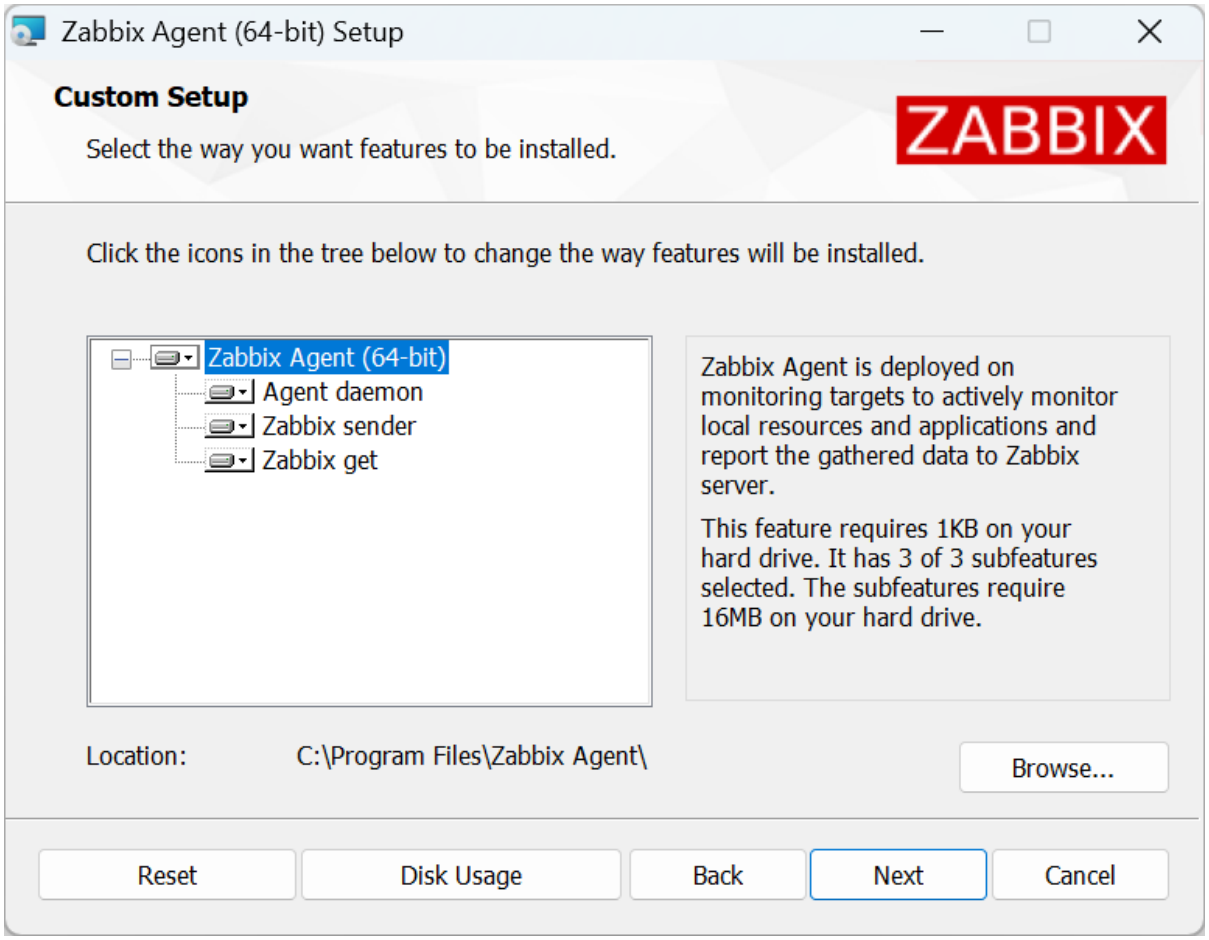


2. Accept the End-User License Agreement:



3. Select the Zabbix components (Agent daemon, Zabbix sender, Zabbix get) to be installed:

Attention:
It is recommended to use the default installation location provided by the installer. Using a custom location without the necessary permissions may compromise the security of the installation.

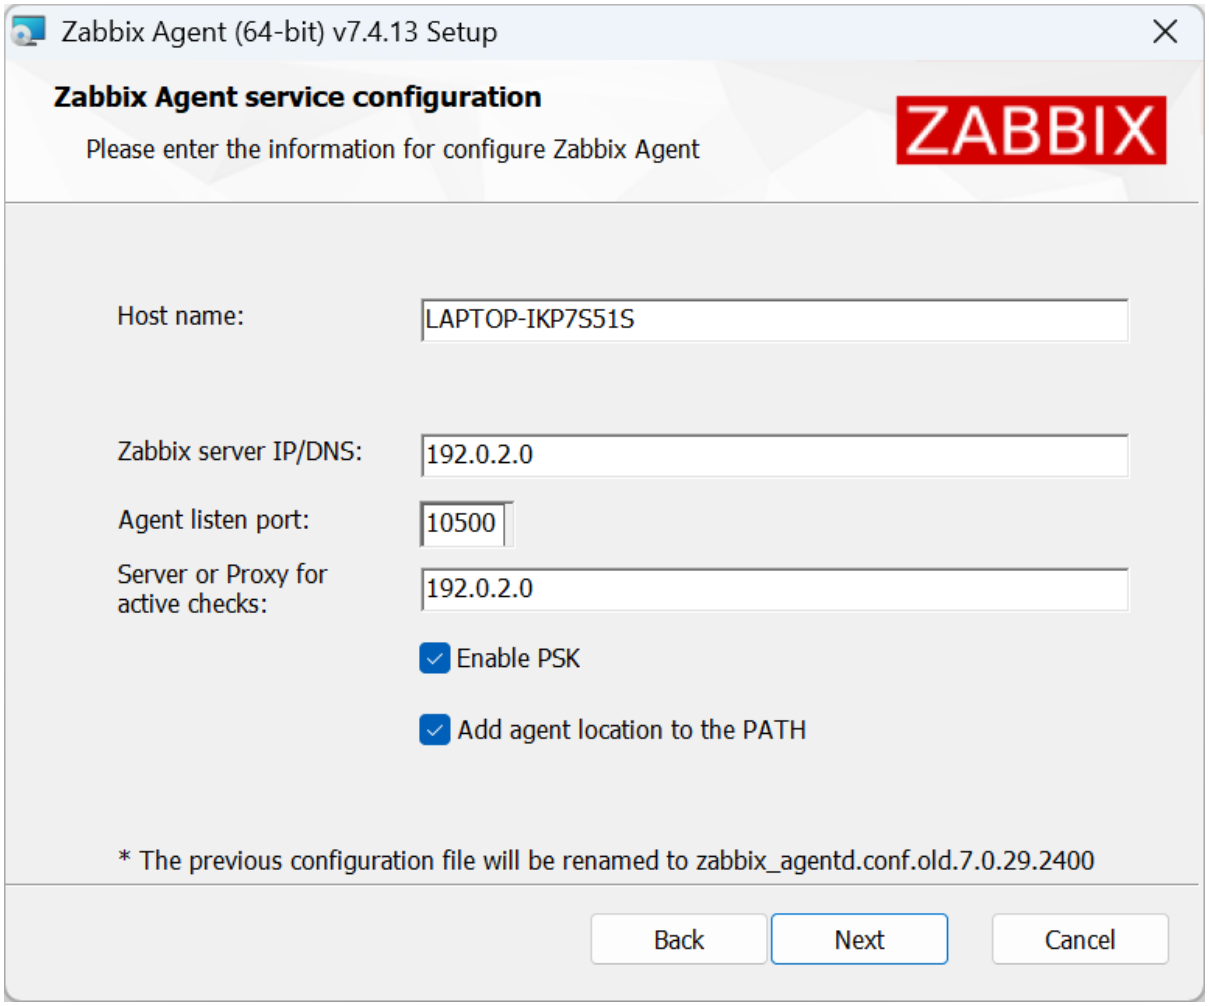


Note:
If a custom location outside the `Program Files` directory is used and the destination folder already exists, the installer will not modify its permissions. During a new installation, you will be prompted to either set the appropriate permissions manually or continue without modifying them.

4. Configure the following parameters. Their values will be set in the Zabbix agent configuration file:

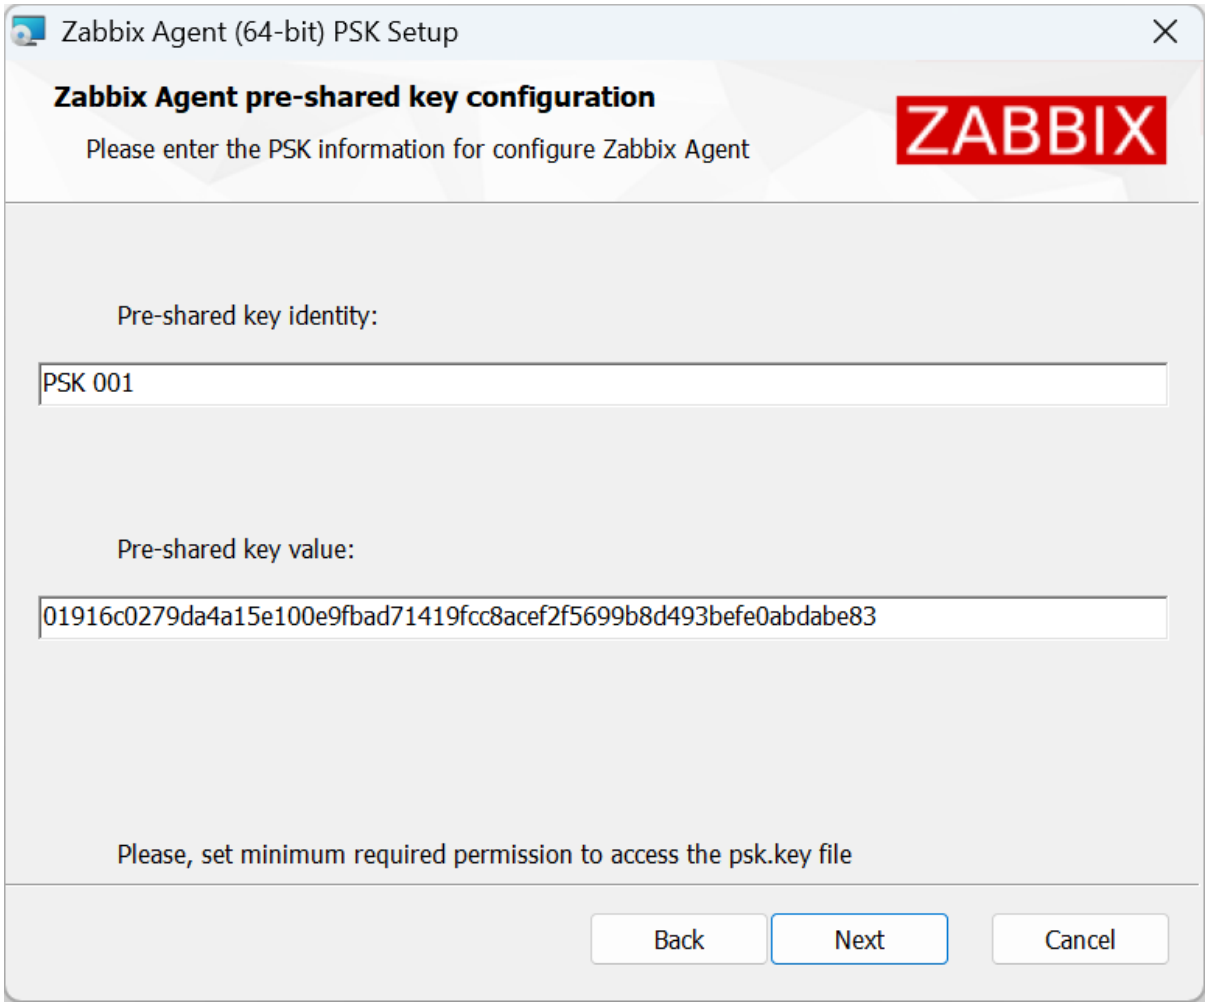
Parameter	Description
Host name	The host name of the machine where Zabbix agent is being installed. Sets the Hostname parameter.
Zabbix server IP/DNS	A list of comma-delimited IP addresses, optionally in CIDR notation, or DNS names of Zabbix servers or Zabbix proxies. This parameter is **mandatory**. Sets the Server parameter.
Agent listen port	The agent will listen on this port for connections from the server. Sets the ListenPort parameter.
Server or Proxy for active checks	The Zabbix server/proxy address or cluster configuration to get active checks from. The server/proxy address is an IP address or DNS name and optional port separated by colon. Sets the ServerActive parameter.
Enable PSK	Mark the checkbox to enable TLS support using pre-shared keys. Sets the TLSConnect and TLSAccept parameters to `psk`.
Add agent location to the PATH	Mark the checkbox to add Zabbix agent location to the system PATH variable.

Note:
If an existing Zabbix agent is detected, the parameter values from its configuration file will be displayed. Additionally the existing configuration file will be renamed during installation, and a new configuration file will be created.



5. Configure PSK parameters if you marked the *Enable PSK* checkbox in the previous step. These parameters will also be set in the Zabbix agent configuration file:

Parameter	Description
Pre-shared key identity	The pre-shared key identity string. Sets the TLSPSKIdentity parameter.
Pre-shared key value	The pre-shared key string value. Creates the `psk.key` file containing the key and sets the TLSPSKFile parameter to the key location (default: `C:\Program Files\Zabbix Agent\psk.key`). It is **recommended** to restrict access to the pre-shared key file by adjusting the file's security settings so that only Zabbix agent (or the user running the agent) can read it.

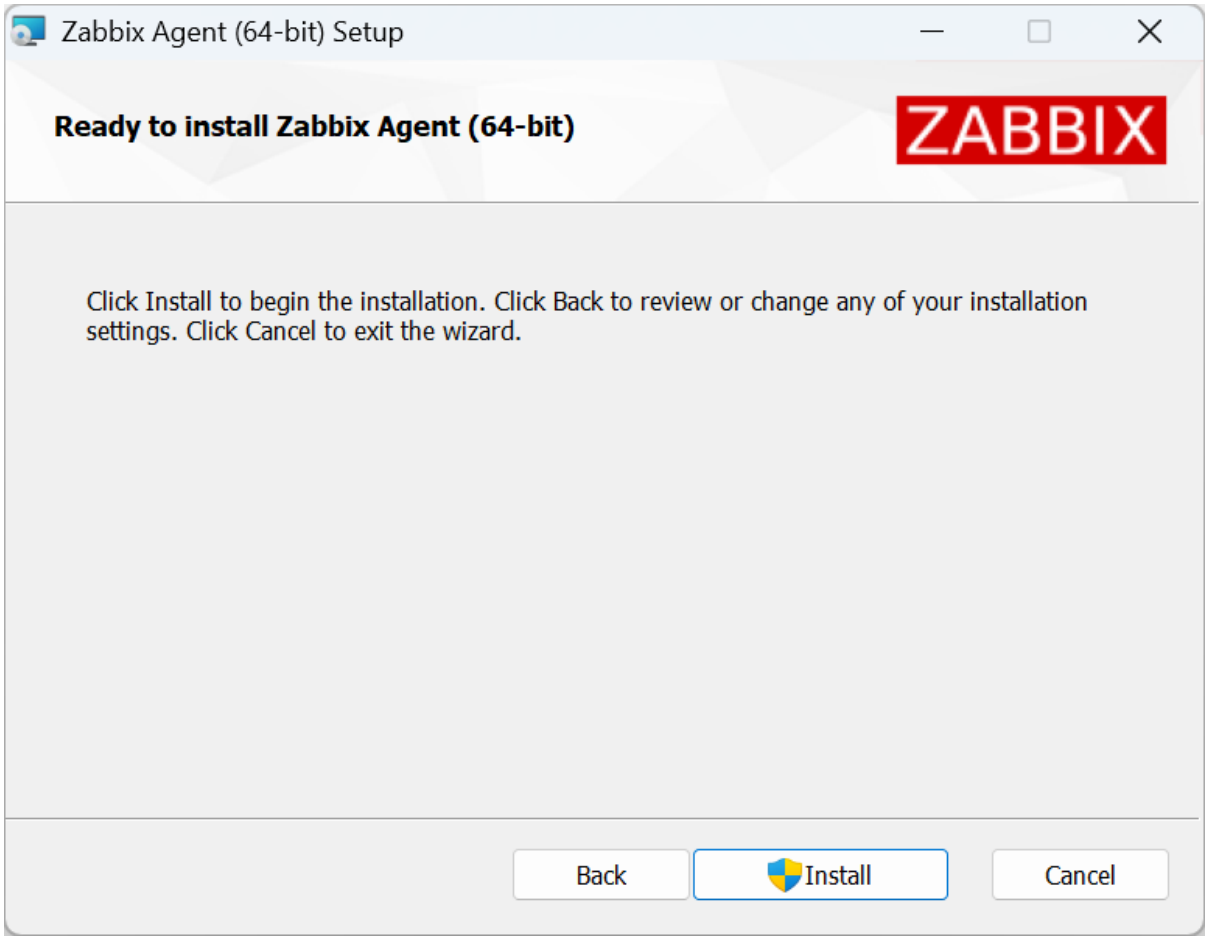


6. Click *Install* to begin the installation.

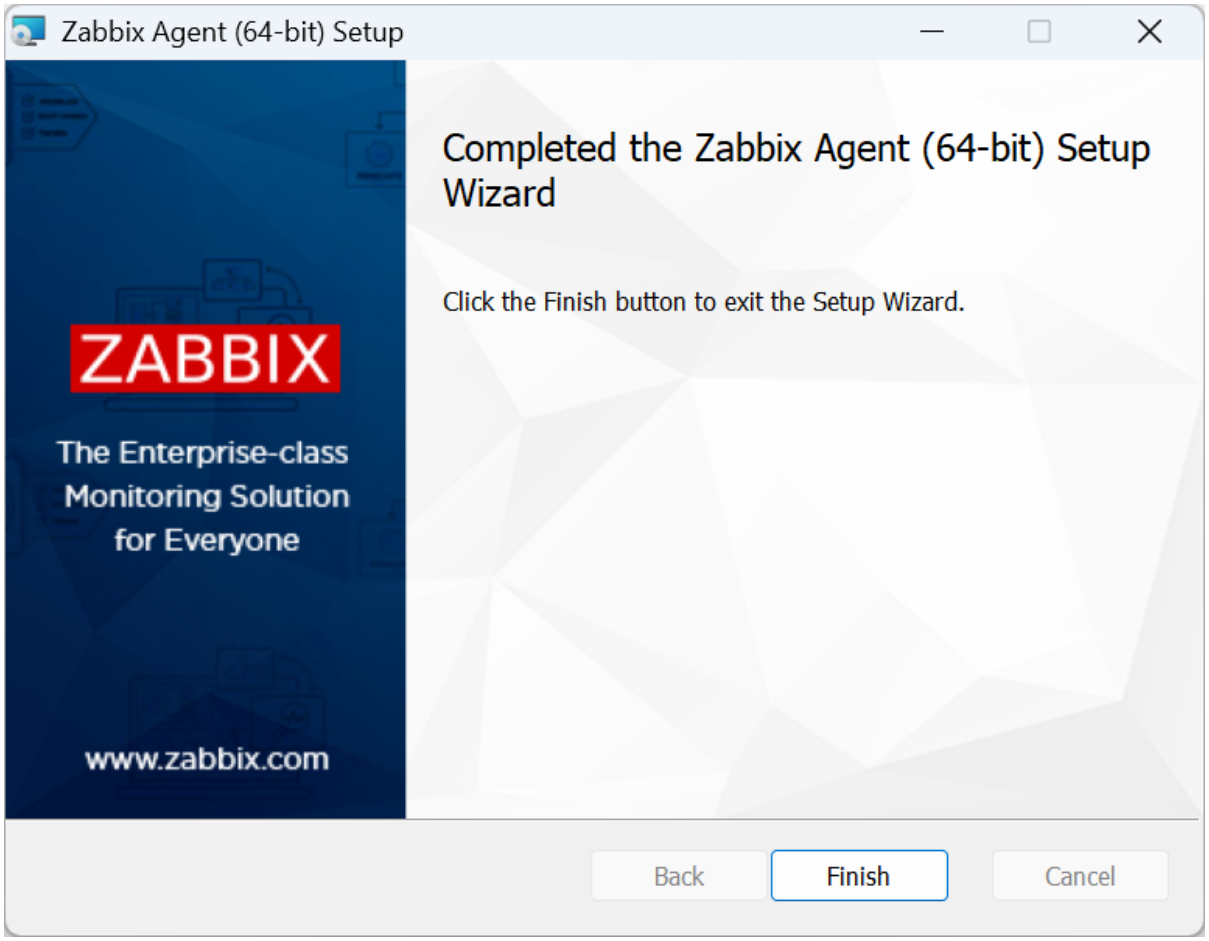
All selected Zabbix components and the Zabbix agent configuration file will be installed in your specified location (default: `C:\Program Files\Zabbix Agent`). The same applies to Zabbix agent 2, except additional configuration files for its built-in plugins will be installed in the `zabbix_agent2.d\plugins.d` subfolder.

Additionally, `zabbix_agentd.exe` (or `zabbix_agent2.exe`) will be set up as a Windows service with delayed automatic startup (or automatic startup on Windows versions before Windows Vista/Server 2008).

If a different version of Zabbix agent is running during installation, you will be prompted to choose either to close the application and try restarting it or leave it open, in which case a reboot will be required.



7. Click the *Finish* button to exit the Setup Wizard.



Installation from command line

Zabbix agent can be installed from the command line by running the MSI installer with `msiexec`. For example:

```
msiexec.exe /l*v "C:\package.log" /i "C:\zabbix_agent-8.2.0-windows-amd64-openssl.msi" /qn+ SERVER=192.0.2
```

This method allows for unattended installations and custom configurations using parameters.

Supported parameters

Zabbix agent MSI installer packages support the following parameters for both Zabbix agent and Zabbix agent 2.

Note:
Zabbix agent/agent2 parameters are set in the configuration file during installation. Click a parameter name to view its detailed description and configuration examples on the Zabbix agent (Windows) page. For Zabbix agent 2, refer to the Zabbix agent 2 (Windows) page.

Parameter	Description
ADDDEFAULT	A list of comma-delimited components to install in their default configuration. For more information, see ADDDEFAULT property. Possible values: `AgentProgram`, `GetProgram`, `SenderProgram`, `ALL` Example: `ADDDEFAULT=AgentProgram,GetProgram`
ADDLOCAL	A list of comma-delimited components to install locally. For more information, see ADDLOCAL property. Possible values: `AgentProgram`, `GetProgram`, `SenderProgram`, `ALL` Example: `ADDLOCAL=AgentProgram,SenderProgram`
ALLOWDENYKEY	A list of semicolon-delimited AllowKey or DenyKey parameters to restrict Zabbix agent checks. If necessary, use a backslash to escape the delimiter (`\;`). Sets the AllowKey and DenyKey parameters in the agent configuration file. Example: `ALLOWDENYKEY="AllowKey=system.run[more C:\Windows\System32\drivers\etc\hosts\\; echo 'File read complete'];DenyKey=system.run[*]"`
CONF	The full pathname to a template configuration file for Zabbix agent. During installation, this file will become the agent configuration file. The file must contain at least the Server and LogFile parameters. Example: `CONF="C:\full\path\to\example.conf"`
DONOTSTART	Use `DONOTSTART=1` to prevent the MSI installer from starting the Zabbix agent service.
ENABLEPATH	Use `ENABLEPATH=1` to add Zabbix agent location to the system PATH variable.
ENABLEPERSISTENTBUFFER	Zabbix agent 2 only. Enable the usage of local persistent storage for active items.
HOSTINTERFACE	An optional parameter that defines the host interface.
HOSTMETADATA	An optional parameter that defines the host metadata.
HOSTMETADATAITEM	An optional parameter that defines an item used for getting the host metadata.
HOSTNAME	An optional parameter that defines the hostname.
HOSTNAMEITEM	An optional parameter that defines an item used for getting the host name.
INCLUDE	A list of semicolon-delimited individual files or all files in a directory to include in the Zabbix agent configuration file.
IGNOREALREADYEXISTSFOLDER	For new installations, use `IGNOREALREADYEXISTSFOLDER=1` to allow installation into an existing destination folder outside the `Program Files` directory. If this parameter is not specified, the installation fails when such a destination folder already exists.
INSTALLFOLDER	The full pathname to a folder where Zabbix components and the Zabbix agent configuration file will be installed. For Zabbix agent 2, additional configuration files for built-in plugins will be installed in the `zabbix_agent2.d\plugins.d` subfolder. Example: `INSTALLFOLDER="C:\Program Files\Zabbix Agent"`
LISTENIP	A list of comma-delimited IP addresses that the agent should listen on.
LISTENPORT	The agent will listen on this port for connections from the server.
LOGFILE	The name of the Zabbix agent log file.
LOGTYPE	The type of the log output.
NONMSICONFNAME	The full pathname to a custom configuration file for Zabbix agent. During installation, any valid agent configuration parameters present in this file (limited to those listed in this table) will be written in the newly created agent configuration file. The file must contain at least the Server parameter. Example: `NONMSICONFNAME="C:\full\path\to\example.conf"`
PERSISTENTBUFFERFILE	Zabbix agent 2 only. The file where Zabbix agent 2 should keep the SQLite database.
PERSISTENTBUFFERPERIOD	Zabbix agent 2 only. The time period for which data should be stored when there is no connection to the server or proxy.

Parameter	Description
SERVER	A list of comma-delimited IP addresses, optionally in CIDR notation, or DNS names of Zabbix servers or Zabbix proxies. This parameter is **mandatory**, except when STARTAGENTS is set to 0.
SERVERACTIVE	The Zabbix server/proxy address or cluster configuration to get active checks from.
SKIP	Use `SKIP=fw` to prevent the MSI installer from adding a Windows Firewall exception rule for Zabbix agent.
SOURCEIP	The source IP address for outgoing connections to Zabbix server or Zabbix proxy, or for making connections while executing some items (web.page.get, net.tcp.port, etc.).
STARTAGENTS	The number of pre-forked instances of zabbix_agentd that process passive checks. If set to 0, passive checks are disabled and the agent will not listen on any TCP port.
STARTUPTYPE	Startup type of the Zabbix agent service. Possible values: **automatic** - start the service automatically at Windows startup; **delayed** - *(default)* delay starting the service after the automatically started services have completed startup (available on Windows Vista/Server 2008 and later versions); **manual** - start the service manually (by a user or application); **disabled** - disable the service so that it cannot be started by a user or application. Example: `STARTUPTYPE=disabled`
STATUSPORT	Zabbix agent 2 only. If set, the agent will listen on this port for HTTP status requests (http://localhost:<port>/status).
TIMEOUT	Specifies how long to wait (in seconds) for establishing connection and exchanging data with Zabbix proxy or server.
TLSACCEPT	The incoming connections to accept (used for passive checks). If set to `psk`, then TLSCONNECT will also be set to `psk` (unless specified otherwise).
TLSCAFILE	The full pathname of a file containing the top-level CA(s) certificates for peer certificate verification.
TLSCERTFILE	The full pathname of a file containing the agent certificate or certificate chain.
TLSCONNECT	How the agent should connect to Zabbix server or proxy (used for active checks). If set to `psk`, then TLSACCEPT will also be set to `psk` (unless specified otherwise).
TLSCRLFILE	The full pathname of a file containing revoked certificates.
TLSKEYFILE	The full pathname of a file containing the Zabbix agent private key.
TLSPSKFILE	The full pathname of a file containing the Zabbix agent pre-shared key. If both TLSPSKFILE and TLSPSKVALUE are set, the value of TLSPSKVALUE will be written into the file specified in TLSPSKFILE. It is **recommended** to restrict access to the pre-shared key file by adjusting the file's security settings so that only Zabbix agent (or the user running the agent) can read it.
TLSPSKIDENTITY	The pre-shared key identity string.
TLSPSKVALUE	The pre-shared key string value. If both TLSPSKFILE and TLSPSKVALUE are set, the value of TLSPSKVALUE will be written into the file specified in TLSPSKFILE. Example: `TLSPSKVALUE=1f87b595725ac58dd977beef14b97461a7c1045b9a1c963065002c5473194952`
TLSSERVERCERTISSUER	The allowed server (proxy) certificate issuer.
TLSSERVERCERTSUBJECT	The allowed server (proxy) certificate subject.
UNSAFEUSERPARAMETERS	Allow all characters to be passed in arguments to user-defined parameters.

Examples

The following example installs Zabbix agent with custom configuration. It also enables TLS support using pre-shared keys.

```
mkdir "C:\Program Files\Zabbix Agent" 2>nul
msiexec.exe /l*v "C:\package.log" /i "C:\zabbix_agent-8.2.0-windows-amd64-openssl.msi" /qn+^
 SERVER=192.0.2.1^
 INSTALLFOLDER="C:\Program Files\Zabbix Agent"^
 HOSTNAME=LAPTOP-IKP7S51S^
 TLSACCEPT=psk^
 TLSCONNECT=psk^
 TLSPSKIDENTITY="PSK 001"^
 TLSPSKFILE="C:\Program Files\Zabbix Agent\psk.key"^
 TLSPSKVALUE=1f87b595725ac58dd977beef14b97461a7c1045b9a1c963065002c5473194952^
 ENABLEPATH=1^
 ALLOWDENYKEY="AllowKey=system.run[type C:\Windows\System32\drivers\etc\hosts];DenyKey=system.run[*]"
```

The next example installs a newer version of Zabbix agent and uses a **template configuration file** (`CONF="C:\agent-template.conf"`). During installation, this file will become the agent configuration file. To inherit parameters from the old configuration file, use parameter placeholders (e.g., `[AllowDenyKey]`).

```
msiexec.exe /l*v "C:\package.log" /i "C:\zabbix_agent-8.2.1-windows-amd64-openssl.msi" /qn+ NONMSICONFNAME

#### agent-template.conf example:
LogFile=[LogFile]
[AllowDenyKey]
Server=192.0.2.8
Hostname=DESKTOP-X9F4A2B
[Include]
[TLSConnect]
[TLSAccept]
[TLSPSKIdentity]
[TLSPSKFile]
```

Alternatively, you can use a **custom configuration file** (`NONMSICONFNAME="C:\agent-custom.conf"`). During installation, any valid agent configuration parameters present in this file (limited to those listed in the table above) will be written in the newly created agent configuration file. To keep existing agent configuration, define parameters to be preserved.

```
msiexec.exe /l*v "C:\package.log" /i "C:\zabbix_agent-8.2.1-windows-amd64-openssl.msi" /qn+ NONMSICONFNAME

#### agent-custom.conf example:
Server=192.0.2.8
Hostname=DESKTOP-X9F4A2B
```

Zabbix agent 2 loadable plugins

Zabbix agent 2 loadable plugins can be installed on Windows using 64-bit MSI installer packages, available for download.

Minimum OS requirements for MSI installation are Windows 10 (64-bit) or Windows Server 2016.

> **Attention:**
> Before installing a plugin, please check its README file. It may contain specific requirements and installation instructions.

Similarly to Zabbix agent/agent2, loadable plugins can be installed using the Setup Wizard or the command line.

Installation from Setup Wizard

1. Double-click the downloaded MSI file to start the installation.
2. Accept the End-User License Agreement.
3. Select the Zabbix agent 2 loadable plugins to be installed.

> **Attention:**
> It is recommended to use the default installation location provided by the installer. Using a custom location without the necessary permissions may compromise the security of the installation.

4. Click *Install* to begin the installation.

All selected Zabbix agent 2 loadable plugins will be installed in your specified location (default: `C:\Program Files\Zabbix Agent 2`), with their configuration files installed in the `zabbix_agent2.d` subfolder.

5. Click the *Finish* button to exit the Setup Wizard.

Installation from command line

Zabbix agent 2 loadable plugins can be installed from the command line by running the MSI installer with `msiexec`. For example:

```
msiexec.exe /l*v "C:\package.log" /i "C:\zabbix_agent2_plugins-8.2.0-windows-amd64.msi" /qn+
```

Zabbix agent 2 loadable plugin MSI installer packages support the following parameters.

Parameter	Description
ADDDEFAULT	A list of comma-delimited components to install in their default configuration. For more information, see ADDDEFAULT property. Possible values: `ALL`, `CephPlugin`, `EmberplusPlugin`, `MongodbPlugin`, `MssqlPlugin`, `NvidiagpuPlugin`, `PostgresqlPlugin` Example: `ADDDEFAULT=MongodbPlugin,PostgresqlPlugin`

Parameter	Description
ADDLOCAL	A list of comma-delimited components to install locally. For more information, see ADDLOCAL property. Possible values: `ALL`, `CephPlugin`, `EmberplusPlugin`, `MongodbPlugin`, `MssqlPlugin`, `NvidiagpuPlugin`, `PostgresqlPlugin` Example: `ADDLOCAL=MongodbPlugin,MssqlPlugin`
INSTALLFOLDER	The full pathname to a folder where Zabbix components will be installed, with their configuration files installed in the `zabbix_agent2.d` subfolder. Example: `INSTALLFOLDER="C:\Program Files\Zabbix Agent 2"`

Building Zabbix agent 2 on Windows

Overview

This page demonstrates how to build Zabbix agent 2 from sources on Windows 10 (64-bit or 32-bit).

Both 32-bit and 64-bit versions can be built on a 64-bit platform, but only the 32-bit version can be built on a 32-bit platform.

Building Zabbix agent 2 requires:

- MinGW build tools
- Go programming language
- OpenSSL (for encryption features in Zabbix)
- PCRE2 (Perl Compatible Regular Expressions; for regular expressions pattern matching features in Zabbix)

You can build Zabbix agent 2 using one of the following methods:

- Using vcpkg - an automated approach that simplifies dependency management using a C++ package manager.
- Manual build - a manual approach that requires installing all dependencies before compiling the agent.

Attention:
Before starting the build process, please keep in mind:

- To execute commands, use the Command Prompt, launched by a user with sufficient permissions to write to protected folders. However, when installing OpenSSL and PCRE2, use the MSYS2 terminal.
- It is recommended to create a working directory at `C:\Zabbix` for all source files and build folders. However, compiled components should be installed in `C:\Zabbix\x64` (or `C:\Zabbix\x86` for 32-bit builds).

Building Zabbix agent 2 with vcpkg

This section contains instructions for building Zabbix agent with vcpkg, a package manager that simplifies dependency management and integration with C++ projects.

1. Download and install Build Tools for Visual Studio 2022. During installation, make sure to select the *Desktop development with C++* workload, which includes the vcpkg package manager.

2. Download and install Go (available as an MSI installer). During installation, make sure to specify `C:\Zabbix\Go` as the installation folder.

3. Download the MinGW distribution that uses the Microsoft Visual C runtime library; for example:

- For 64-bit builds: `x86_64-15.1.0-release-win32-seh-msvcrt-rt_v12-rev0.7z`
- For 32-bit builds: `i686-15.1.0-release-win32-dwarf-msvcrt-rt_v12-rev0.7z`

Attention:
The MinGW version must meet the minimum requirements.

Then, extract it to `C:\Zabbix\mingw64` (or `C:\Zabbix\mingw32` for 32-bit builds).

4. Initialize vcpkg and install the dependencies required for building Zabbix agent 2 (note that this may take some time):

```
cd C:\Zabbix

set PATH=%PATH%;"C:\Program Files (x86)\Microsoft Visual Studio\2022\BuildTools\VC\vcpkg"
vcpkg new --application
vcpkg add port pcre2
vcpkg add port libiconv
vcpkg add port openssl
```

```
#### For 64-bit builds:
set PATH=C:\Zabbix\mingw64\bin;%PATH%
vcpkg install --triplet x64-mingw-static --x-install-root=x64

#### For 32-bit builds:
set PATH=C:\Zabbix\mingw32\bin;%PATH%
vcpkg install --triplet x86-mingw-static --x-install-root=x86
```

5. Download the Zabbix source archive and extract it to `C:\Zabbix\zabbix-8.2.0`.

6. Navigate to the Zabbix build directory (`C:\Zabbix\zabbix-8.2.0\build\mingw`) and create the following `build.bat` script:

- For 64-bit builds:

```
:: Add MinGW and Go to the system `PATH` variable for the current session:
set PATH=C:\Zabbix\mingw64\bin;%PATH%
set PATH=C:\Zabbix\Go\bin;%PATH%

:: Set vcpkg installation path:
set vcpkg="C:\Zabbix\x64\x64-mingw-static"

:: Set linker flags for Crypt32 library:
SET CGO_LDFLAGS="-lCrypt32"

:: Run the build process:
mingw32-make GOFLAGS="-buildvcs=false" ARCH=AMD64 ^
    PCRE2="%vcpkg%" ^
    OPENSSL="%vcpkg%" ^
    all
```

- For 32-bit builds:

```
:: Add MinGW and Go to the system `PATH` variable for the current session:
set PATH=C:\Zabbix\mingw32\bin;%PATH%
set PATH=C:\Zabbix\Go\bin;%PATH%

:: Set vcpkg installation path:
set vcpkg="C:\Zabbix\x86\x86-mingw-static"

:: Set linker flags for Crypt32 library:
SET CGO_LDFLAGS="-lCrypt32"

:: Run the build process:
mingw32-make GOFLAGS="-buildvcs=false" ARCH=x86 ^
    PCRE2="%vcpkg%" ^
    OPENSSL="%vcpkg%" ^
    all
```

7. Compile Zabbix agent 2 by executing the script:

```
build.bat
```

After compilation, the Zabbix agent 2 binary will be located in `C:\Zabbix\zabbix-8.2.0\bin\win64` (for 64-bit builds) or `C:\Zabbix\zabbix-8.2.0\bin\win32` (for 32-bit builds). Zabbix agent 2 configuration files are located in `C:\Zabbix\zabbix-8.2.0\s`

To run the agent, `zabbix_agent2.exe` and its configuration files to a dedicated folder (e.g., `C:\Zabbix\agent2`) and then run the agent:

```
mkdir C:\Zabbix\agent2

#### For 64-bit builds:
copy C:\Zabbix\zabbix-8.2.0\bin\win64\zabbix_agent2.exe C:\Zabbix\agent2\

#### For 32-bit builds:
copy C:\Zabbix\zabbix-8.2.0\bin\win32\zabbix_agent2.exe C:\Zabbix\agent2\
```

```
copy C:\Zabbix\zabbix-8.2.0\src\go\conf\zabbix_agent2.conf C:\Zabbix\agent2\
xcopy /E /I C:\Zabbix\zabbix-8.2.0\src\go\conf\zabbix_agent2.d C:\Zabbix\agent2\zabbix_agent2.d\

C:\Zabbix\agent2\zabbix_agent2.exe -c C:\Zabbix\agent2\zabbix_agent2.conf
```

If necessary, continue with compiling Zabbix agent 2 loadable plugins.

Compiling Zabbix agent 2 loadable plugins

1. Download the Zabbix plugin source that matches your Zabbix agent 2 version (e.g., `zabbix-agent2-plugin-postgresql-8.2.0.tar.`) and extract it to `C:\Zabbix`.

> **Attention:**
> Before installing a plugin, please check its README file. It may contain specific requirements and installation instructions.

2. Navigate to the extracted plugin directory and compile the plugin:

```
cd C:\Zabbix\zabbix-agent2-plugin-ember-plus-8.2.0

#### For 64-bit builds:
mingw32-make ARCH=AMD64

#### For 32-bit builds:
mingw32-make ARCH=x86
```

After compilation, the Zabbix agent 2 plugin binary and its configuration file will be located in the same plugin directory.

The plugin executable may be placed anywhere as long as it is loadable by Zabbix agent 2. Specify the path to the plugin binary in the plugin configuration file, e.g., in postgresql.conf for the PostgreSQL plugin:

```
Plugins.PostgreSQL.System.Path=/path/to/executable/zabbix-agent2-plugin-postgresql
```

The path to the plugin configuration file must be specified in the Include parameter of the Zabbix agent 2 configuration file:

```
Include=/path/to/plugin/configuration/file/postgresql.conf
```

Proceed to setup for more details on configuring plugins.

Building Zabbix agent 2 manually

> **Attention:**
> This method of building Zabbix agent 2 is suitable for users who require full control over the build environment or are in a restricted environment where using vcpkg is not possible.

This section contains instructions for building Zabbix agent 2 manually, which includes installing the required build tools and dependencies, and then compiling the agent.

Setting up build tools

1. Download and install MSYS2 (available as an MSI installer). During installation, make sure to specify `C:\Zabbix\msys64` as the installation folder.

2. Download and install Go (available as an MSI installer; see currently supported Go versions). During installation, make sure to specify `C:\Zabbix\Go` as the installation folder.

3. Download the MinGW distribution that uses the Microsoft Visual C runtime library; for example:

- For 64-bit builds: `x86_64-15.1.0-release-win32-seh-msvcrt-rt_v12-rev0.7z`
- For 32-bit builds: `i686-15.1.0-release-win32-dwarf-msvcrt-rt_v12-rev0.7z`

> **Attention:**
> The MinGW version must meet the minimum requirements.

Then, extract it to `C:\Zabbix\mingw64` (or `C:\Zabbix\mingw32` for 32-bit builds).

Installing OpenSSL

Note:
To compile Zabbix agent without TLS support, proceed to the Installing PCRE2 section.

1. Open the MSYS2 MSYS terminal with administrator privileges and run the following commands:

```
pacman -S perl-Locale-Maketext-Simple
pacman -S nasm
pacman -S make
pacman -S cmake
```

2. Download the OpenSSL source archive and extract it to `C:\Zabbix\openssl-3.5.0`.

3. Navigate to the extracted OpenSSL directory and create the following `build.sh` script:

- For 64-bit builds:

```
####!/usr/bin/env bash

export PATH="/c/Zabbix/mingw64/bin:$PATH"
export d="/c/Zabbix/x64/OpenSSL-Win64-350-static"

perl Configure mingw64 no-shared no-capieng no-winstore no-srp no-gost no-dgram no-dtls1-method no-dtls1_2

make
make install
```

- For 32-bit builds:

```
####!/usr/bin/env bash

export PATH="/c/Zabbix/mingw32/bin:$PATH"
export d="/c/Zabbix/x86/OpenSSL-Win64-350-static"

perl Configure mingw no-shared no-capieng no-winstore no-srp no-gost no-dgram no-dtls1-method no-dtls1_2-m

make
make install
```

Attention:
Make sure to revoke write access from non-administrator users to the `C:\Zabbix\x64\OpenSSL-Win64-350-static` directory. Otherwise, the agent will load SSL settings from a path that can be modified by unprivileged users, resulting in a potential security vulnerability.

- The `no-shared` option makes libcrypto.lib and libssl.lib OpenSSL static libraries self-contained, so Zabbix binaries include OpenSSL without needing external DLLs. This means that Zabbix binaries can be copied to other Windows machines without OpenSSL libraries; however, when a new OpenSSL bugfix version is released, Zabbix agent will need to be recompiled.
- Without the `no-shared` option, Zabbix relies on OpenSSL DLLs at runtime. This means that OpenSSL updates may not require recompiling Zabbix agent; however, when copying it to other machines, the OpenSSL DLLs must be also be copied.

For more information about other OpenSSL configuration options, refer to OpenSSL documentation.

4. Configure and install OpenSSL by executing the script (note that this may take some time):

```
cd /c/Zabbix/openssl-3.5.0
./build.sh
```

Installing PCRE2

1. Download PCRE2 source archive and extract it to `C:\Zabbix\pcre2-10.45`.

2. Open the MSYS2 MSYS terminal with administrator privileges. Then, create a `build` directory in the extracted PCRE2 directory and navigate to it:

```
mkdir /c/Zabbix/pcre2-10.45/build
cd /c/Zabbix/pcre2-10.45/build
```

3. Configure PCRE2:

```
#### For 64-bit builds:
export PATH="/c/Zabbix/mingw64/bin:$PATH"
cmake -DCMAKE_C_COMPILER=gcc -DCMAKE_C_FLAGS="-O2 -g" -DCMAKE_INSTALL_PREFIX="/c/Zabbix/x64/PCRE2" ..

#### For 32-bit builds:
export PATH="/c/Zabbix/mingw32/bin:$PATH"
cmake -DCMAKE_C_COMPILER=gcc -DCMAKE_C_FLAGS="-m32 -O2 -g" -DCMAKE_EXE_LINKER_FLAGS="-Wl,-mi386pe" -DCMAKE
```

Note:
If any errors occur, it is recommended to delete the CMake cache before attempting to repeat the CMake build process. The cache (`CMakeCachecache.txt`) can be located in the build directory of the extracted PCRE2 directory.

4. Install PCRE2:

```
make install
```

Compiling Zabbix agent 2

1. Download the Zabbix source archive and extract it to `C:\Zabbix\zabbix-8.2.0`.

If you need to generate a source archive from the raw source repository (e.g., to apply custom patches or build from the latest source code), run the following commands on a **Linux** machine with Go installed (required for configuring Zabbix agent 2):

```
git clone https://git.zabbix.com/scm/zbx/zabbix.git
cd zabbix
./bootstrap.sh
./configure
make dist
```

This will create a source archive, which can then be copied to a Windows machine.

2. Open the Command Prompt with administrator privileges. Then, navigate to the Zabbix build directory and compile Zabbix agent; make sure to correctly specify the directories where OpenSSL and PCRE2 are installed:

- For 64-bit builds:

```
cd C:\Zabbix\zabbix-8.2.0\build\mingw
set PATH=C:\Zabbix\mingw64\bin;%PATH%
mklink /D C:\Zabbix\x64\OpenSSL-Win64-350-static\lib C:\Zabbix\x64\OpenSSL-Win64-350-static\lib64

#### With TLS support:
mingw32-make ARCH=AMD64 PCRE2="C:\Zabbix\x64\PCRE2" OPENSSL="C:\Zabbix\x64\OpenSSL-Win64-350-static"

#### Without TLS support:
mingw32-make ARCH=AMD64 PCRE2="C:\Zabbix\x64\PCRE2"
```

- For 32-bit builds:

```
cd C:\Zabbix\zabbix-8.2.0\build\mingw
set PATH=C:\Zabbix\mingw32\bin;%PATH%

#### With TLS support:
mingw32-make ARCH=x86 PCRE2="C:\Zabbix\x86\PCRE2" OPENSSL="C:\Zabbix\x86\OpenSSL-Win32-350-static"

#### Without TLS support:
mingw32-make ARCH=x86 PCRE2="C:\Zabbix\x86\PCRE2"
```

After compilation, the Zabbix agent 2 binary will be located in `C:\Zabbix\zabbix-8.2.0\bin\win64` (or `C:\Zabbix\zabbix-8.2.0\bin` for 32-bit builds). Zabbix agent 2 configuration files are located in `C:\Zabbix\zabbix-8.2.0\src\go\conf`.

To run the agent, copy the `zabbix_agent2.exe` binary and its configuration files to a dedicated folder (e.g., `C:\Zabbix\agent2`) and then run the agent:

```
mkdir C:\Zabbix\agent2
copy C:\Zabbix\zabbix-8.2.0\bin\win64\zabbix_agent2.exe C:\Zabbix\agent2\
copy C:\Zabbix\zabbix-8.2.0\src\go\conf\zabbix_agent2.conf C:\Zabbix\agent2\
xcopy /E /I C:\Zabbix\zabbix-8.2.0\src\go\conf\zabbix_agent2.d C:\Zabbix\agent2\zabbix_agent2.d\
```

```
C:\Zabbix\agent2\zabbix_agent2.exe -c C:\Zabbix\agent2\zabbix_agent2.conf
```

If necessary, continue with compiling Zabbix agent 2 loadable plugins.

Building Zabbix agent on Windows

Overview

This page demonstrates how to build Zabbix agent from source on Windows 10 (64-bit).

These instructions apply to Windows versions that support Visual Studio 2022.

Building Zabbix agent requires:

- C compiler (included in Build Tools for Visual Studio 2022)
- OpenSSL (for encryption features in Zabbix)
- PCRE2 (Perl Compatible Regular Expressions; for regular expressions pattern matching features in Zabbix)

You can build Zabbix agent using one of the following methods:

- Using vcpkg - an automated approach that simplifies dependency management using a C++ package manager.
- Manual build - a manual approach that requires installing all dependencies before compiling the agent.

Depending on your monitoring needs, additional libraries may be required. For more information, see Requirements.

Attention:
Before starting the build process, please keep in mind:

- To execute commands, use the "x64 Native Tools Command Prompt" (included in Build Tools for Visual Studio 2022), launched by a user with sufficient permissions to write to protected folders.
- It is recommended to create a working directory at `C:\Zabbix` for all source files and build folders. However, compiled components should be installed in `C:\Program Files\Zabbix\x64`.

Building Zabbix agent with vcpkg

This section contains instructions for building Zabbix agent with vcpkg, a package manager that simplifies dependency management and integration with C++ projects.

1. Download and install Build Tools for Visual Studio 2022. During installation, make sure to select the *Desktop development with C++* workload, which includes the tools required for building the agent with vcpkg:

- C compiler (Microsoft Visual C++)
- NMake command-line tool
- vcpkg package manager
- "x64 Native Tools Command Prompt"

2. Open `cmd.exe` with administrator privileges, initialize vcpkg, and install the dependencies required for building Zabbix agent (note that this may take some time):

```
cd C:\Zabbix
vcpkg new --application
vcpkg add port pcre2
vcpkg add port openssl
vcpkg install --triplet x64-windows-static --x-install-root="C:\Program Files\Zabbix\x64"
```

3. Download the Zabbix source archive and extract it to `C:\Zabbix\zabbix-8.2.0`.

4. Navigate to the Zabbix build directory (`C:\Zabbix\zabbix-8.2.0\build\win32\project`) and create the following `build.bat` script; make sure to correctly specify the directories where OpenSSL and PCRE2 are installed:

```
:: Set vcpkg installation path:
set vcpkg=C:\Program Files\Zabbix\x64\x64-windows-static

:: Run the build process:
nmake -f Makefile CPU=AMD64 ^
    PCRE2INCDIR="%vcpkg%\include" ^
    PCRE2LIBDIR="%vcpkg%\lib" ^
    TLS=openssl ^
    TLSINCDIR="%vcpkg%\include" ^
```

```
    TLSLIBDIR="%vcpkg%\lib" ^
    LIBS="$(LIBS) Crypt32.lib" ^
    all
```

5. Open the "x64 Native Tools Command Prompt" and compile Zabbix agent by executing the script:

```
cd C:\Zabbix\zabbix-8.2.0\build\win32\project
build.bat
```

If you opened a different command prompt, initialize the Visual Studio build environment first:

```
"%ProgramFiles(x86)%\Microsoft Visual Studio\2022\BuildTools\VC\Auxiliary\Build\vcvarsall.bat" x64
cd C:\Zabbix\zabbix-8.2.0\build\win32\project
build.bat
```

Wrap paths that contain spaces in quotes.

After compilation, Zabbix component binaries will be located in `C:\Zabbix\zabbix-8.2.0\bin\win64`. Zabbix agent configuration file is located in `C:\Zabbix\zabbix-8.2.0\conf`.

To run the agent, copy `zabbix_agentd.exe` and its configuration file to a dedicated folder (e.g., `C:\Zabbix\agent`) and then run the agent:

```
mkdir C:\Zabbix\agent
copy C:\Zabbix\zabbix-8.2.0\bin\win64\zabbix_agentd.exe C:\Zabbix\agent\
copy C:\Zabbix\zabbix-8.2.0\conf\zabbix_agentd.conf C:\Zabbix\agent\

C:\Zabbix\agent\zabbix_agentd.exe -c C:\Zabbix\agent\zabbix_agentd.conf
```

Building Zabbix agent manually

> **Attention:**
> This method of building Zabbix agent is suitable for users who require full control over the build environment or are in a restricted environment where using vcpkg is not possible.

This section contains instructions for building Zabbix agent manually, which includes installing the required build tools and dependencies (Perl, OpenSSL, PCRE2), and then compiling the agent.

Installing build tools

1. Download and install Build Tools for Visual Studio 2022. During installation, make sure to select the *Desktop development with C++* workload, which includes the tools required for building the agent manually:

- C compiler (Microsoft Visual C++)
- NMake command-line tool
- "x64 Native Tools Command Prompt"

Installing OpenSSL

> **Note:**
> To compile Zabbix agent without TLS support, proceed to the Installing PCRE2 section.

1. Download and install Strawberry Perl (available as an MSI installer). During installation, make sure to specify `C:\Zabbix\Strawberry` as the installation folder.

2. Install the `Text::Template` Perl module:

```
cpanm Text::Template
```

3. Verify that the Netwide Assembler (NASM; required for compiling OpenSSL), was compiled during the installation of Strawberry Perl:

```
nasm -v
#### NASM version 2.16.01 compiled on May  3 2024
```

If NASM is not compiled, install it manually. For more information, refer to NASM documentation.

4. Download the OpenSSL source archive and extract it to `C:\Zabbix\openssl-3.5.0`.

5. Navigate to the extracted directory and configure OpenSSL, for example:

```
cd C:\Zabbix\openssl-3.5.0
perl Configure VC-WIN64A no-shared no-capieng no-winstore no-srp no-gost no-dgram no-dtls1-method no-dtls1
```

Attention:
If you choose a custom directory for OpenSSL when compiling Zabbix agent on Windows (e.g., `C:\zabbix` or `C:\openssl-64bit`), make sure to revoke write access from non-administrator users to this directory. Otherwise, the agent will load SSL settings from a path that can be modified by unprivileged users, resulting in a potential security vulnerability.

- The `no-shared` option makes libcrypto.lib and libssl.lib OpenSSL static libraries self-contained, so Zabbix binaries include OpenSSL without needing external DLLs. This means that Zabbix binaries can be copied to other Windows machines without OpenSSL libraries; however, when a new OpenSSL bugfix version is released, Zabbix agent will need to be recompiled.
- Without the `no-shared` option, Zabbix relies on OpenSSL DLLs at runtime. This means that OpenSSL updates may not require recompiling Zabbix agent; however, when copying it to other machines, the OpenSSL DLLs must be also be copied.

For more information about other OpenSSL configuration options, refer to OpenSSL documentation.

6. Compile OpenSSL and run tests (note that this may take some time):

Attention:
Run the tests without administrative privileges; otherwise, it may lead to unexpected results or security risks. If some tests fail, refer to OpenSSL documentation for troubleshooting.

```
nmake
nmake test
...
All tests successful.
Files=325, Tests=3101, 822 wallclock secs ( 4.81 usr +  0.81 sys =  5.62 CPU)
Result: PASS
```

7. Install OpenSSL:

```
nmake install
```

To install only software components (libraries, header files, but no documentation), you may use `nmake install_sw`.

Installing PCRE2

1. Download and install CMake (available as an MSI installer). During installation, make sure to specify `C:\Zabbix\CMake` as the installation folder and select the *Add CMake to the PATH environment variable* option.

2. Download the PCRE2 source archive and extract it to `C:\Zabbix\pcre2-10.45`.

3. Create a `build` directory in the extracted PCRE2 directory and navigate to it:

```
mkdir C:\Zabbix\pcre2-10.45\build
cd C:\Zabbix\pcre2-10.45\build
```

4. Configure PCRE2:

```
cmake -G "NMake Makefiles" -DPCRE_SUPPORT_UNICODE_PROPERTIES=ON -DCMAKE_BUILD_TYPE=Release -DCMAKE_INSTALL
```

Note:
If any errors occur, it is recommended to delete the CMake cache before attempting to repeat the CMake build process. The cache (`CMakeCachecache.txt`) can be located in the build directory of the extracted PCRE2 directory.

5. Build PCRE2 using NMake:

```
nmake
```

6. Install PCRE2:

```
cmake --install .
```

Compiling Zabbix agent

1. Download the Zabbix source archive and extract it to `C:\Zabbix\zabbix-8.2.0`.

If you need to generate a source archive from the raw source repository (e.g., to apply custom patches or build from the latest source code), run the following commands on a **Linux** machine:

```
git clone https://git.zabbix.com/scm/zbx/zabbix.git
cd zabbix
./bootstrap.sh
./configure
make dist
```

This will create a source archive, which can then be copied to a Windows machine.

2. Navigate to the Zabbix build directory, open the "x64 Native Tools Command Prompt", and compile Zabbix agent (or other components); make sure to correctly specify the directories where OpenSSL and PCRE2 are installed:

```
"%ProgramFiles(x86)%\Microsoft Visual Studio\2022\BuildTools\VC\Auxiliary\Build\vcvarsall.bat" x64

cd C:\Zabbix\zabbix-8.2.0\build\win32\project

#### With TLS support:
nmake /K -f Makefile_agent PCRE2INCDIR="C:\Program Files\Zabbix\x64\PCRE2\include" PCRE2LIBDIR="C:\Program
nmake /K -f Makefile_get PCRE2INCDIR="C:\Program Files\Zabbix\x64\PCRE2\include" PCRE2LIBDIR="C:\Program F
nmake /K -f Makefile_sender PCRE2INCDIR="C:\Program Files\Zabbix\x64\PCRE2\include" PCRE2LIBDIR="C:\Progra

#### Without TLS support:
nmake /K -f Makefile_agent PCRE2INCDIR="C:\Program Files\Zabbix\x64\PCRE2\include" PCRE2LIBDIR="C:\Program
nmake /K -f Makefile_get PCRE2INCDIR="C:\Program Files\Zabbix\x64\PCRE2\include" PCRE2LIBDIR="C:\Program F
nmake /K -f Makefile_sender PCRE2INCDIR="C:\Program Files\Zabbix\x64\PCRE2\include" PCRE2LIBDIR="C:\Progra
```

Wrap paths that contain spaces in quotes.

After compilation, Zabbix component binaries will be located in `C:\Zabbix\zabbix-8.2.0\bin\win64`. Zabbix agent configuration file is located in `C:\Zabbix\zabbix-8.2.0\conf`.

To run the agent, copy `zabbix_agentd.exe` and its configuration file to a dedicated folder (e.g., `C:\Zabbix\agent`) and then run the agent:

```
mkdir C:\Zabbix\agent
copy C:\Zabbix\zabbix-8.2.0\bin\win64\zabbix_agentd.exe C:\Zabbix\agent\
copy C:\Zabbix\zabbix-8.2.0\conf\zabbix_agentd.conf C:\Zabbix\agent\

C:\Zabbix\agent\zabbix_agentd.exe -c C:\Zabbix\agent\zabbix_agentd.conf -f
```

macOS agent installation from PKG

Overview

Zabbix agent can be installed on macOS using PKG installer packages, available for download.

Zabbix agent packages are available with or without encryption.

Installing agent

The agent can be installed using the graphical user interface or from the command line, for example:

```
sudo installer -pkg zabbix_agent-8.2.0-macos-arm64-openssl.pkg -target /
```

Make sure to use the correct Zabbix package version in the command. It must match the name of the downloaded package.

Running agent

The agent will start automatically after installation or restart.

You may edit the configuration file at `/usr/local/etc/zabbix/zabbix_agentd.conf` if necessary.

To start the agent manually, you may run:

```
sudo launchctl start com.zabbix.zabbix_agentd
```

To stop the agent manually:

```
sudo launchctl stop com.zabbix.zabbix_agentd
```

During upgrade, the existing configuration file is not overwritten. Instead a new `zabbix_agentd.conf.NEW` file is created to be used for reviewing and updating the existing configuration file, if necessary. Remember to restart the agent after manual changes to the configuration file.

Troubleshooting and removing agent

This section lists some useful commands that can be used for troubleshooting and removing Zabbix agent installation.

See if Zabbix agent is running:

```
ps aux | grep zabbix_agentd
```

See if Zabbix agent has been installed from packages:

```
pkgutil --pkgs | grep zabbix
com.zabbix.pkg.ZabbixAgent
```

See the files that were installed from the installer package (note that the initial / is not displayed in this view):

```
pkgutil --only-files --files com.zabbix.pkg.ZabbixAgent
Library/LaunchDaemons/com.zabbix.zabbix_agentd.plist
usr/local/bin/zabbix_get
usr/local/bin/zabbix_sender
usr/local/etc/zabbix/zabbix_agentd/userparameter_examples.conf.NEW
usr/local/etc/zabbix/zabbix_agentd/userparameter_mysql.conf.NEW
usr/local/etc/zabbix/zabbix_agentd.conf.NEW
usr/local/sbin/zabbix_agentd
```

Stop Zabbix agent if it was launched with `launchctl`:

```
sudo launchctl unload /Library/LaunchDaemons/com.zabbix.zabbix_agentd.plist
```

Remove files (including configuration and logs) that were installed with installer package:

```
sudo rm -f /Library/LaunchDaemons/com.zabbix.zabbix_agentd.plist
sudo rm -f /usr/local/sbin/zabbix_agentd
sudo rm -f /usr/local/bin/zabbix_get
sudo rm -f /usr/local/bin/zabbix_sender
sudo rm -rf /usr/local/etc/zabbix
sudo rm -rf /var/log/zabbix
```

Forget that Zabbix agent has been installed:

```
sudo pkgutil --forget com.zabbix.pkg.ZabbixAgent
```

Building Zabbix agent on macOS

Overview

This section demonstrates how to build Zabbix macOS agent binaries from sources with or without TLS.

Prerequisites

You will need command line developer tools (Xcode is not required), Automake, pkg-config and PCRE (v8.x) or PCRE2 (v10.x). If you want to build agent binaries with TLS, you will also need OpenSSL or GnuTLS.

To install Automake and pkg-config, you will need a Homebrew package manager from https://brew.sh/. To install it, open terminal and run the following command:

```
/bin/bash -c "$(curl -fsSL https://raw.githubusercontent.com/Homebrew/install/HEAD/install.sh)"
```

Then install Automake and pkg-config:

```
brew install automake
brew install pkg-config
```

Preparing PCRE, OpenSSL and GnuTLS libraries depends on the way how they are going to be linked to the agent.

If you intend to run agent binaries on a macOS machine that already has these libraries, you can use precompiled libraries that are provided by Homebrew. These are typically macOS machines that use Homebrew for building Zabbix agent binaries or for other purposes.

If agent binaries will be used on macOS machines that don't have the shared version of libraries, you should compile static libraries from sources and link Zabbix agent with them.

Building agent binaries with shared libraries

Install PCRE2 (replace *pcre2* with *pcre* in the commands below, if needed):

```
brew install pcre2
```

When building with TLS, install OpenSSL and/or GnuTLS:

```
brew install openssl
brew install gnutls
```

Download Zabbix source:

```
git clone https://git.zabbix.com/scm/zbx/zabbix.git
```

Build agent without TLS:

```
cd zabbix
./bootstrap.sh
./configure --sysconfdir=/usr/local/etc/zabbix --enable-agent --enable-ipv6
make
make install
```

Build agent with OpenSSL:

```
cd zabbix
./bootstrap.sh
./configure --sysconfdir=/usr/local/etc/zabbix --enable-agent --enable-ipv6 --with-openssl=/usr/local/opt/
make
make install
```

Build agent with GnuTLS:

```
cd zabbix-source/
./bootstrap.sh
./configure --sysconfdir=/usr/local/etc/zabbix --enable-agent --enable-ipv6 --with-gnutls=/usr/local/opt/g
make
make install
```

Building agent binaries with static libraries without TLS

Let's assume that PCRE static libraries will be installed in `$HOME/static-libs`. We will use PCRE2 10.39.

```
PCRE_PREFIX="$HOME/static-libs/pcre2-10.39"
```

Download and build PCRE with Unicode properties support:

```
mkdir static-libs-source
cd static-libs-source
curl --remote-name https://github.com/PhilipHazel/pcre2/releases/download/pcre2-10.39/pcre2-10.39.tar.gz
tar xf pcre2-10.39.tar.gz
cd pcre2-10.39
./configure --prefix="$PCRE_PREFIX" --disable-shared --enable-static --enable-unicode-properties
make
make check
make install
```

Download Zabbix source and build agent:

```
git clone https://git.zabbix.com/scm/zbx/zabbix.git
cd zabbix
./bootstrap.sh
./configure --sysconfdir=/usr/local/etc/zabbix --enable-agent --enable-ipv6 --with-libpcre2="$PCRE_PREFIX"
make
make install
```

Building agent binaries with static libraries with OpenSSL

When building OpenSSL, it's recommended to run `make test` after successful building. Even if building was successful, tests sometimes fail. If this is the case, problems should be researched and resolved before continuing.

Let's assume that PCRE and OpenSSL static libraries will be installed in $HOME/static-libs. We will use PCRE2 10.39 and OpenSSL 1.1.1a.

```
PCRE_PREFIX="$HOME/static-libs/pcre2-10.39"
OPENSSL_PREFIX="$HOME/static-libs/openssl-1.1.1a"
```

Let's build static libraries in static-libs-source:

```
mkdir static-libs-source
cd static-libs-source
```

Download and build PCRE with Unicode properties support:

```
curl --remote-name https://github.com/PhilipHazel/pcre2/releases/download/pcre2-10.39/pcre2-10.39.tar.gz
tar xf pcre2-10.39.tar.gz
cd pcre2-10.39
./configure --prefix="$PCRE_PREFIX" --disable-shared --enable-static --enable-unicode-properties
make
make check
make install
cd ..
```

Download and build OpenSSL:

```
curl --remote-name https://www.openssl.org/source/openssl-1.1.1a.tar.gz
tar xf openssl-1.1.1a.tar.gz
cd openssl-1.1.1a
./Configure --prefix="$OPENSSL_PREFIX" --openssldir="$OPENSSL_PREFIX" --api=1.1.0 no-shared no-capieng no-
make
make test
make install_sw
cd ..
```

Download Zabbix source and build agent:

```
git clone https://git.zabbix.com/scm/zbx/zabbix.git
cd zabbix
./bootstrap.sh
./configure --sysconfdir=/usr/local/etc/zabbix --enable-agent --enable-ipv6 --with-libpcre2="$PCRE_PREFIX"
make
make install
```

Building agent binaries with static libraries with GnuTLS

GnuTLS depends on the Nettle crypto backend and GMP arithmetic library. Instead of using full GMP library, this guide will use mini-gmp which is included in Nettle.

When building GnuTLS and Nettle, it's recommended to run make check after successful building. Even if building was successful, tests sometimes fail. If this is the case, problems should be researched and resolved before continuing.

Let's assume that PCRE, Nettle and GnuTLS static libraries will be installed in $HOME/static-libs. We will use PCRE2 10.39, Nettle 3.4.1 and GnuTLS 3.6.5.

```
PCRE_PREFIX="$HOME/static-libs/pcre2-10.39"
NETTLE_PREFIX="$HOME/static-libs/nettle-3.4.1"
GNUTLS_PREFIX="$HOME/static-libs/gnutls-3.6.5"
```

Let's build static libraries in static-libs-source:

```
mkdir static-libs-source
cd static-libs-source
```

Download and build Nettle:

```
curl --remote-name https://ftp.gnu.org/gnu/nettle/nettle-3.4.1.tar.gz
tar xf nettle-3.4.1.tar.gz
cd nettle-3.4.1
./configure --prefix="$NETTLE_PREFIX" --enable-static --disable-shared --disable-documentation --disable-a
make
make check
make install
cd ..
```

Download and build GnuTLS:

```
curl --remote-name https://www.gnupg.org/ftp/gcrypt/gnutls/v3.6/gnutls-3.6.5.tar.xz
tar xf gnutls-3.6.5.tar.xz
cd gnutls-3.6.5
PKG_CONFIG_PATH="$NETTLE_PREFIX/lib/pkgconfig" ./configure --prefix="$GNUTLS_PREFIX" --enable-static --dis
make
make check
make install
cd ..
```

Download Zabbix source and build agent:

```
git clone https://git.zabbix.com/scm/zbx/zabbix.git
cd zabbix
./bootstrap.sh
CFLAGS="-Wno-unused-command-line-argument -framework Foundation -framework Security" \
> LIBS="-lgnutls -lhogweed -lnettle" \
> LDFLAGS="-L$GNUTLS_PREFIX/lib -L$NETTLE_PREFIX/lib" \
> ./configure --sysconfdir=/usr/local/etc/zabbix --enable-agent --enable-ipv6 --with-libpcre2="$PCRE_PREFI
make
make install
```

AWS deployment guide for Zabbix server

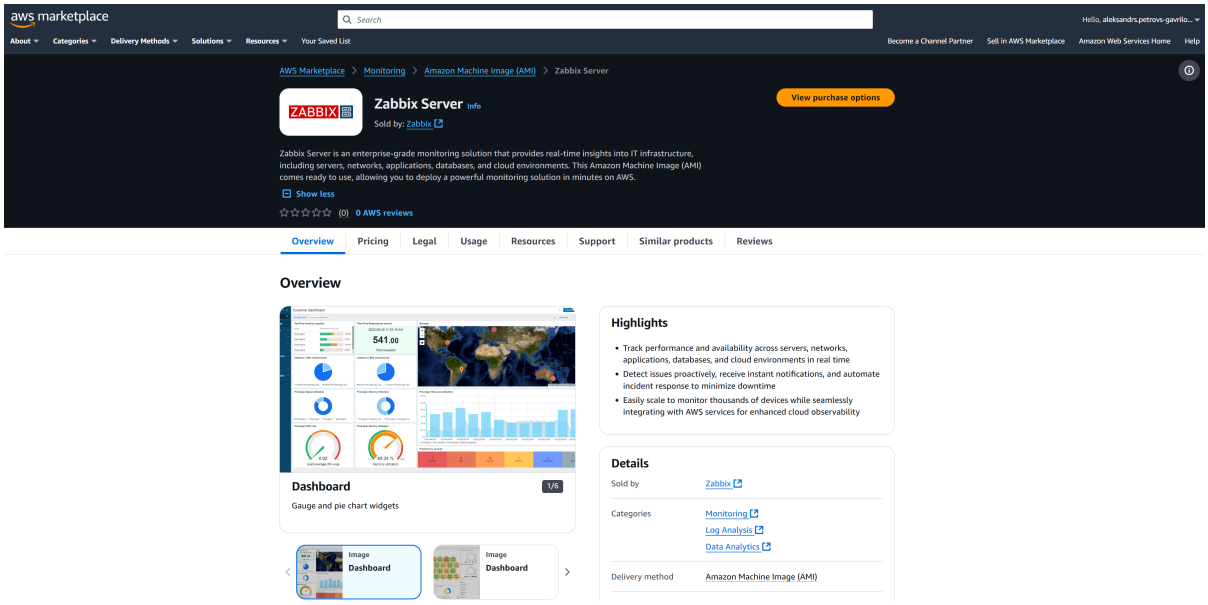
Overview

This page describes how to deploy Zabbix server using AWS.

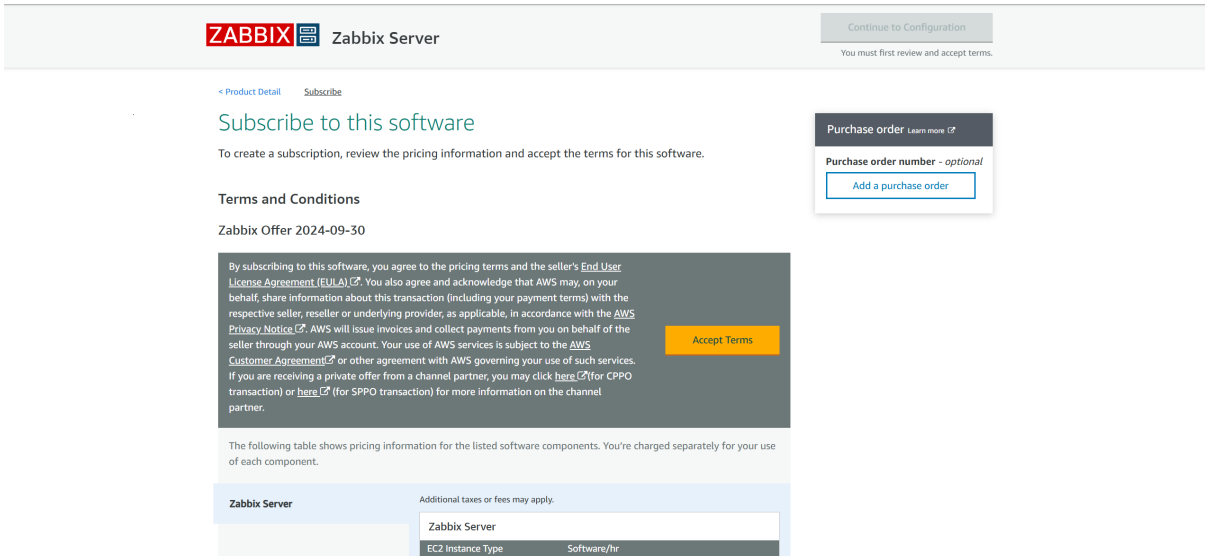
> **Attention:**
> Several Zabbix versions are available for deployment on the platform. Make sure to read the documentation for the same Zabbix version that you have deployed. Documentation for different Zabbix versions can be selected from the version dropdown menu.

Deploying the appliance

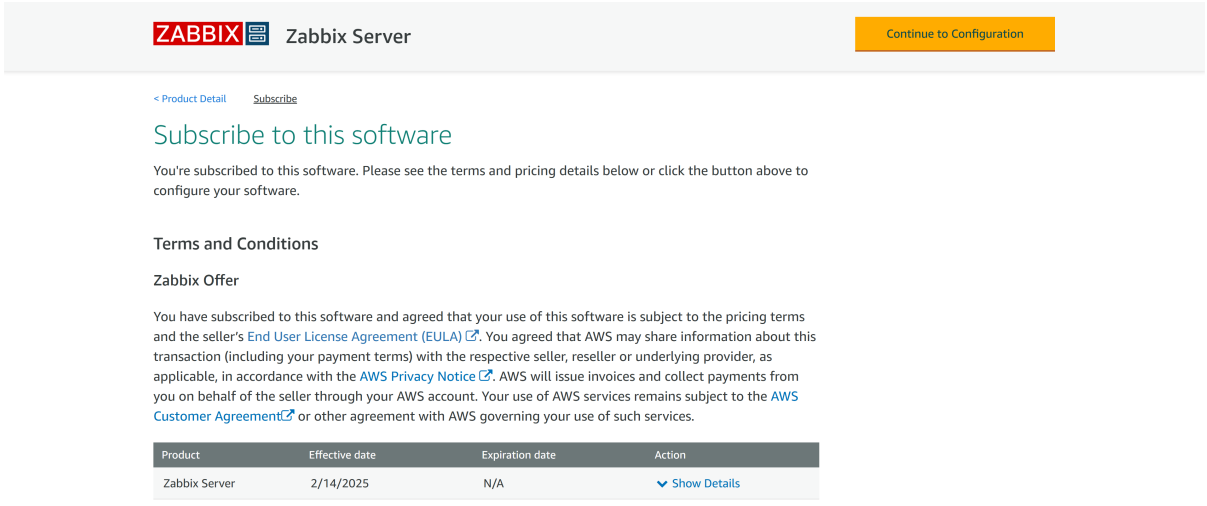
To deploy Zabbix appliance, first, let's navigate to the Zabbix AWS marketplace page. Next, click on the *View purchase options* to deploy the image and proceed to the subscription page.



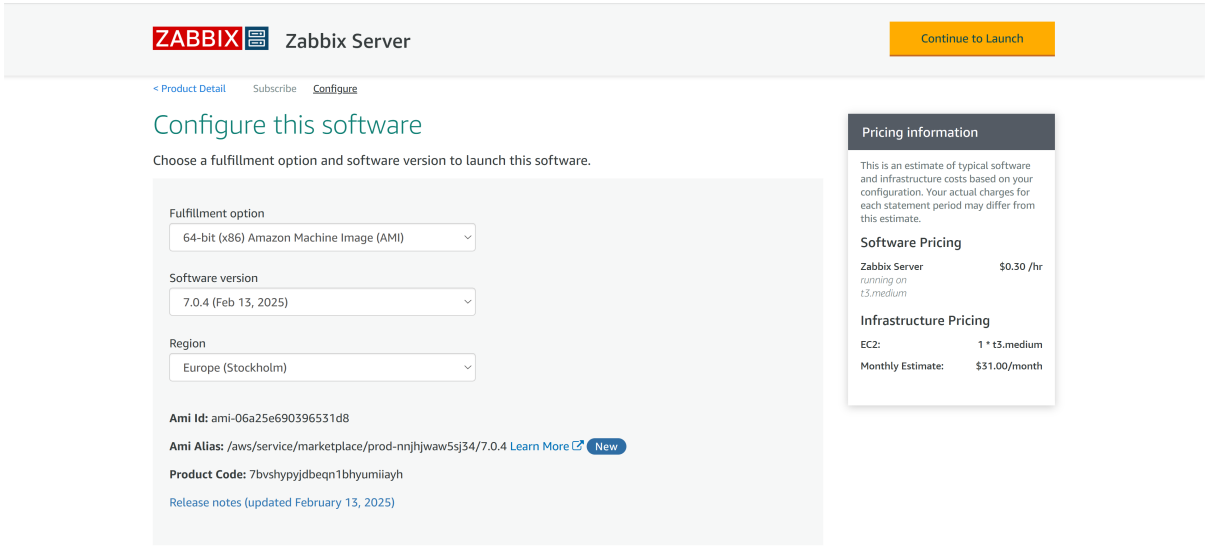
Click the *Accept terms* button to accept *Terms and Conditions*.



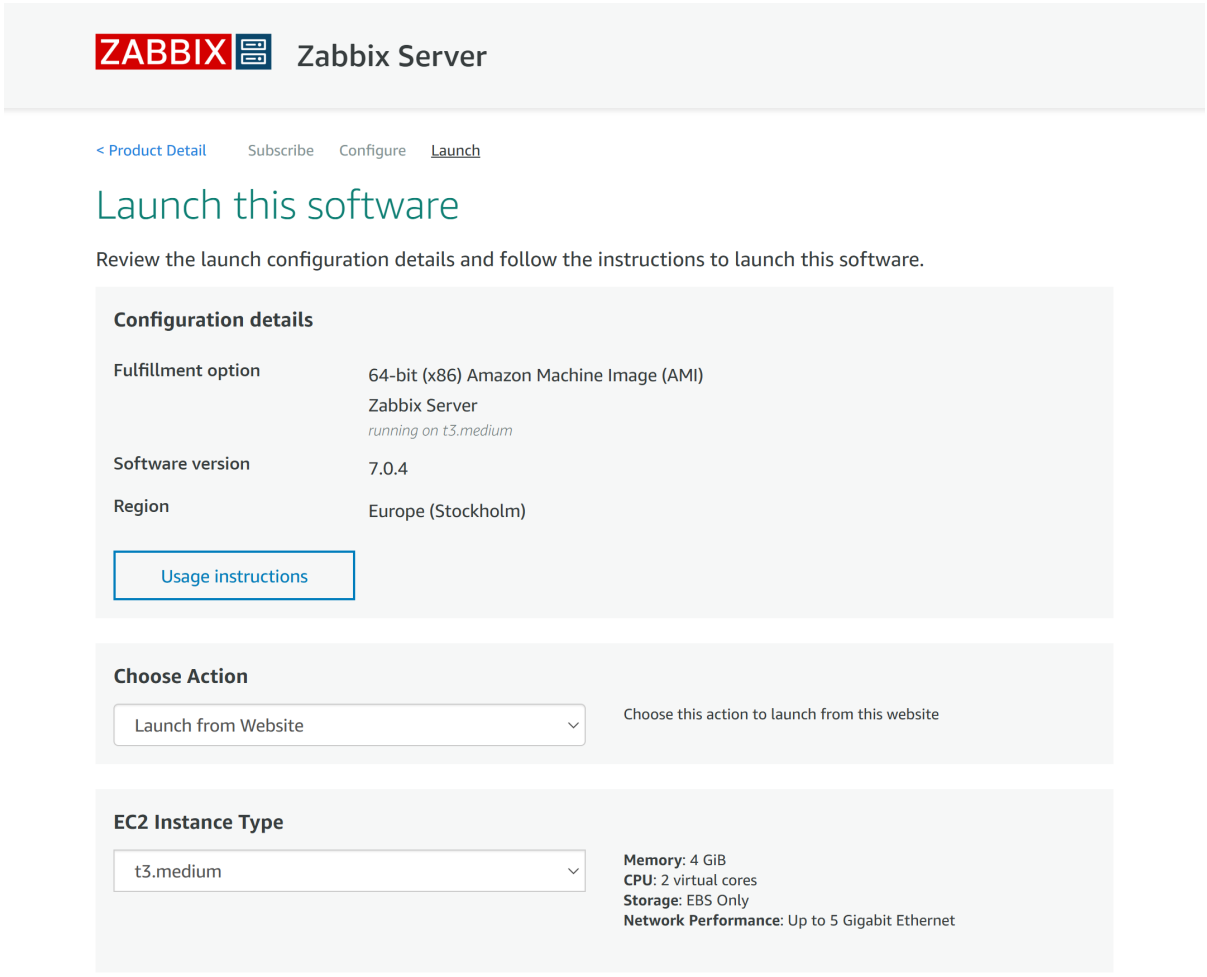
Click the *Continue to Configuration* button to move to the Configuration section for Zabbix appliance.



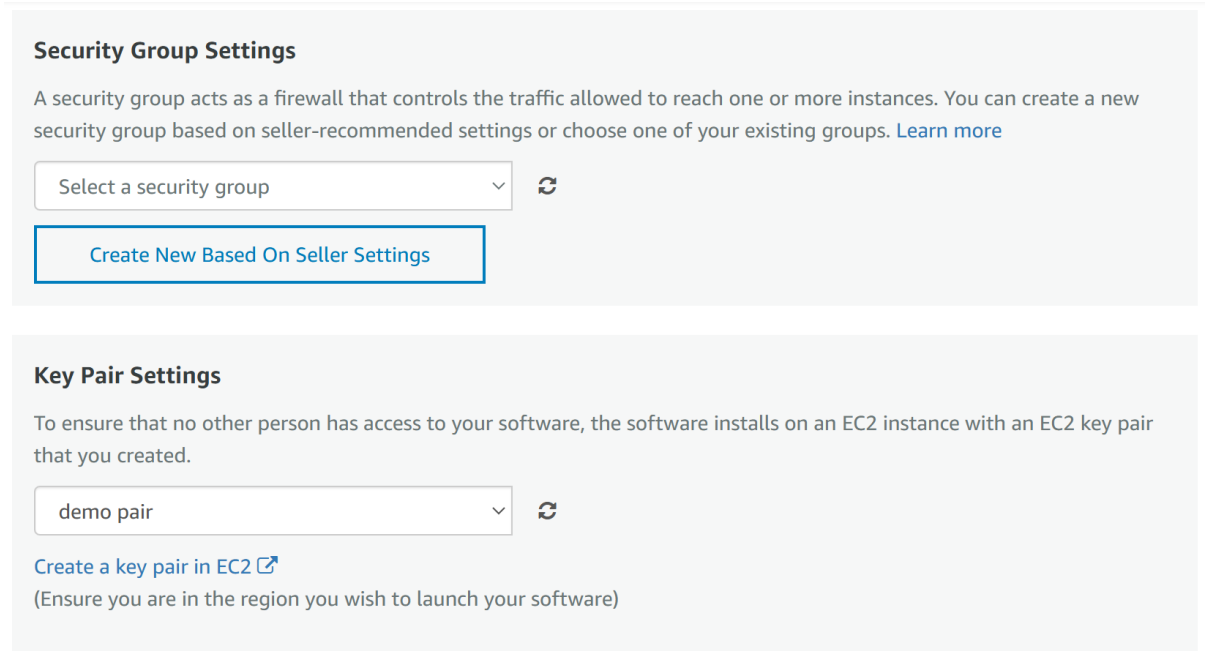
Select the Fulfilment option, Zabbix version, and your preferred Region.



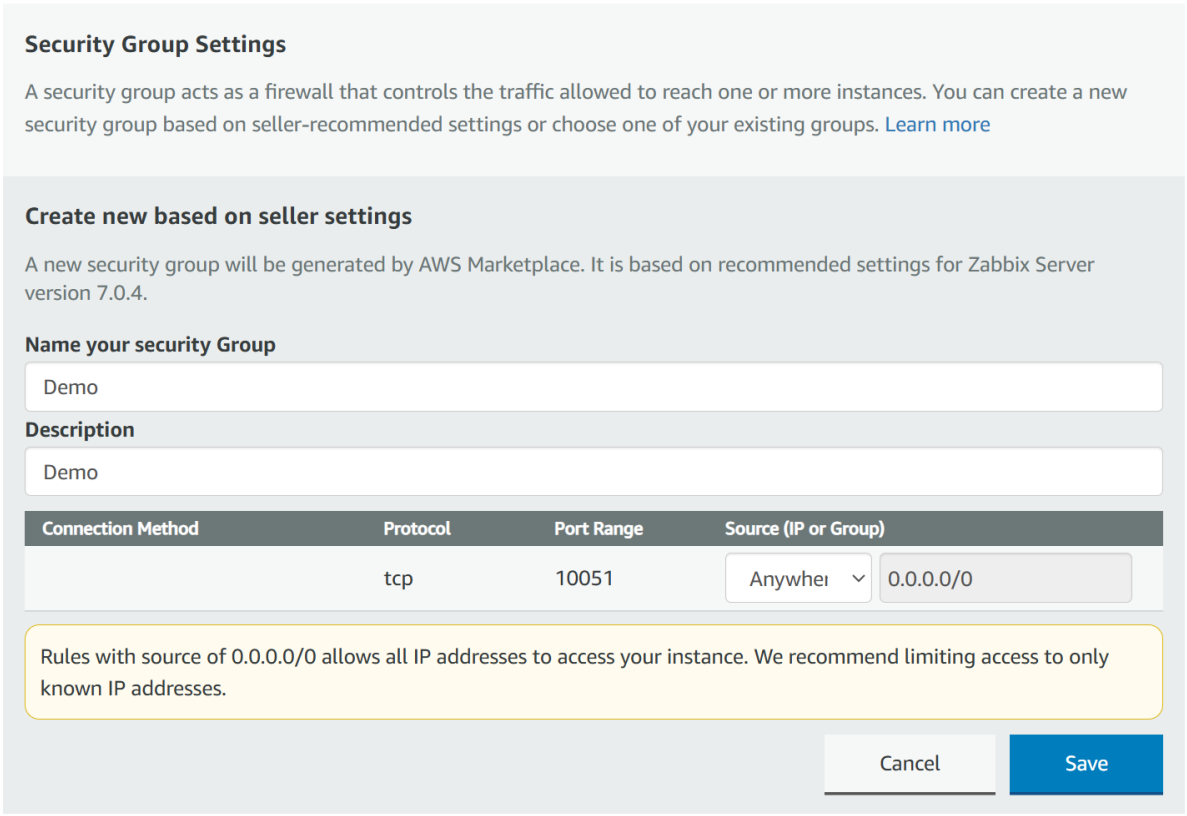
Click the *Continue to Launch* button. On the Launch page, select the preferred Launch option, EC2 Instance Type, VPC, and Subnet settings.



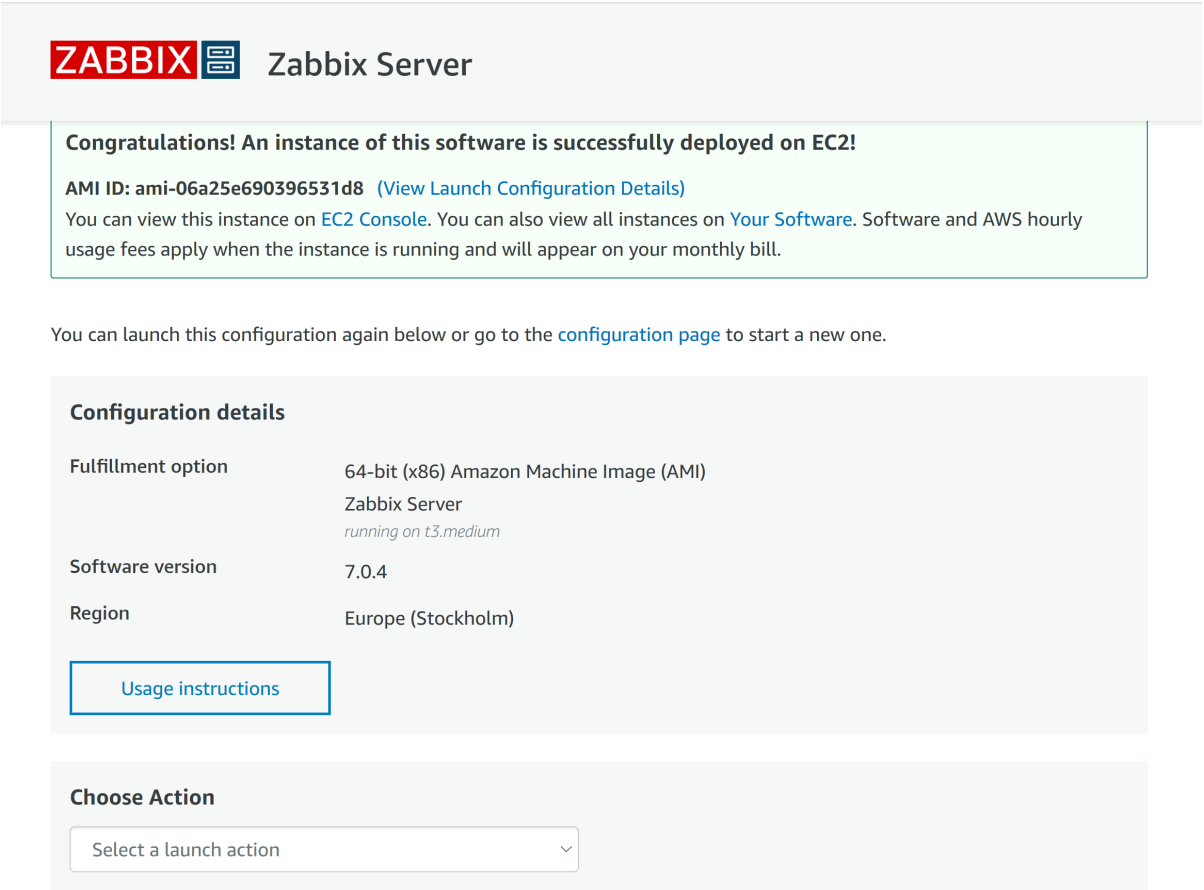
At the bottom of the Launch page, specify the Security Group settings and select or create an EC2 key pair - make sure to save your private key in a safe location!



If you select *Create a new Security Group based on seller settings*, TCP port 10051 (Zabbix server) will be allowed. Note: If you wish to access the instance via SSH and HTTP/HTTPS (Web frontend), you will also need to permit SSH and HTTP/HTTPS connections in your Security Group settings.

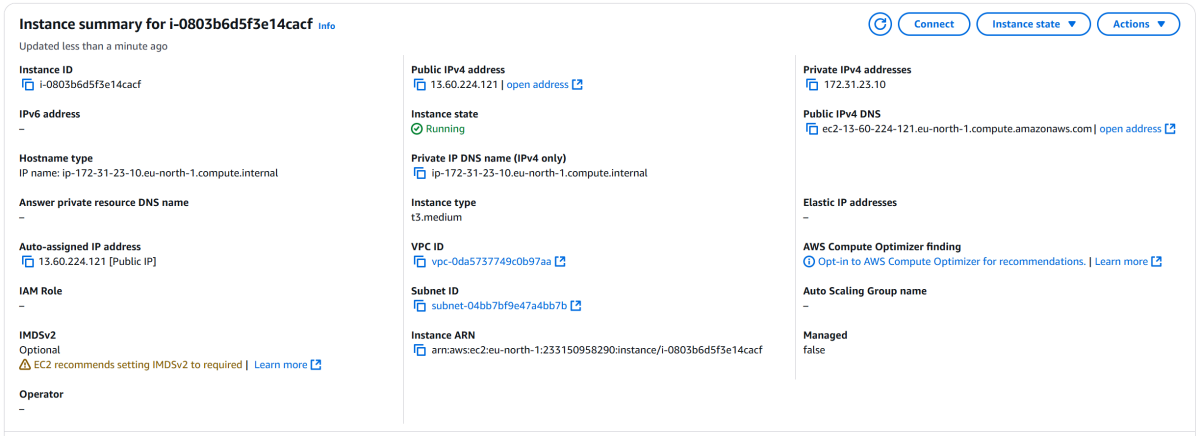


Congratulations, your Zabbix appliance has been successfully deployed.

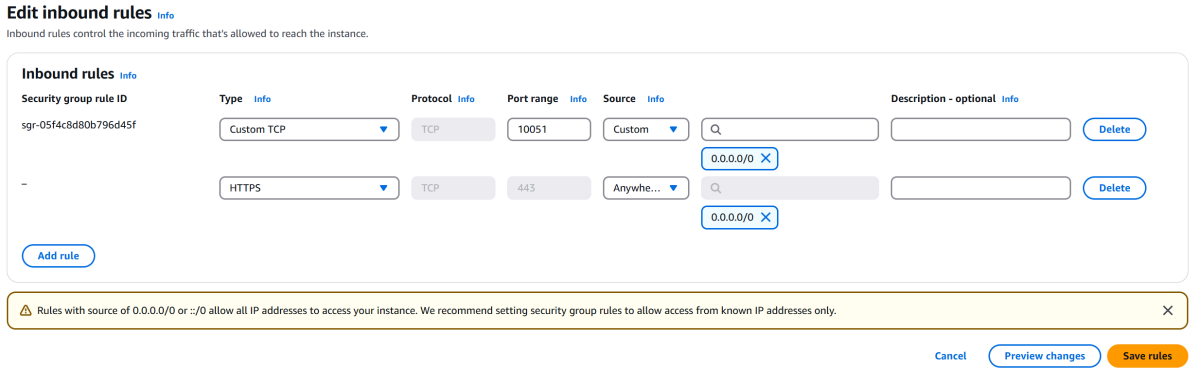


Accessing the appliance

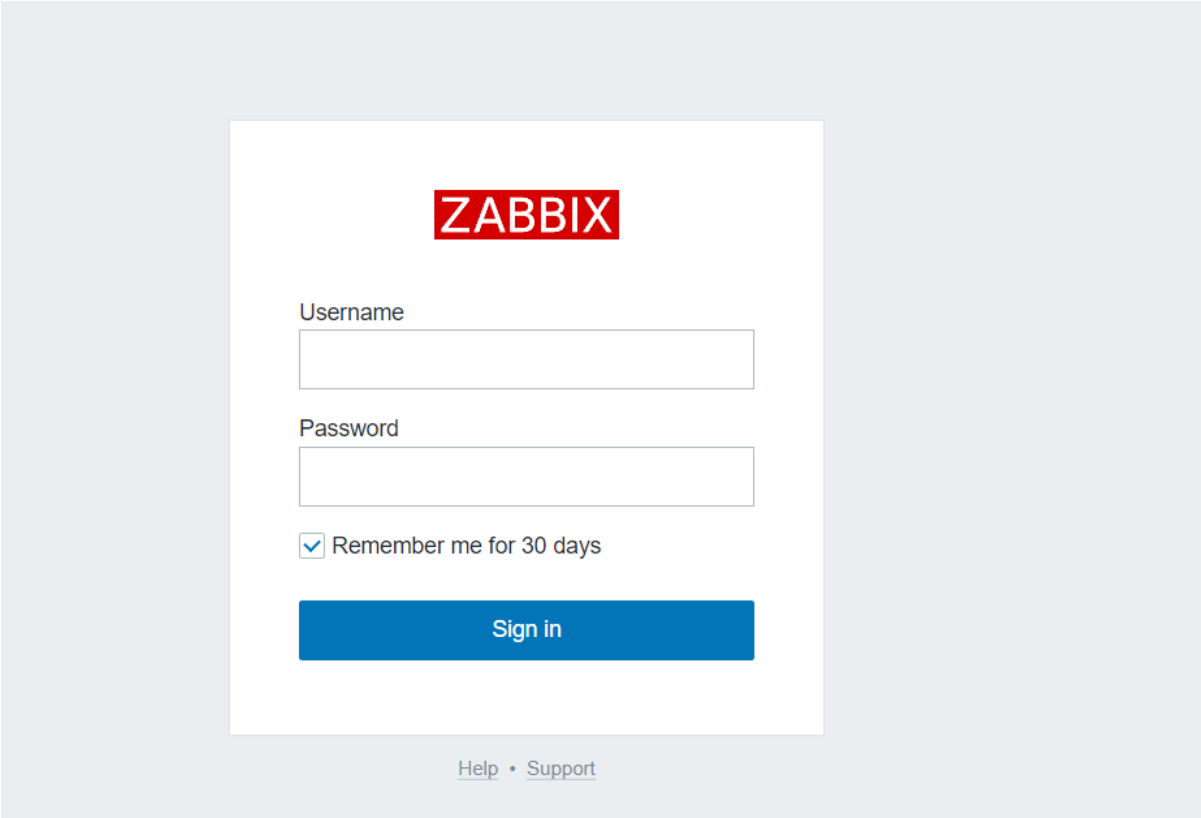
Open the *Instances* section and select the newly deployed Zabbix appliance instance.



To access Zabbix frontend using the instance's Public IPv4 address or Public IPv4 DNS, allow inbound HTTPS traffic on port 443 in the Security Group.



Now we can access Zabbix frontend by opening the Public IPv4 address or Public IPv4 DNS of Zabbix appliance.



To obtain the password, we need to connect to the machine either using the serial console or SSH. The default SSH username is rocky. Don't forget to use the key pair you selected or generated previously. Please note that generating credentials may take time. Credential generation may take a few minutes. Wait a few minutes before attempting to connect.

```
Last login: Tue Nov 16 15:16:07 2021 from ec2-13-48-4-200.eu-north-1.compute.amazonaws.com
********************************************************************************

Zabbix frontend credentials:

Username: Admin

Password: gJ8f7XpgJbsnc2BU48g

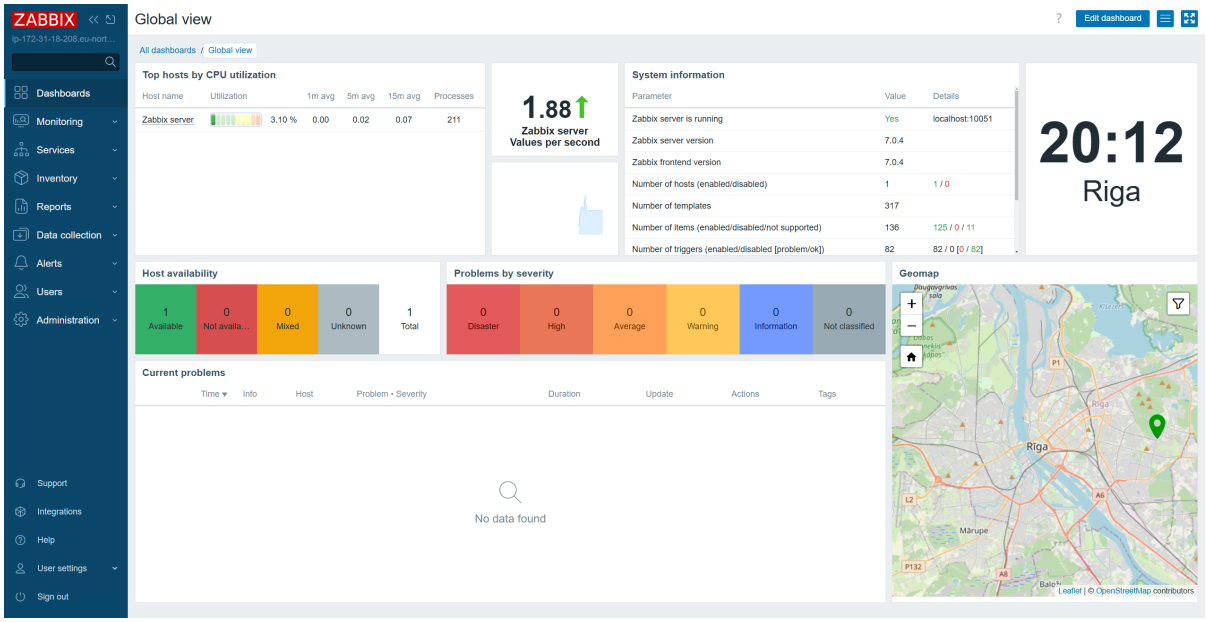
To learn about available professional services, including technical suppport and training, please visit https://www.zabbix.com/services

Official Zabbix documentation available at https://www.zabbix.com/documentation/current/

Note! Do not forget to change timezone PHP variable in /etc/php.d/99-zabbix.ini file.

********************************************************************************
```

Now you can use this information to log in to your Zabbix frontend.



Accessing the Database

You can access Zabbix database using the `root` user without entering a password.

```
[root@ip-172-31-29-138 ~]# mysql -uroot
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 124
Server version: 8.0.26 MySQL Community Server - GPL

Copyright (c) 2000, 2021, Oracle and/or its affiliates.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

mysql>
```

The credentials are stored in the `/root/.my.cnf` configuration file.

The default Zabbix frontend timezone can be changed in *Administration > General > GUI*

AWS deployment guide for Zabbix proxy

Overview

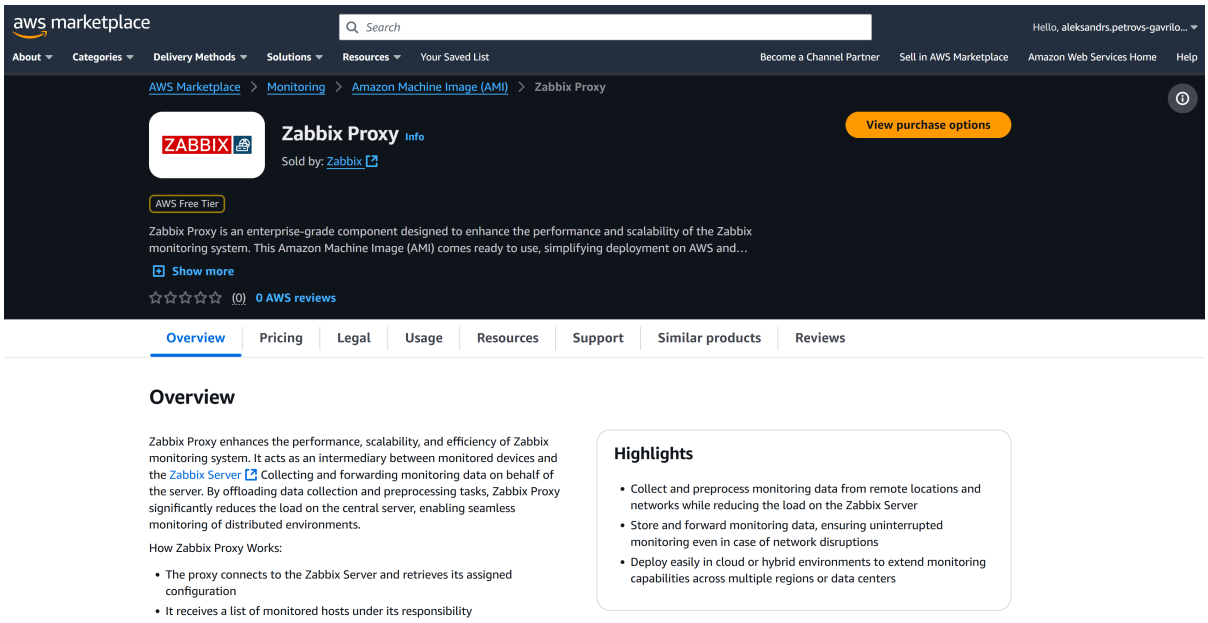
This page describes how to deploy Zabbix proxy using AWS.

Attention:
Several Zabbix versions are available for deployment on the platform. Make sure to read the documentation for the same Zabbix version that you have deployed. Documentation for different Zabbix versions can be selected from the version dropdown menu.

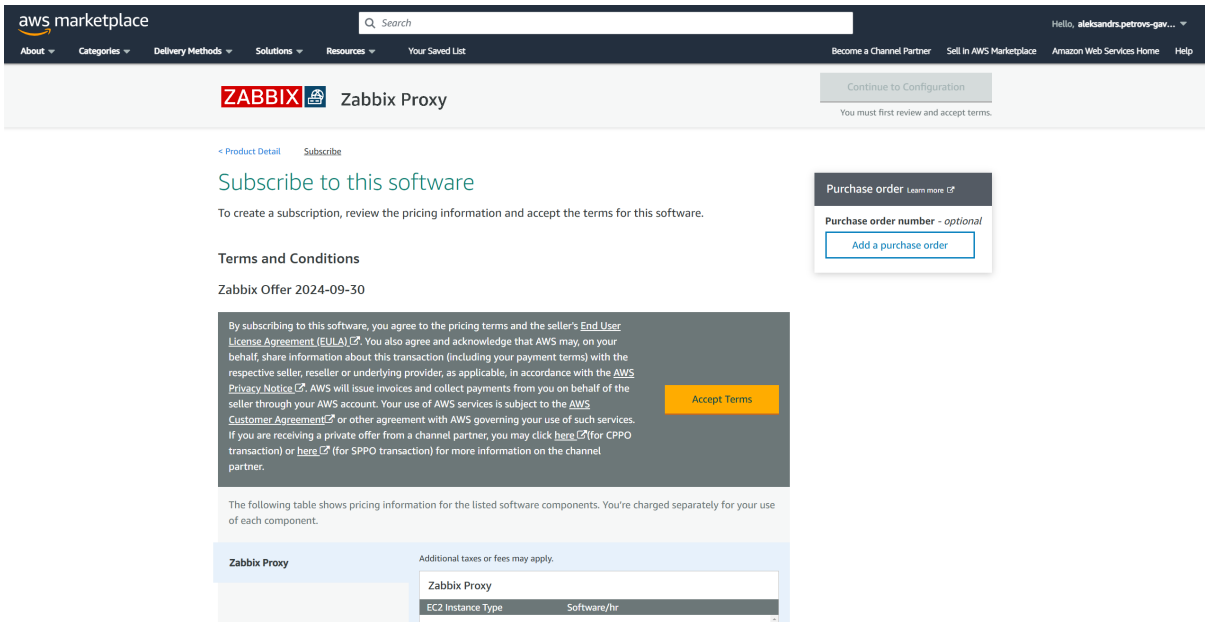
Deploying Zabbix proxy

Attention:
A working instance of Zabbix server is required to use Zabbix proxy component.

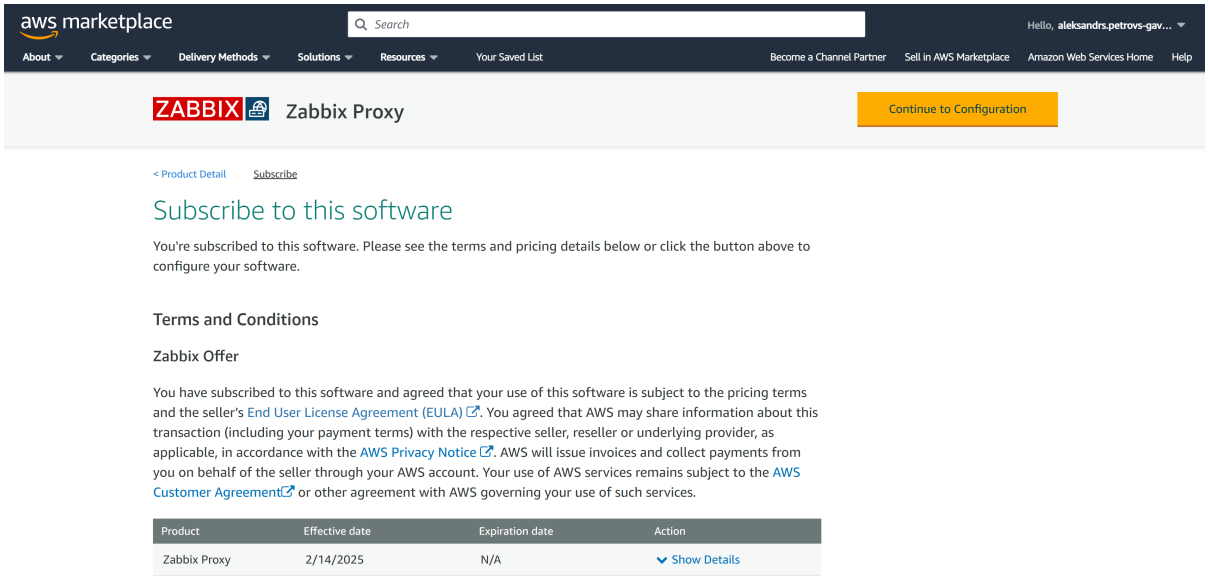
To deploy Zabbix proxy, first, navigate to the Zabbix AWS marketplace page. Next, click on the *View purchase options* to deploy the image and proceed to the subscription page.



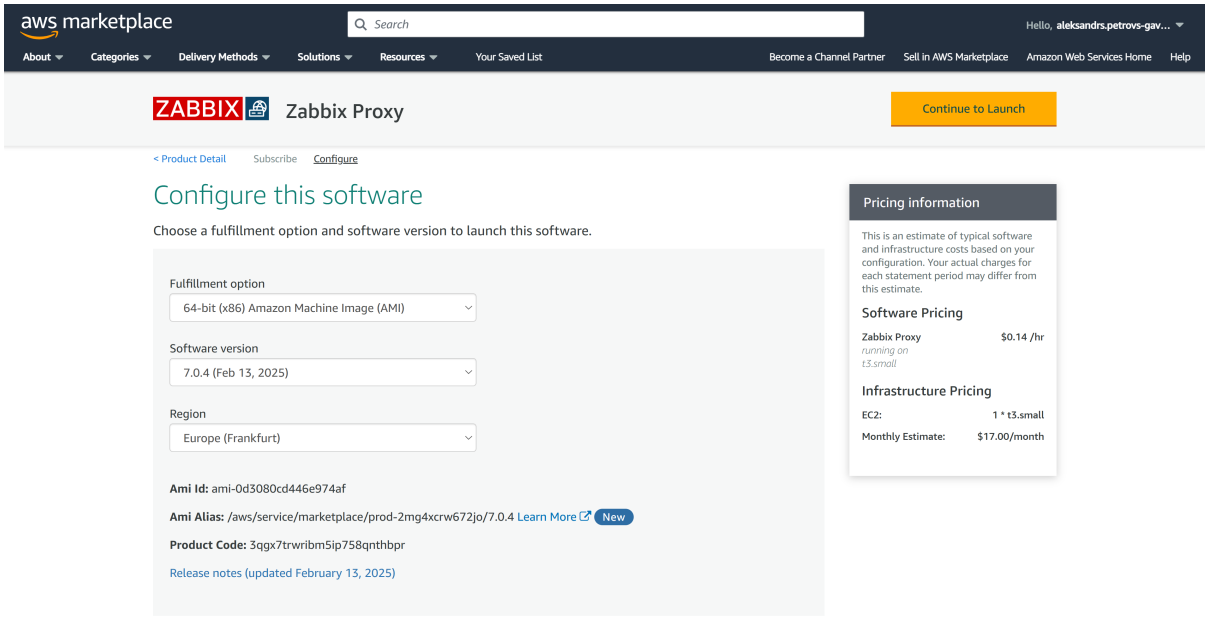
Click the *Accept terms* button to accept *Terms and Conditions*.



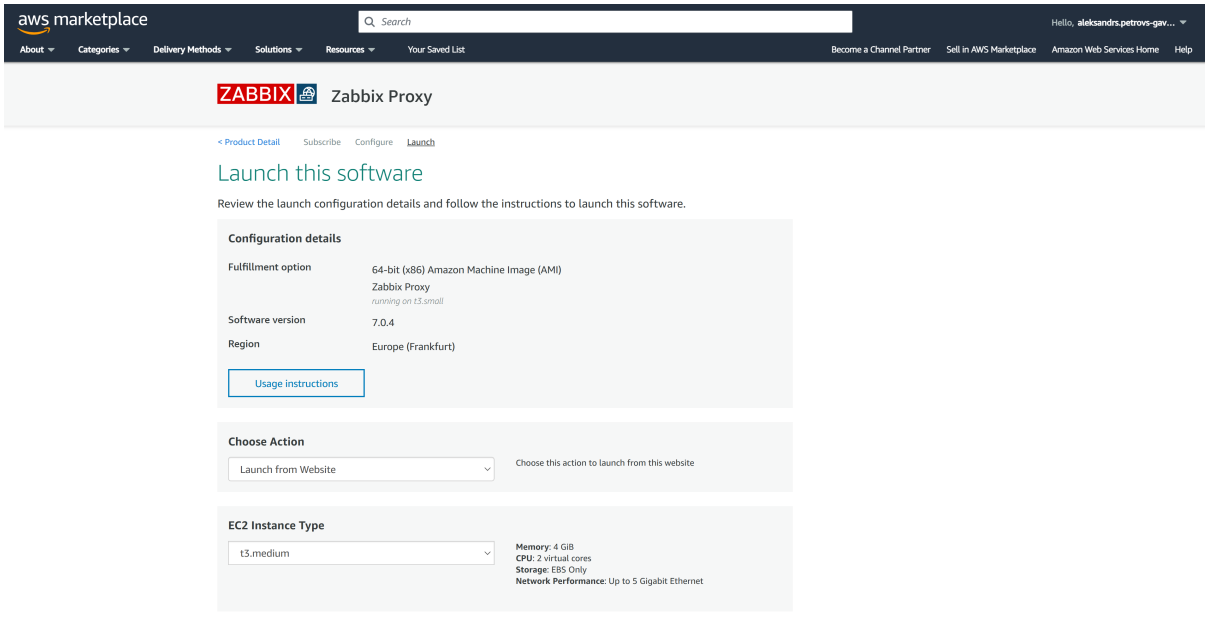
Click the *Continue to Configuration* button to move to the Configuration section for Zabbix proxy.



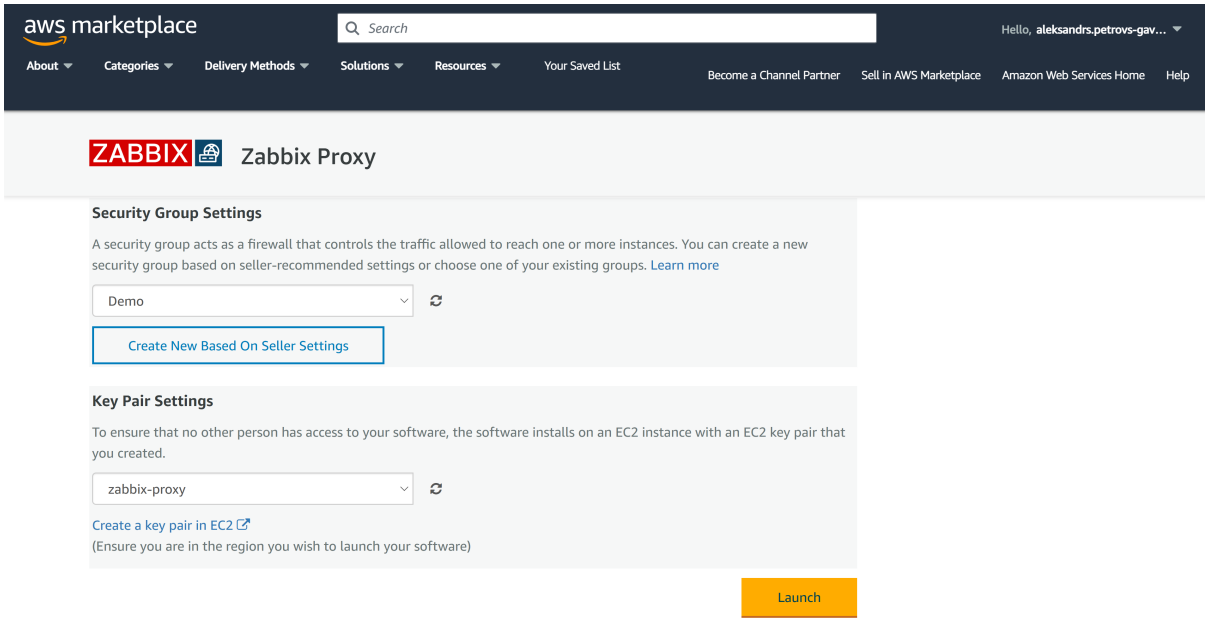
Select the Fulfilment option, Zabbix version, and your preferred Region.



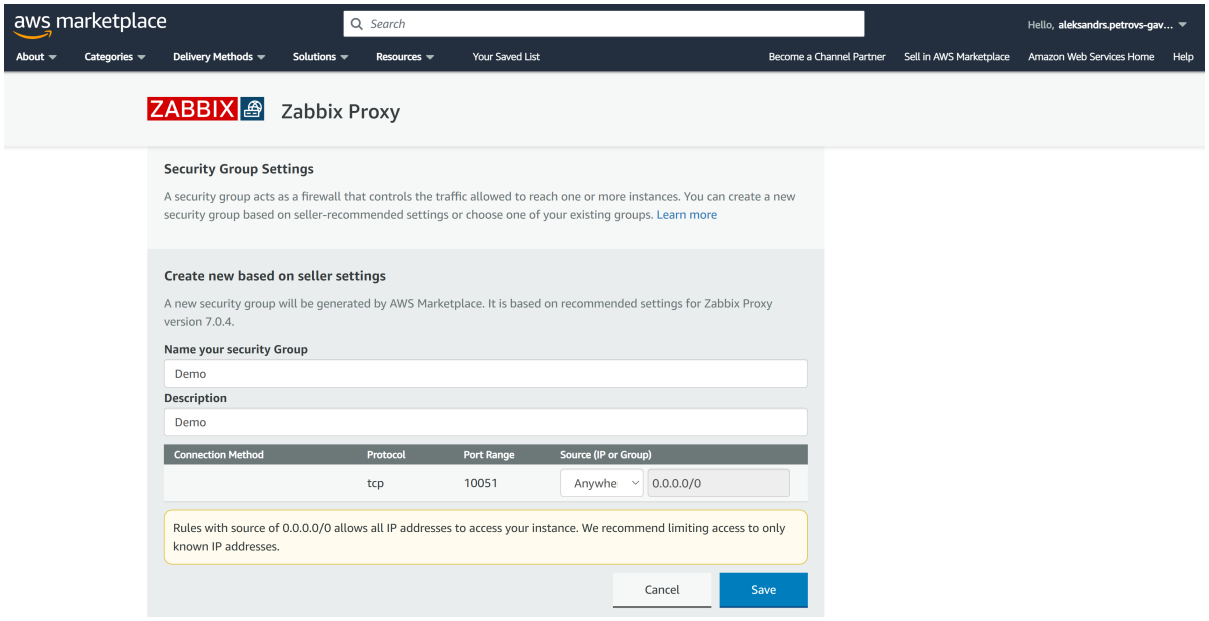
Click the *Continue to Launch* button. On the Launch page, select the preferred Launch option, EC2 Instance Type, VPC, and Subnet settings.



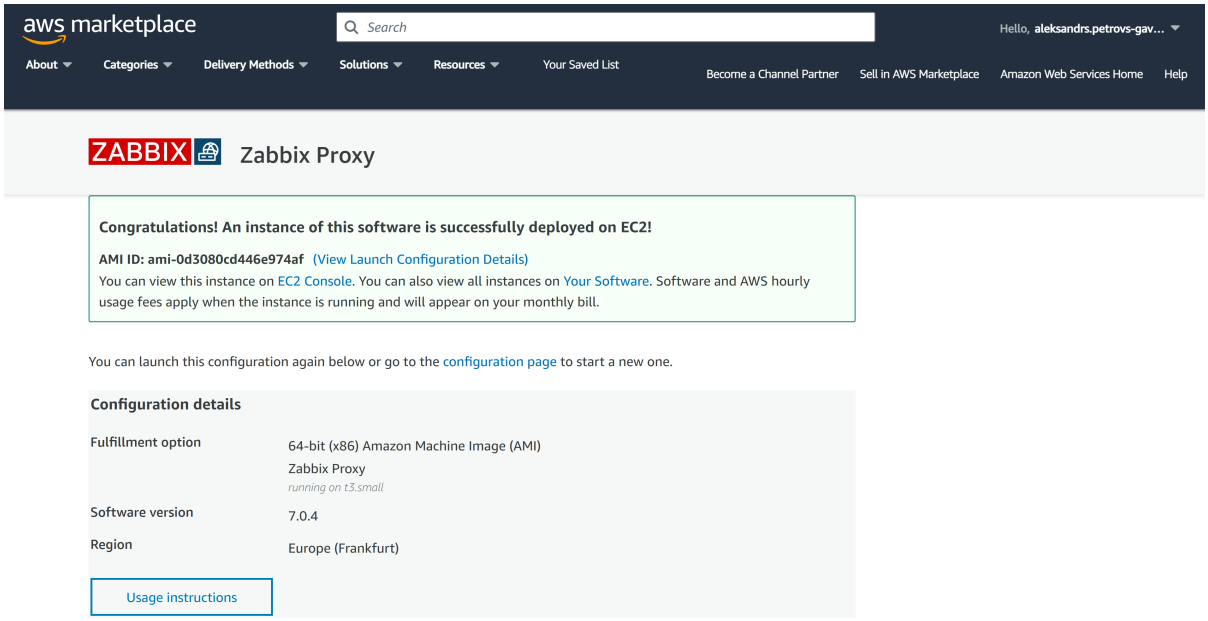
At the bottom of the Launch page, specify the Security Group settings and sselect or create an EC2 key pair. Make sure to save your private key in a safe location.



If you select *Create a new Security Group based on the Seller settings*, TCP port 10051 (Zabbix server) will be allowed. Note: If you wish to access the instance via SSH, you will also need to permit SSH connections in your Security Group settings.

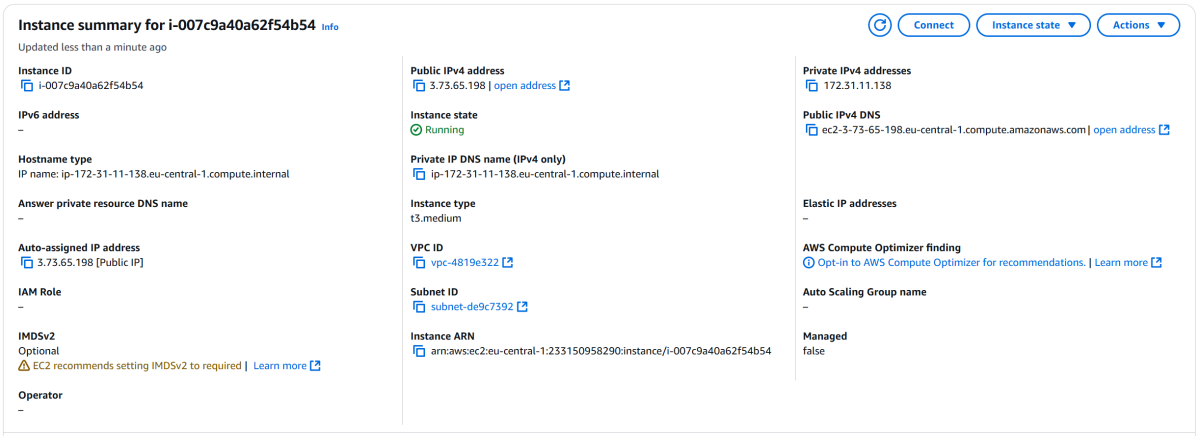


Finally, click the *Launch* button. Congratulations, your Zabbix proxy is now being deployed.



Configuring and accessing Zabbix proxy

Open the Instances section in the EC2 console and select the newly deployed Zabbix proxy instance.



Connect to the Public IPv4 address or Public IPv4 DNS of the Zabbix proxy using SSH. The default SSH username is `rocky`. Don't forget to use the key pair that you generated previously.

Credential generation may take a few minutes. Wait a few minutes before attempting to connect.



Zabbix proxy configuration file can be found in `/etc/zabbix/zabbix_proxy.conf`.

The configuration file can be opened with a text editor of your choice, for example vi:

```
sudo vi /etc/zabbix/zabbix_proxy.conf
```

The initial proxy configuration consists of modifying the `ProxyMode`, `Server`, and `Hostname` parameters in Zabbix proxy configuration file.

The proxy mode can be changed by editing the `ProxyMode` parameter. Set `ProxyMode=0` for active mode or `ProxyMode=1` for passive mode.

The `Server` parameter should contain Zabbix server IP address. For example:

```
Server=192.168.50.141
```

The `Hostname` parameter should contain Zabbix proxy hostname. The hostname defined here should match the proxy name provided during proxy configuration in Zabbix frontend. For example:

```
Hostname=EC2-Frankfurt
```

Once the configuration is complete, restart the Zabbix proxy with the following command:

```
sudo systemctl restart zabbix-proxy
```

Note: use sudo to edit configuration files and execute commands from the command line.

```
### Option: ProxyMode
#       Proxy operating mode.
#       0 - proxy in the active mode
#       1 - proxy in the passive mode
#
# Mandatory: no
# Default:
# ProxyMode=0

### Option: Server
#       If ProxyMode is set to active mode:
#               IP address or DNS name of Zabbix server to get configuration data from and send data to.
#       If ProxyMode is set to passive mode:
#               List of comma delimited IP addresses, optionally in CIDR notation, or DNS names of Zabbix
server.
#               Incoming connections will be accepted only from the addresses listed here.
#               If IPv6 support is enabled then '127.0.0.1', '::127.0.0.1', '::ffff:127.0.0.1' are treated
 equally
#               and '::/0' will allow any IPv4 or IPv6 address.
#               '0.0.0.0/0' can be used to allow any IPv4 address.
#               Example: Server=127.0.0.1,192.168.1.0/24,::1,2001:db8::/32,zabbix.example.com
#
# Mandatory: yes
# Default:
# Server=

Server=127.0.0.1
```

```
### Option: Hostname
#       Unique, case sensitive Proxy name. Make sure the Proxy name is known to the server!
#       Value is acquired from HostnameItem if undefined.
#
# Mandatory: no
# Default:
# Hostname=

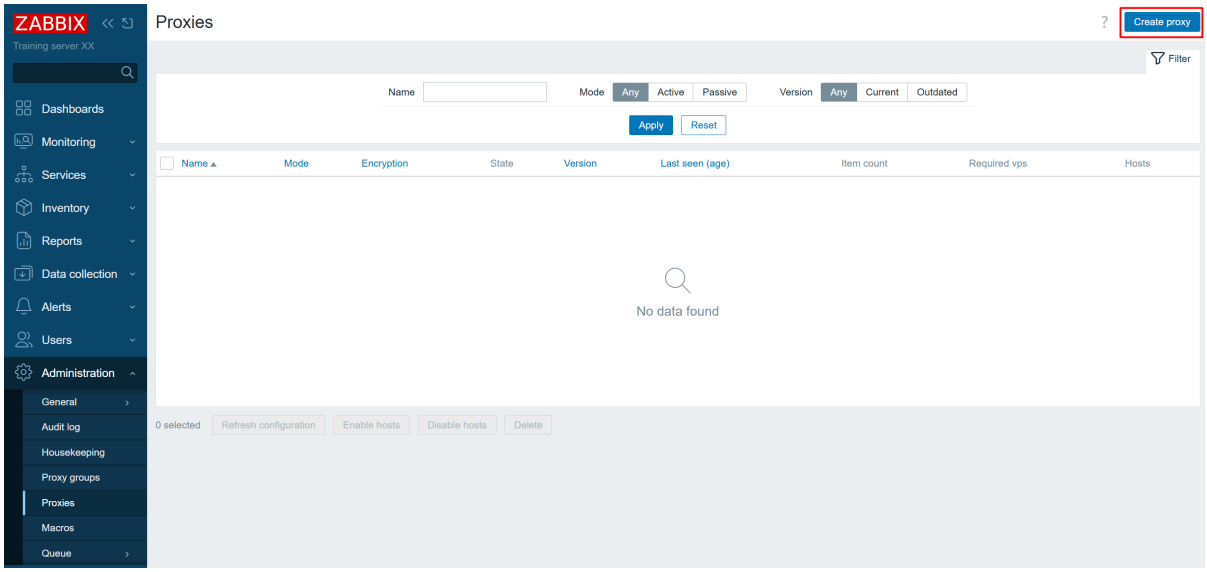
Hostname=EC2-Frankfurt
```

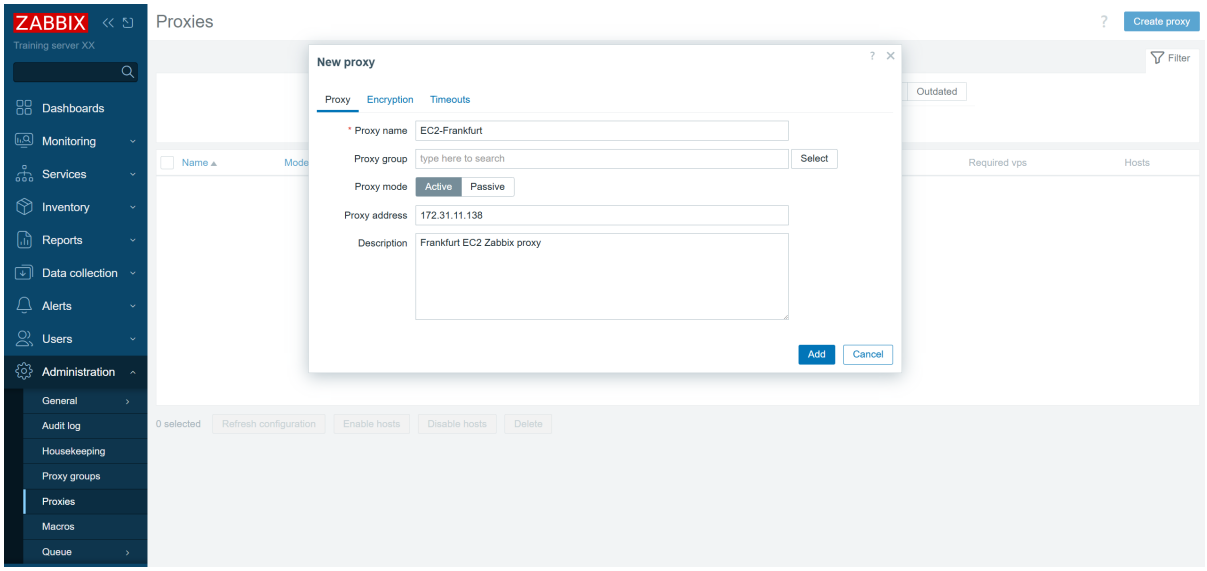
Open Zabbix frontend and navigate to *Administration* > *Proxies*. Click the *Create proxy* button.

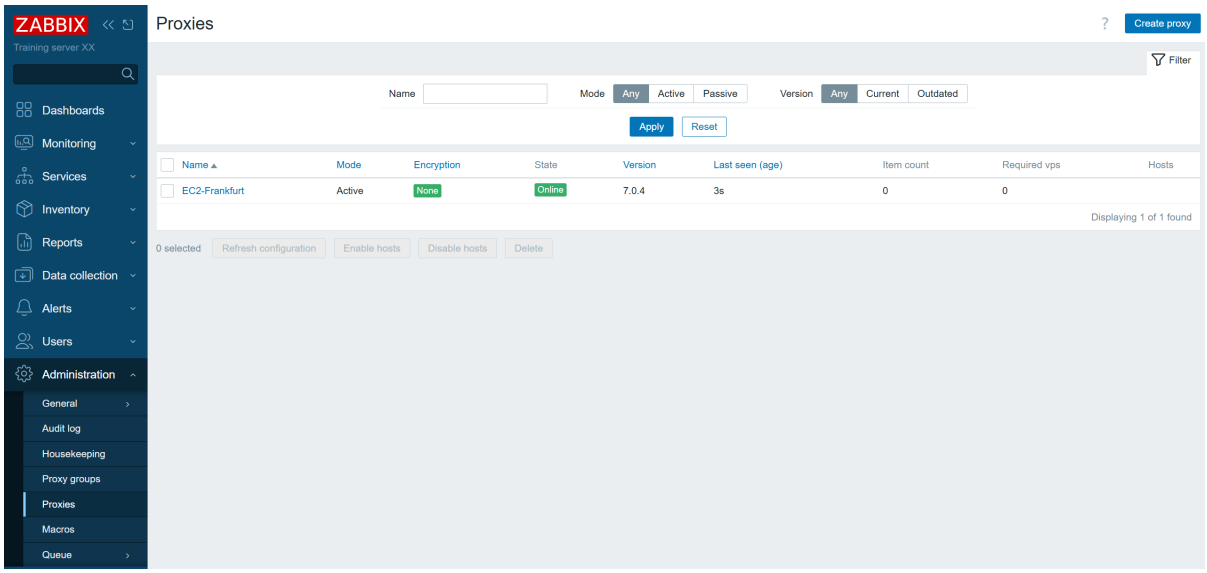
Enter the same proxy name as the value of the `Hostname` parameter in the proxy configuration file.

Select the proxy mode and provide the proxy IP address (optional for active proxies).

Click the *Create proxy*.







Proxy database location

The SQLite proxy database is stored in `/tmp/zabbix_proxy.sqlite3` and is specified by the `DBName` parameter in the Zabbix proxy configuration file.

Azure deployment guide for Zabbix server

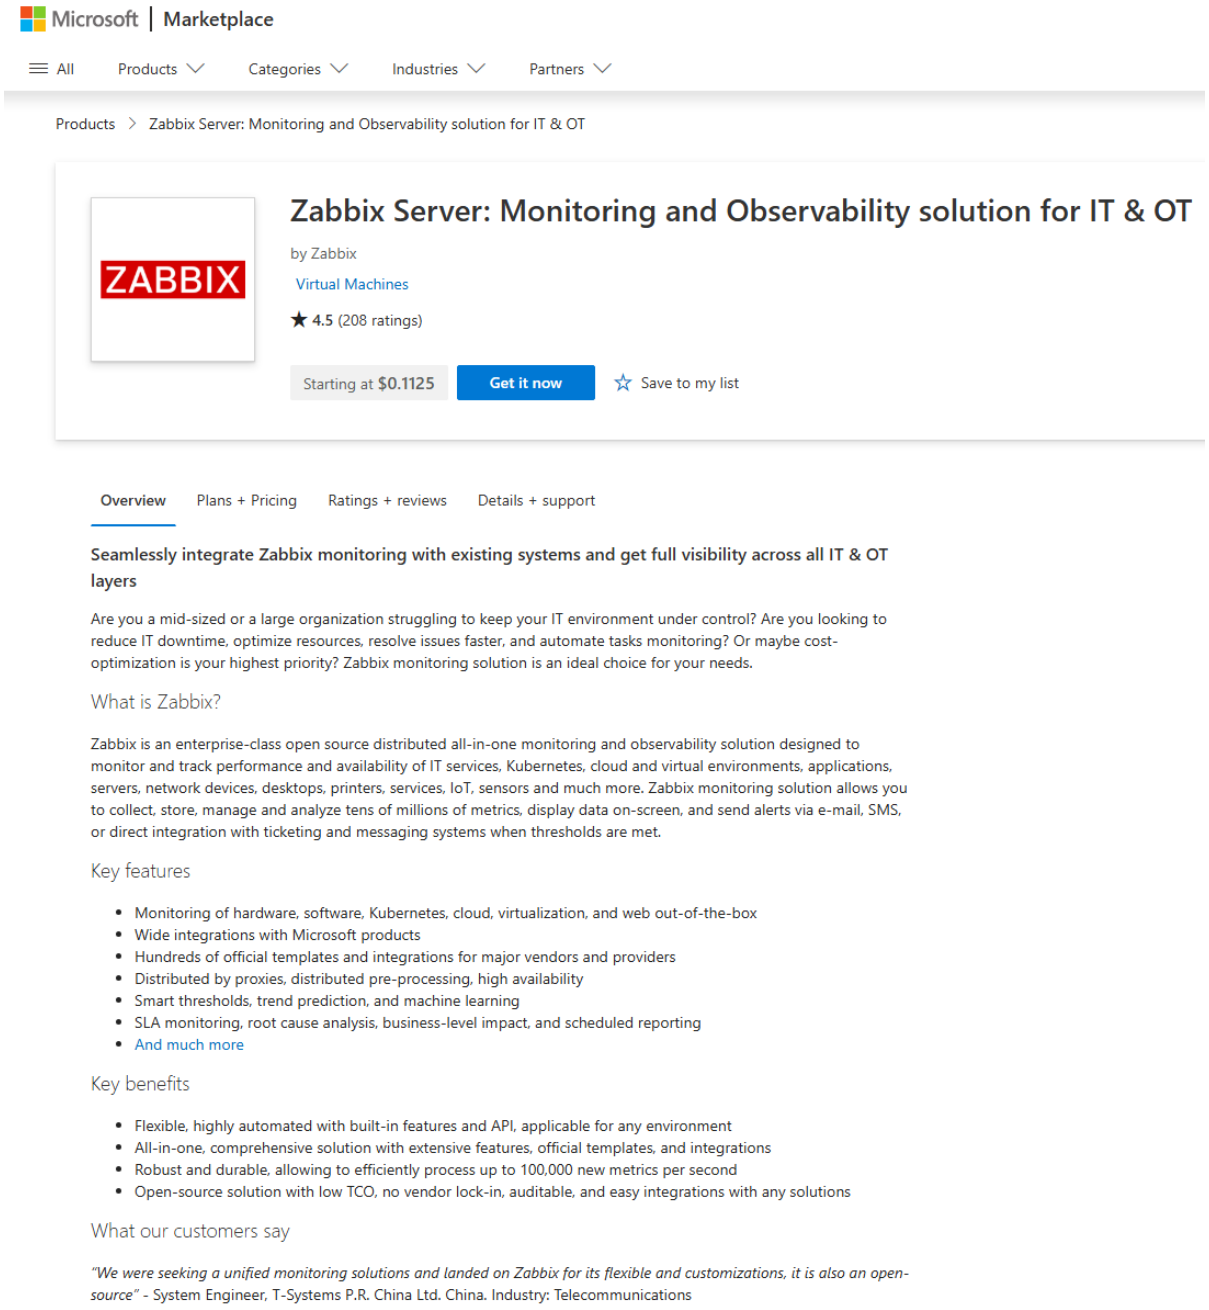
Overview

This page describes how to deploy Zabbix server using Azure.

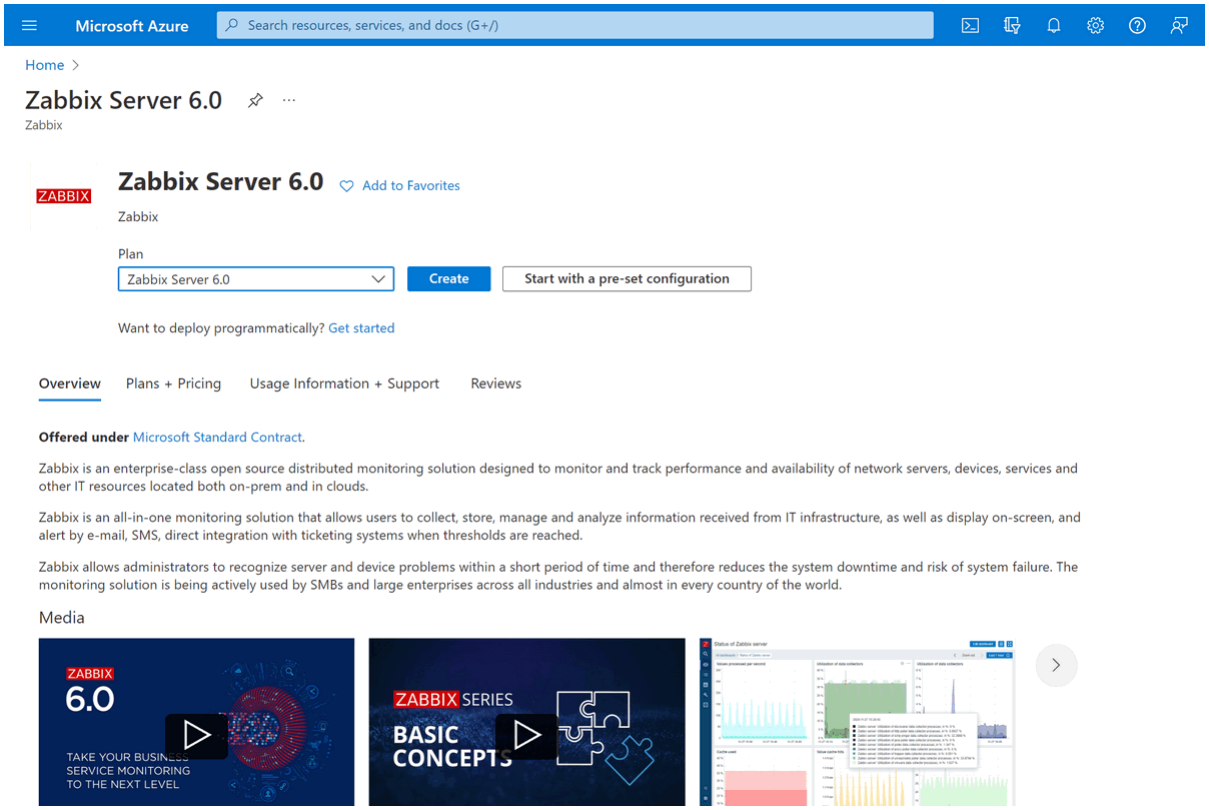
Attention:
Several Zabbix versions are available for deployment on the platform. Make sure to read the documentation for the same Zabbix version that you have deployed. Documentation for different Zabbix versions can be selected from the version dropdown menu.

Deploying the appliance

To deploy Zabbix appliance, first, navigate to Zabbix Azure Marketplace page. Review the overview, then click *Get it now*.



Review the deployment overview, pricing plans, and other details. Next, click either *Create* or *Start with a pre-set configuration* to deploy your Zabbix appliance.

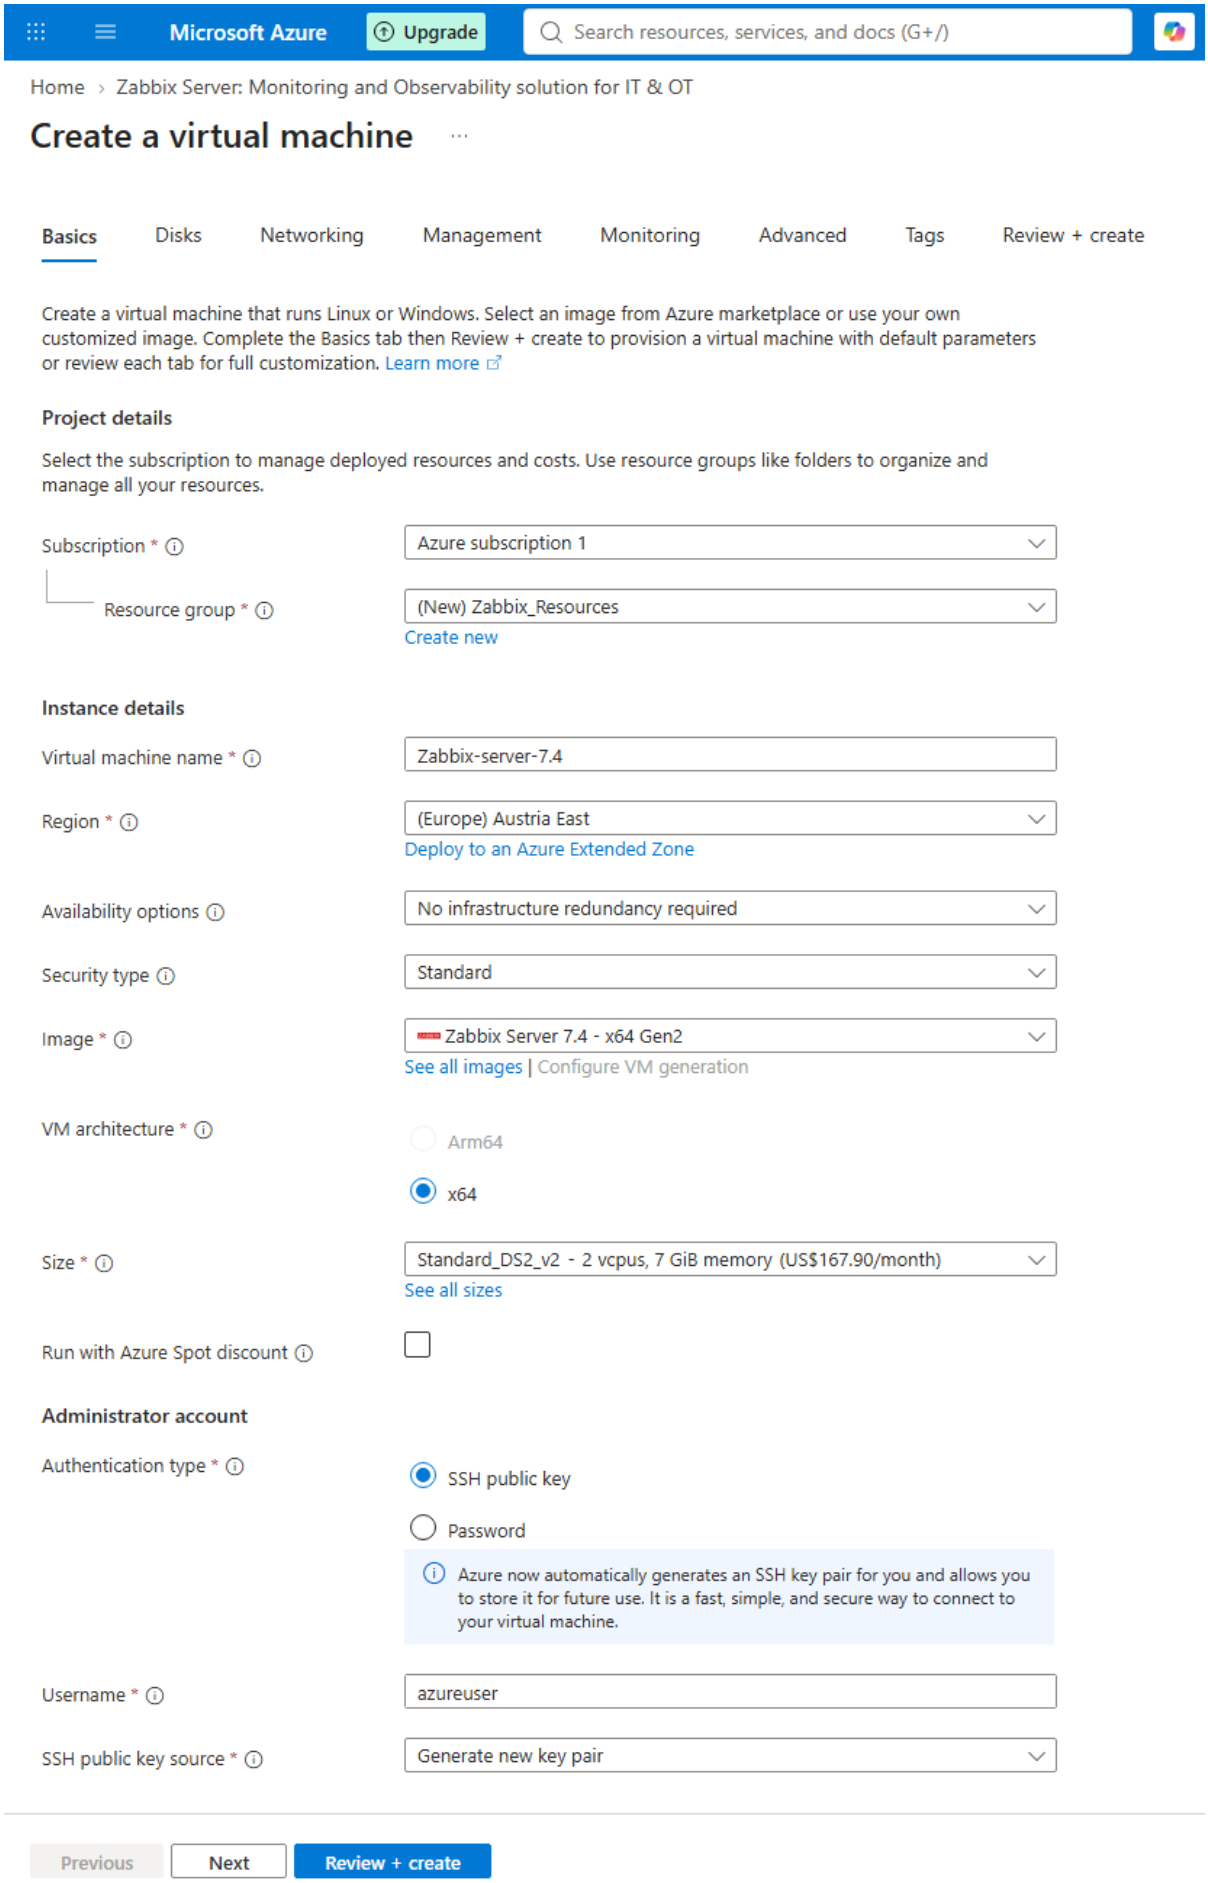


Configure the VM by selecting the VM size, creating an administrator account, and defining disk, network, management, and other settings. These settings may vary depending on your infrastructure requirements.

You can use the following configuration as an example:

Basics tab:

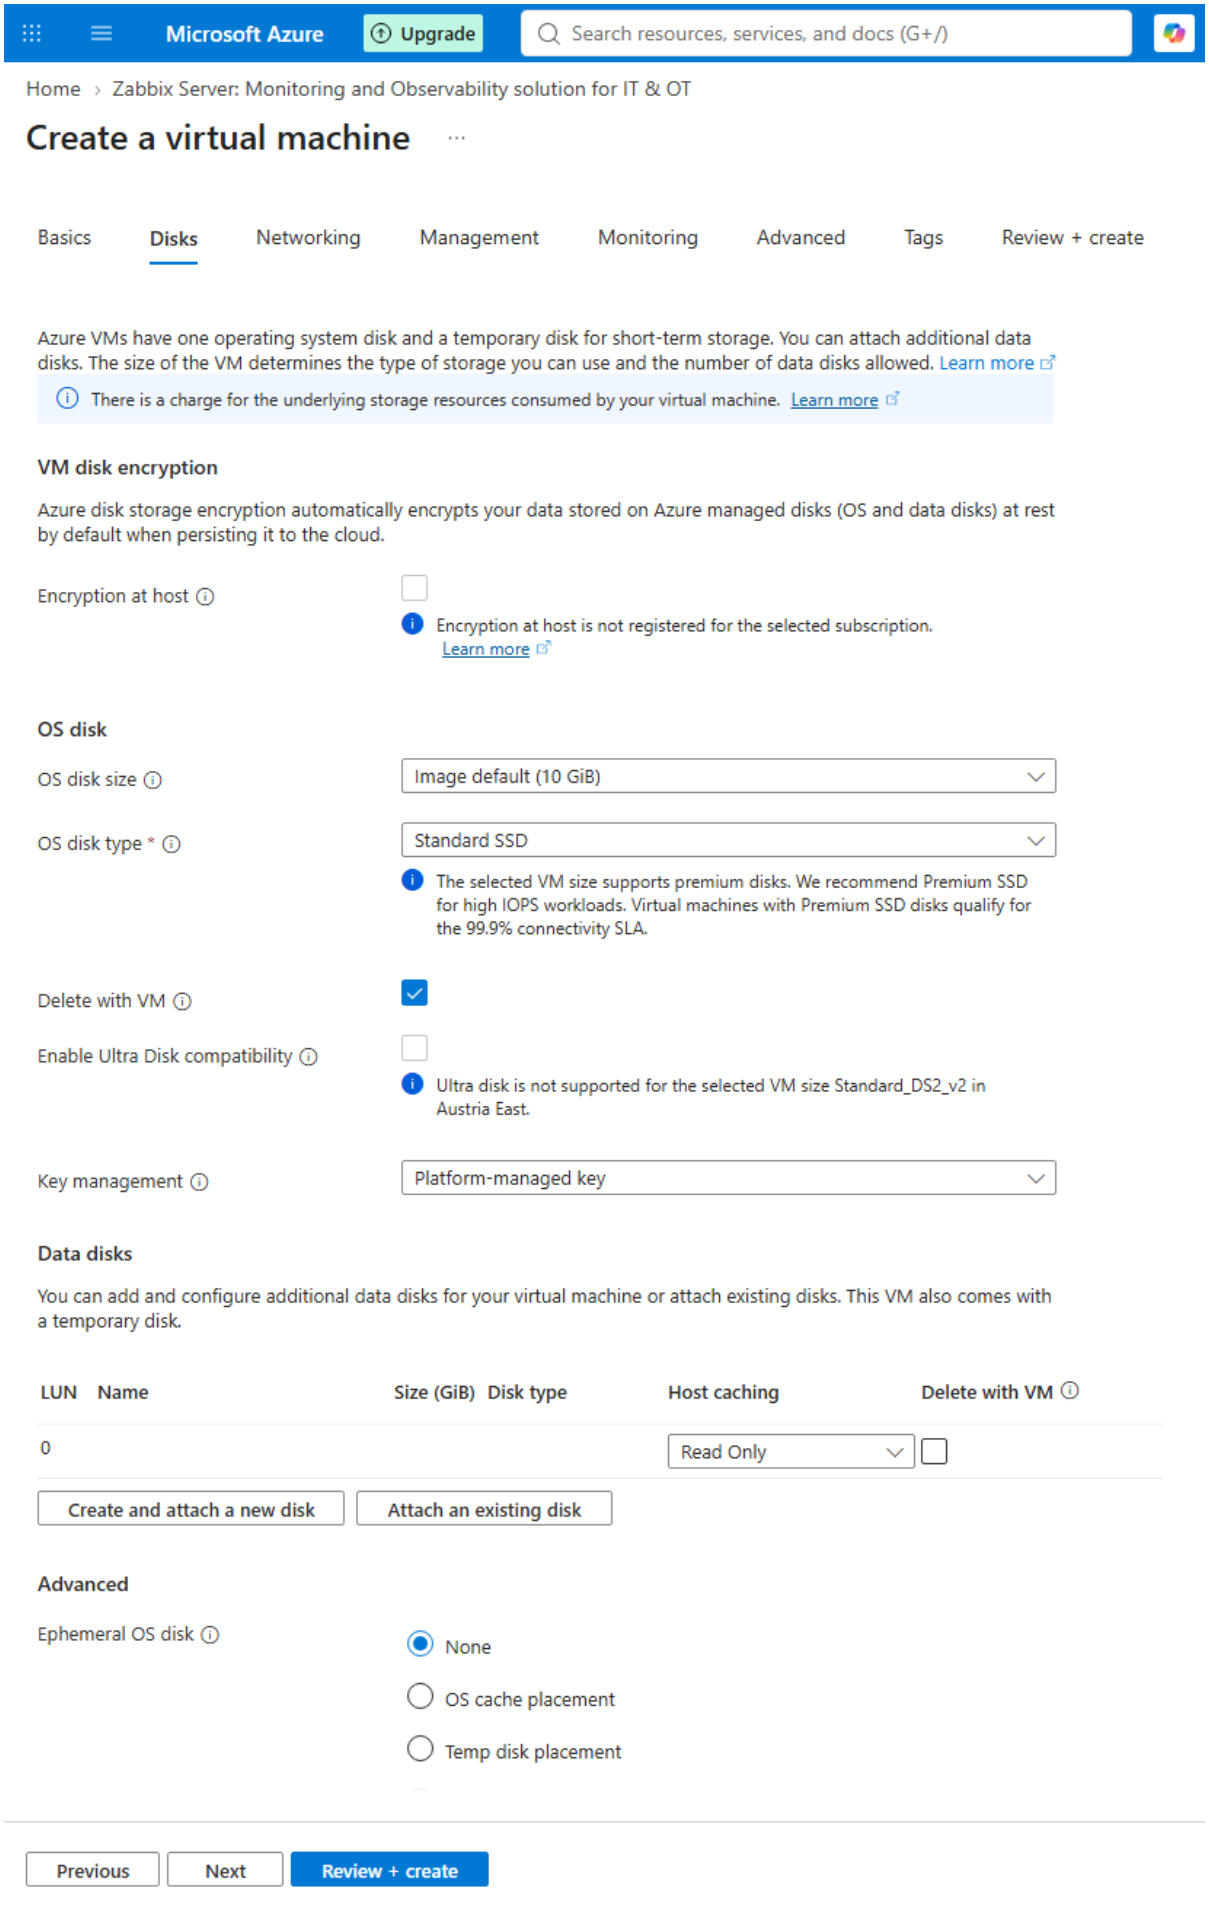
- *Subscription*: Azure subscription 1
- *Resource group*: Zabbix_Resources
- *Virtual machine name*: Zabbix-server--7.4
- *Availability options*: No infrastructure redundancy required
- *Security type*: Standard
- *Image*: Zabbix Server 7.4 - x64 Gen2
- *Size*: Standard_DS2_V2 - 2 vcpus, 7 GiB memory
- *Authentication type*: SSH public key
- *Username*: azureuser
- *SSH public key source*: Generate new key pair



Disks tab:

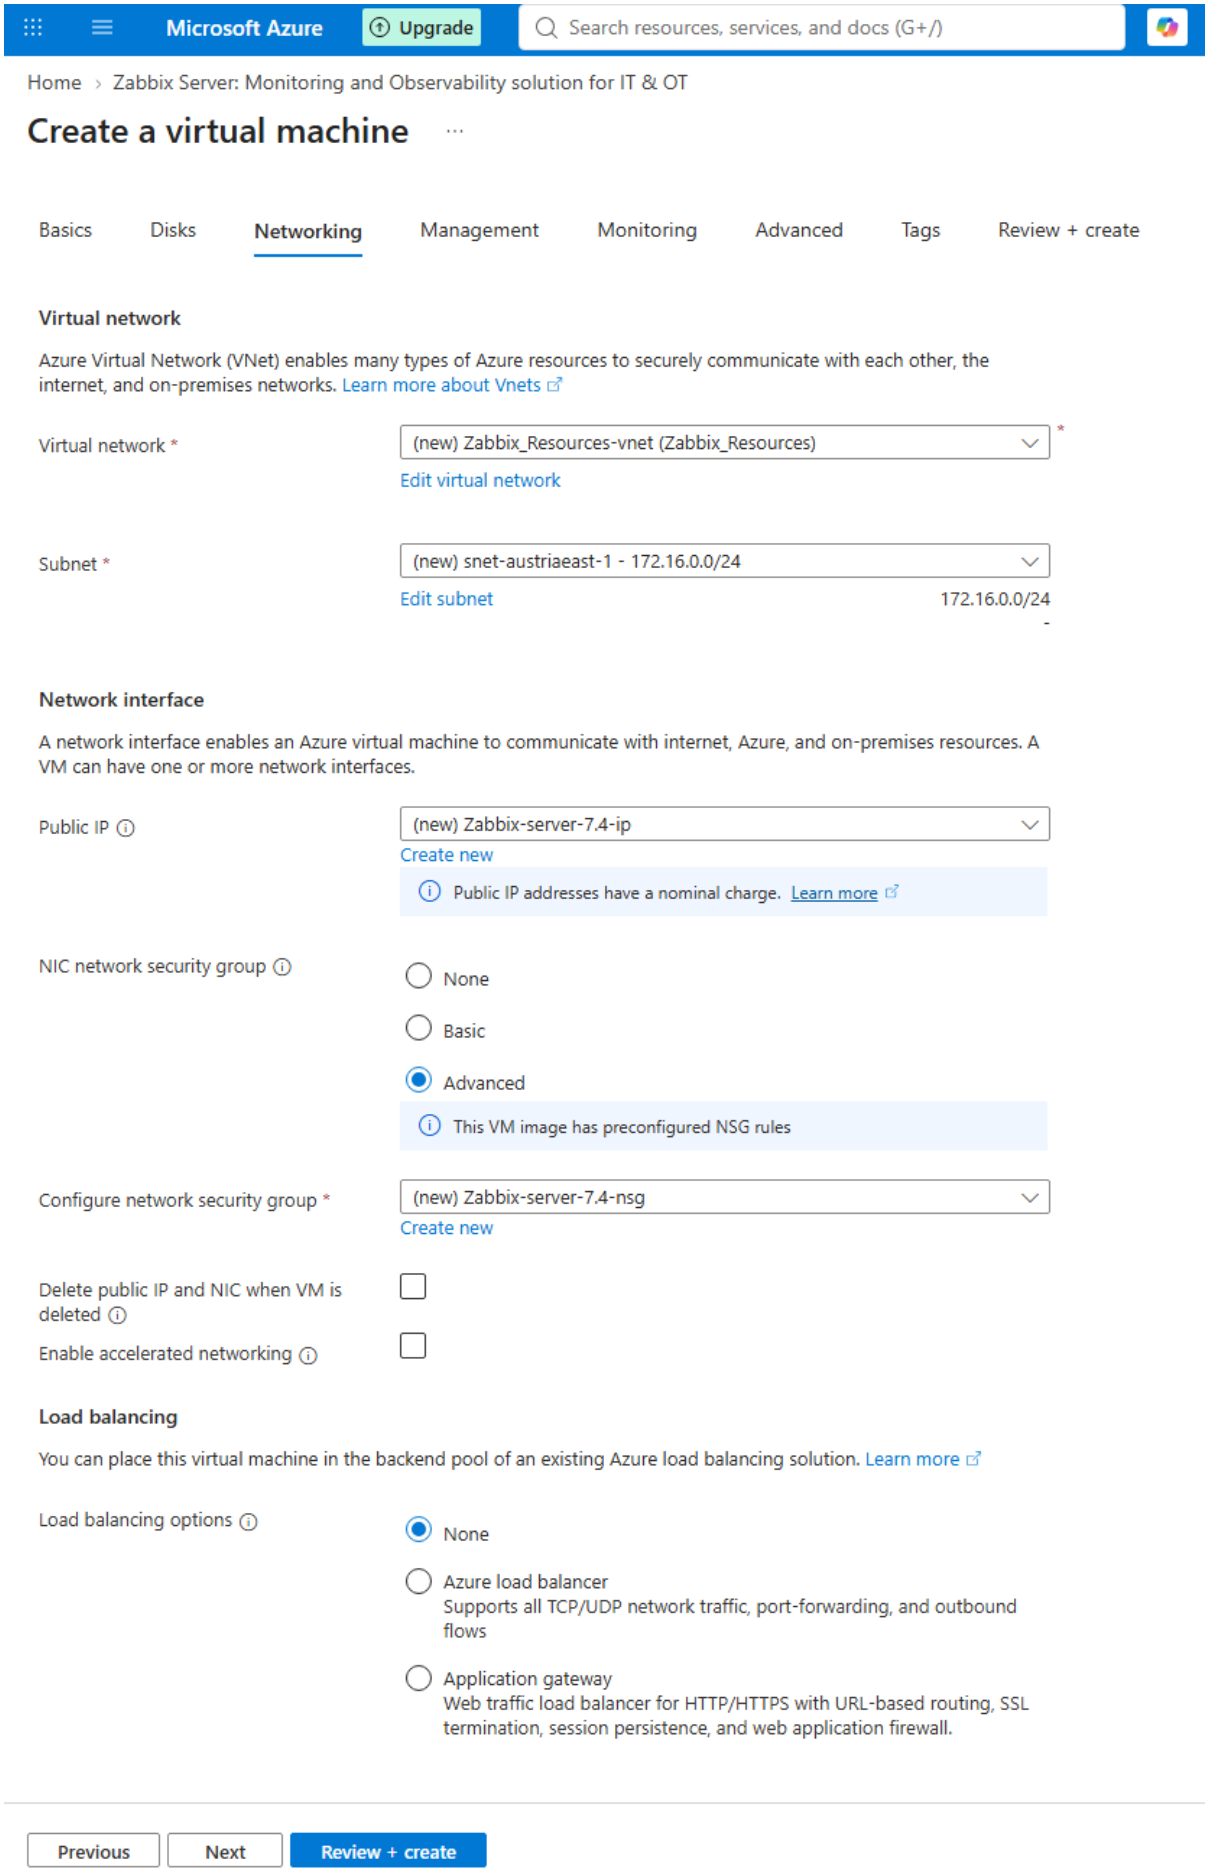
- *Encryption at host*: unchecked
- *OS disk type*: Standard SSD

- *Delete with VM*: checked
- *Enable Ultra Disk compatibility*: unchecked

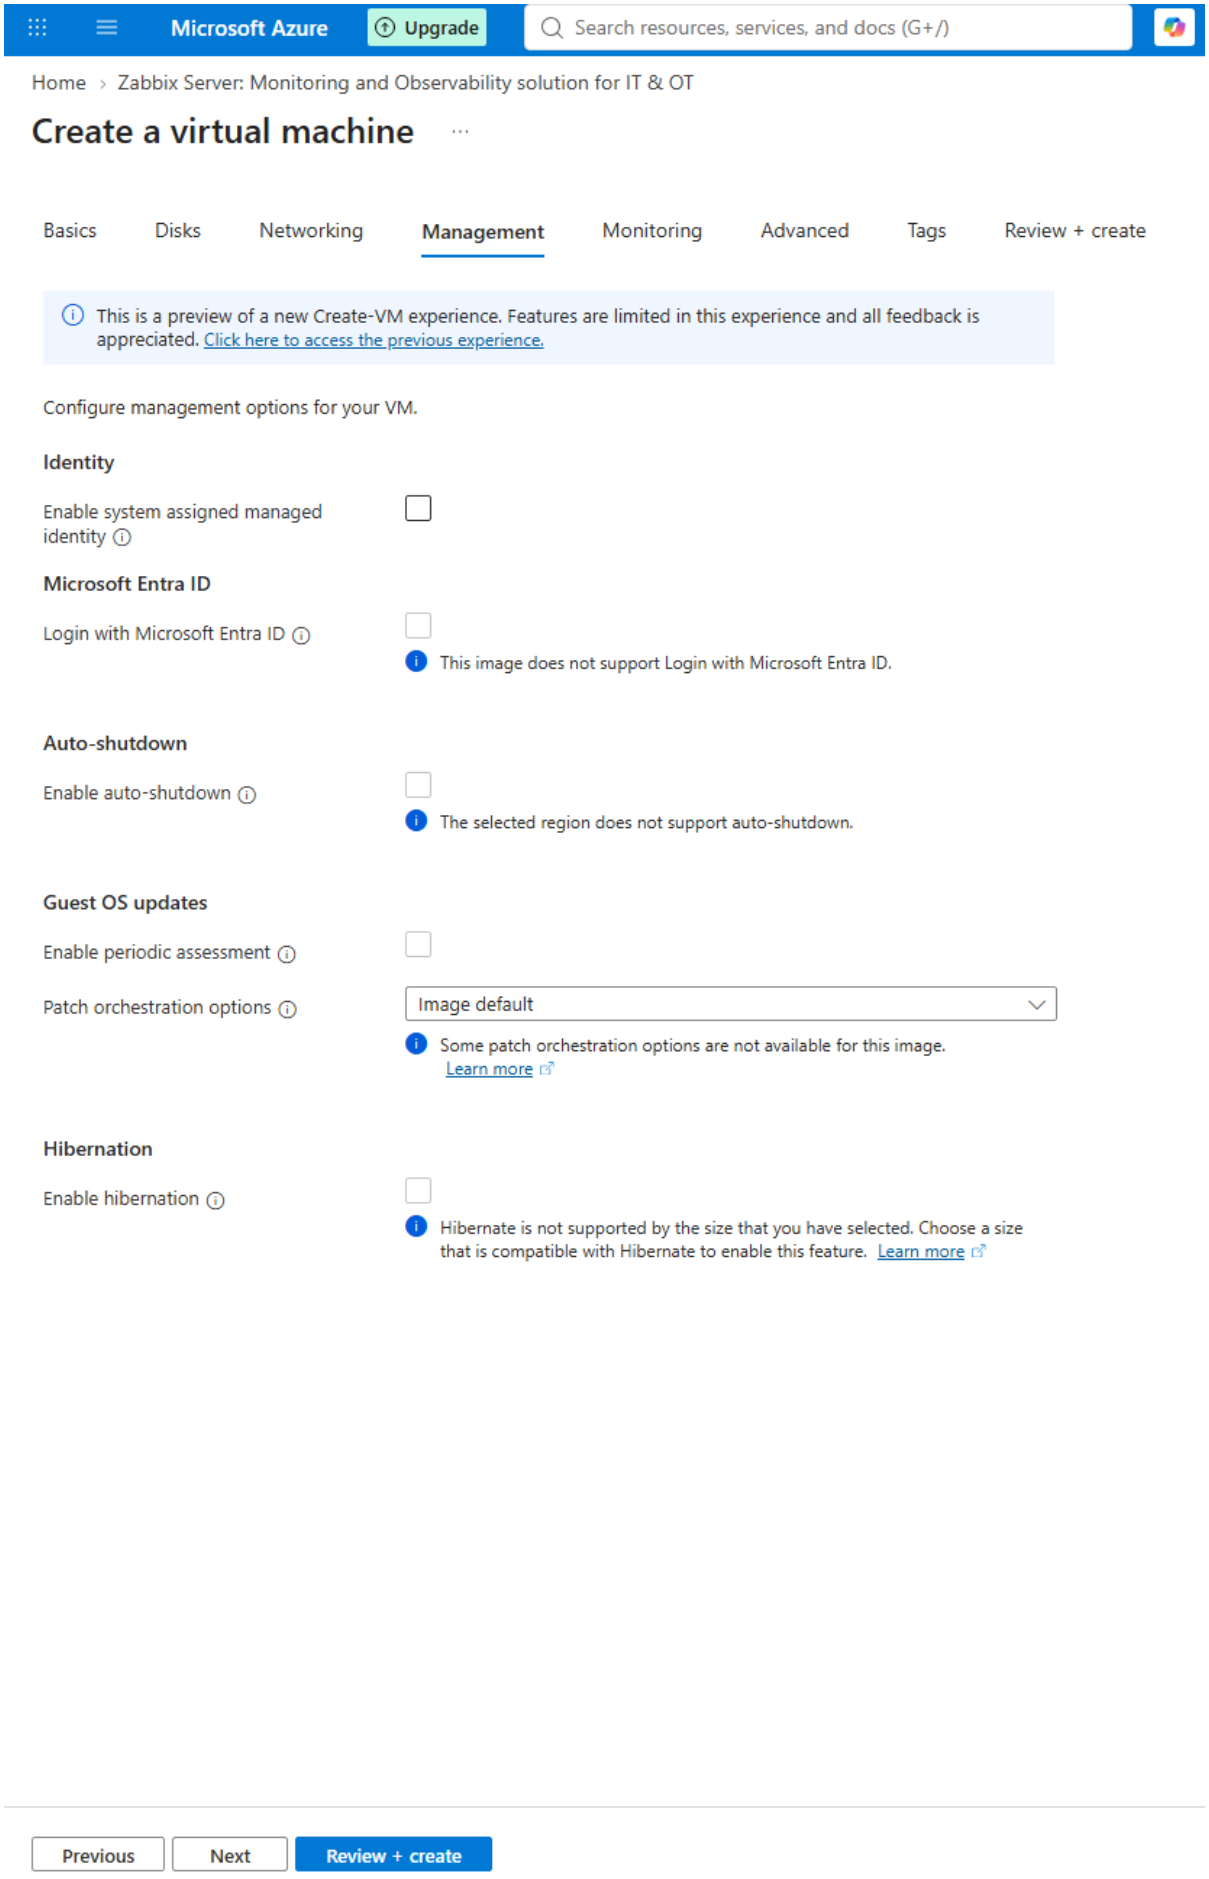


Networking tab:

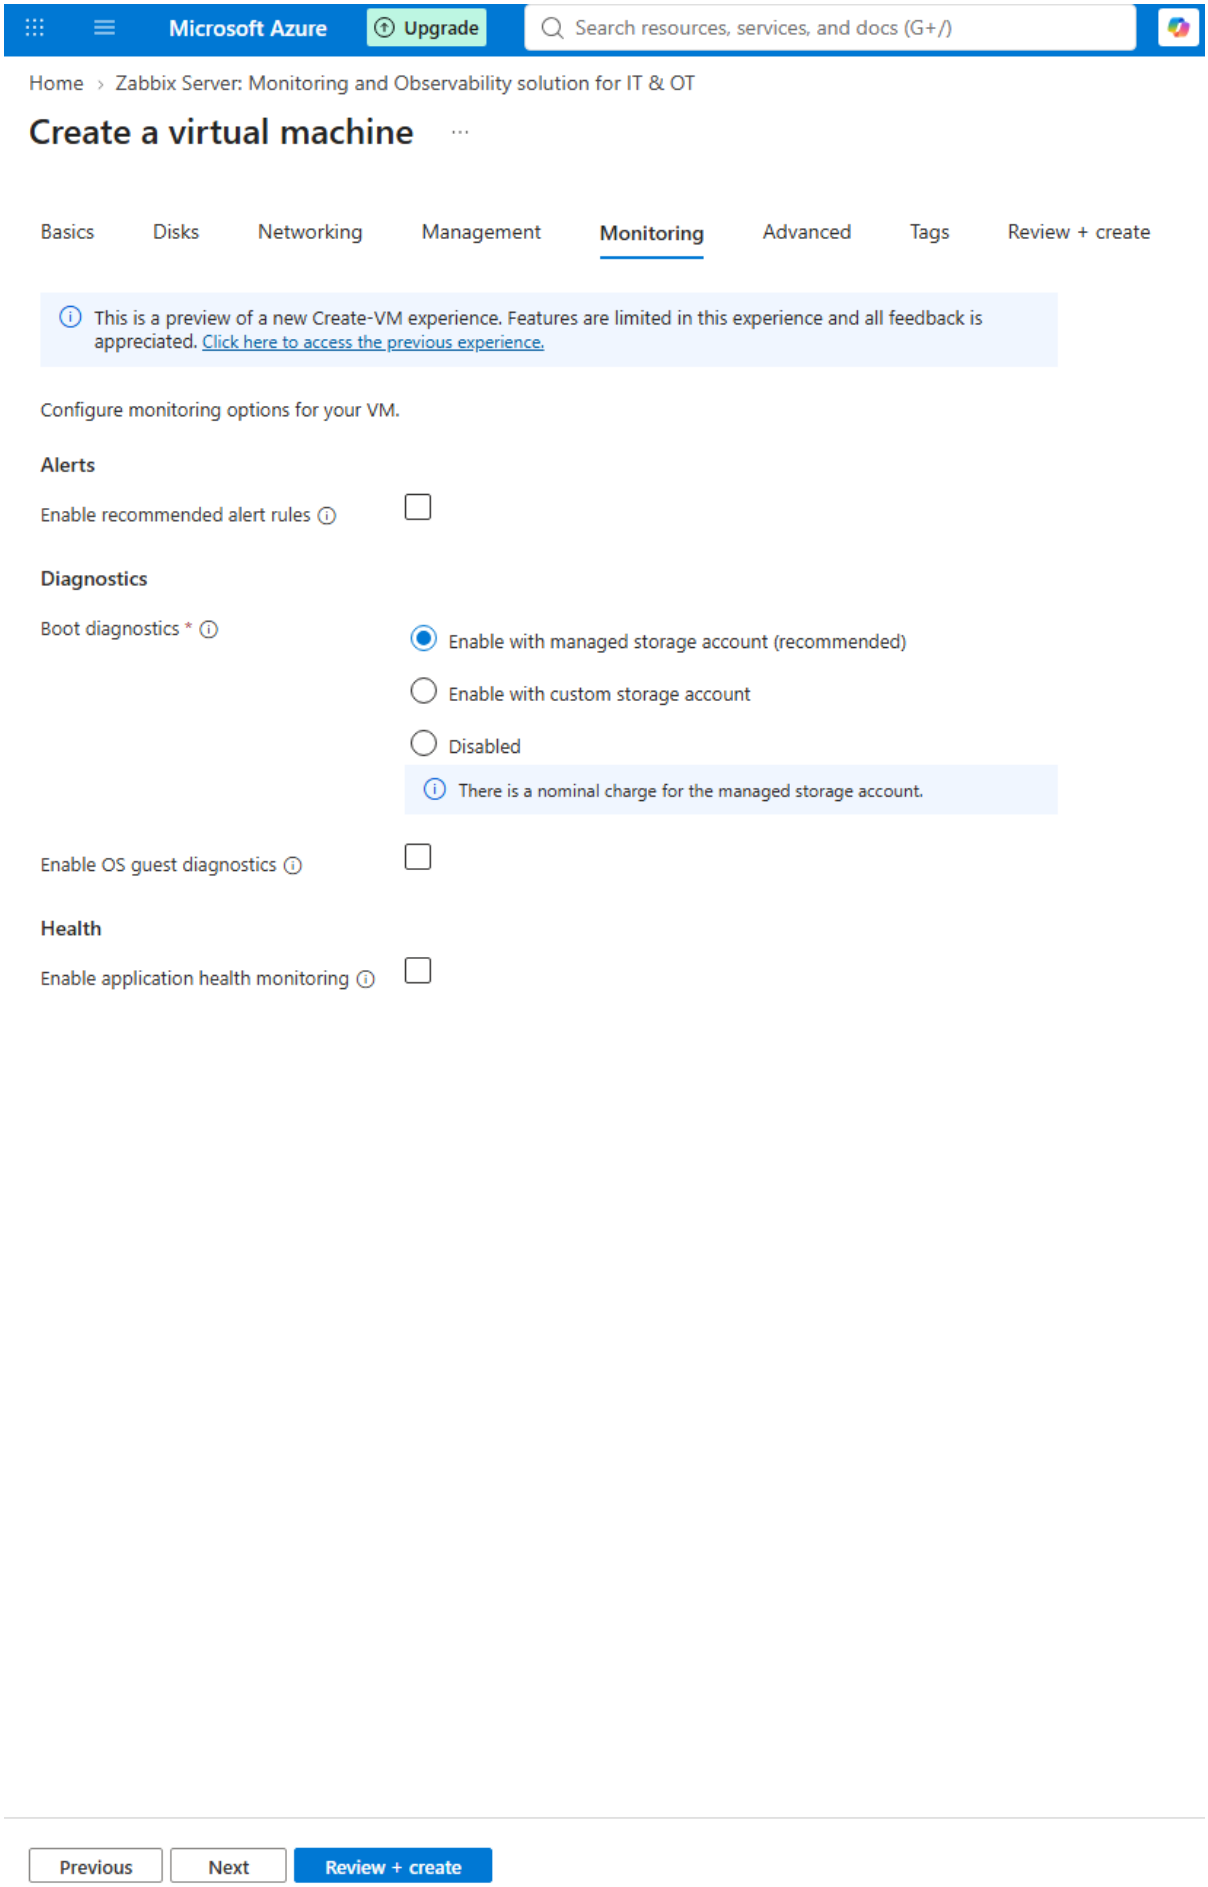
- *Virtual network*: Zabbix_Resources-vnet
- *Public IP*: Zabbix-server-7.4-ip
- *NIC network security group*: Advanced
- *Configure network security group*: Zabbix-server-7.4-nsg



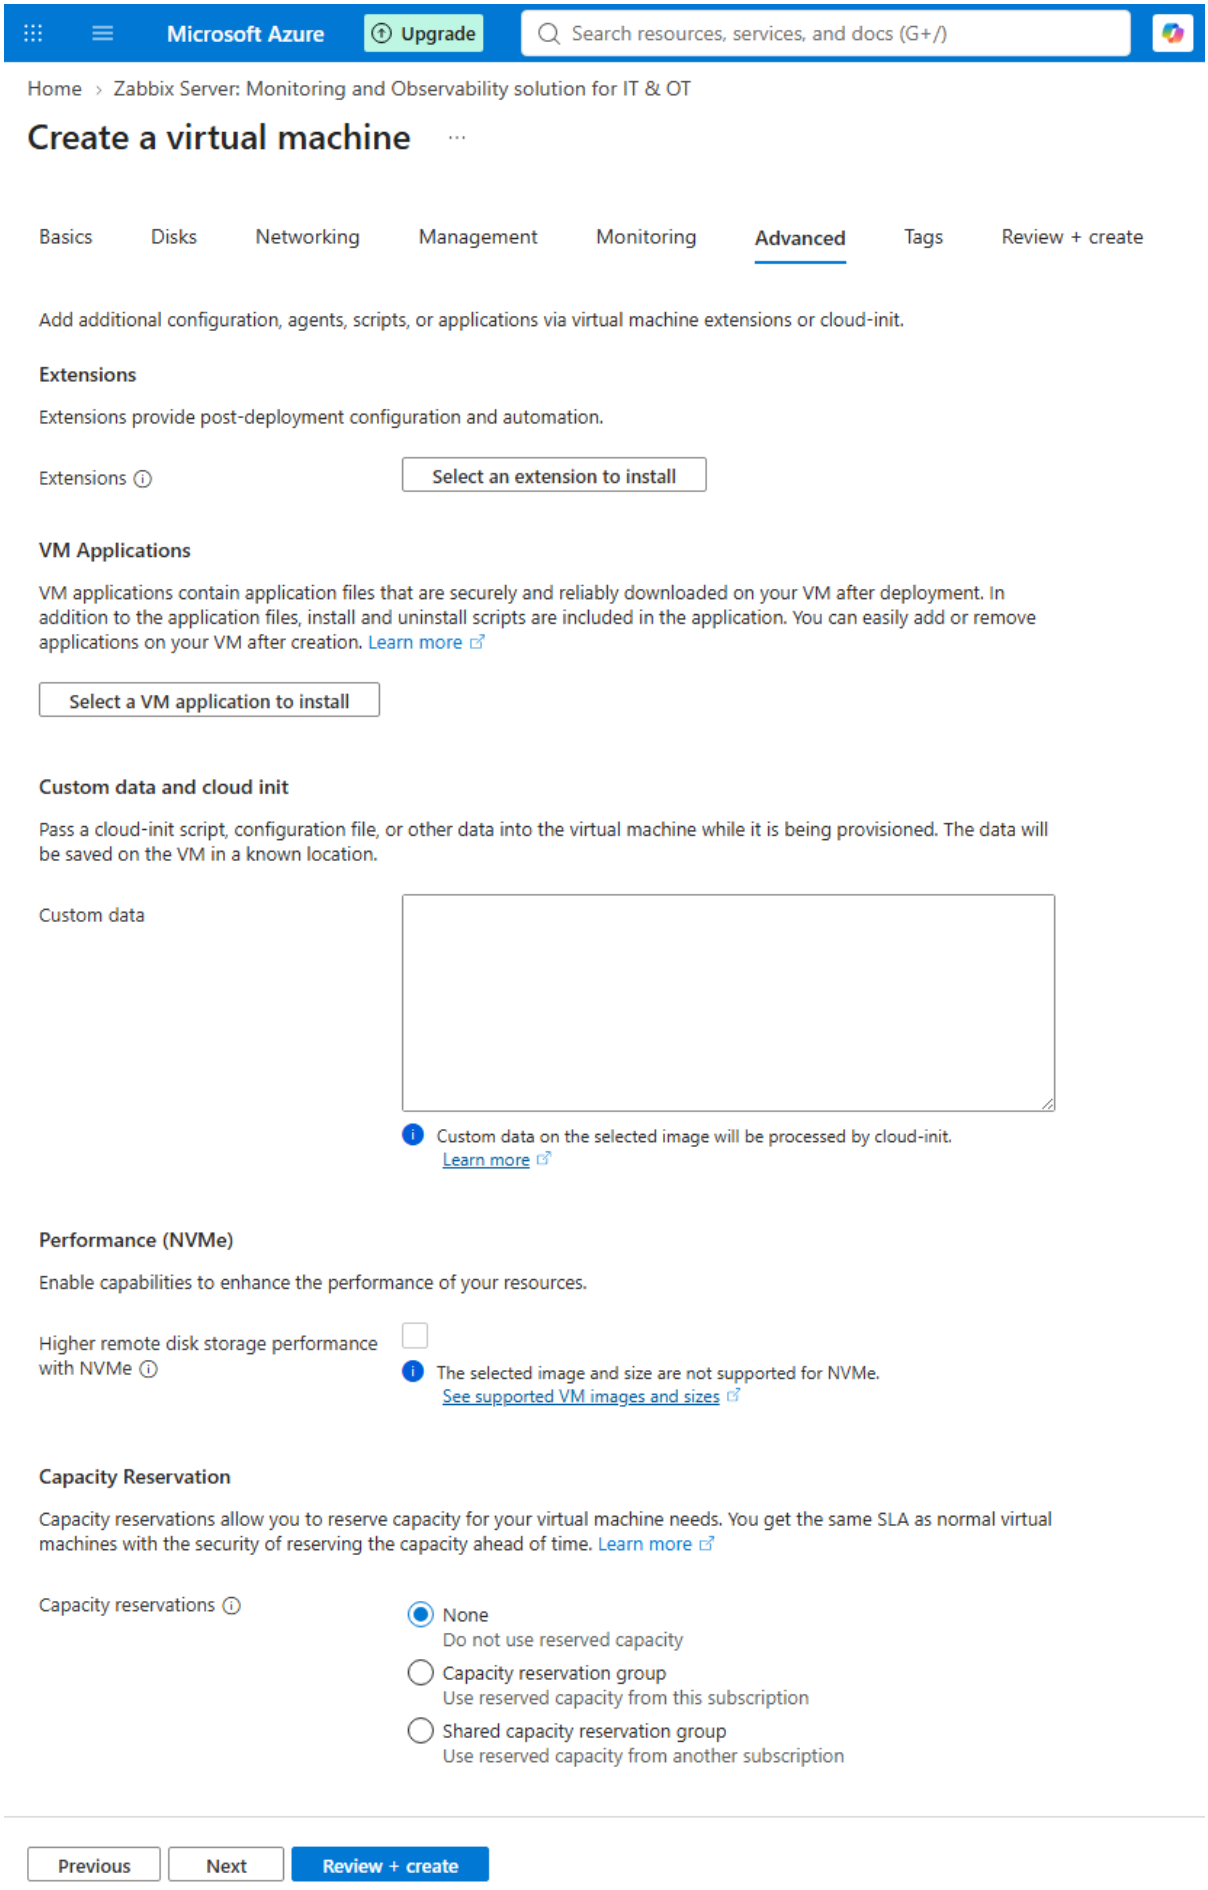
Management tab: parameters can be left by default.



Monitoring tab: parameters can be left by default.



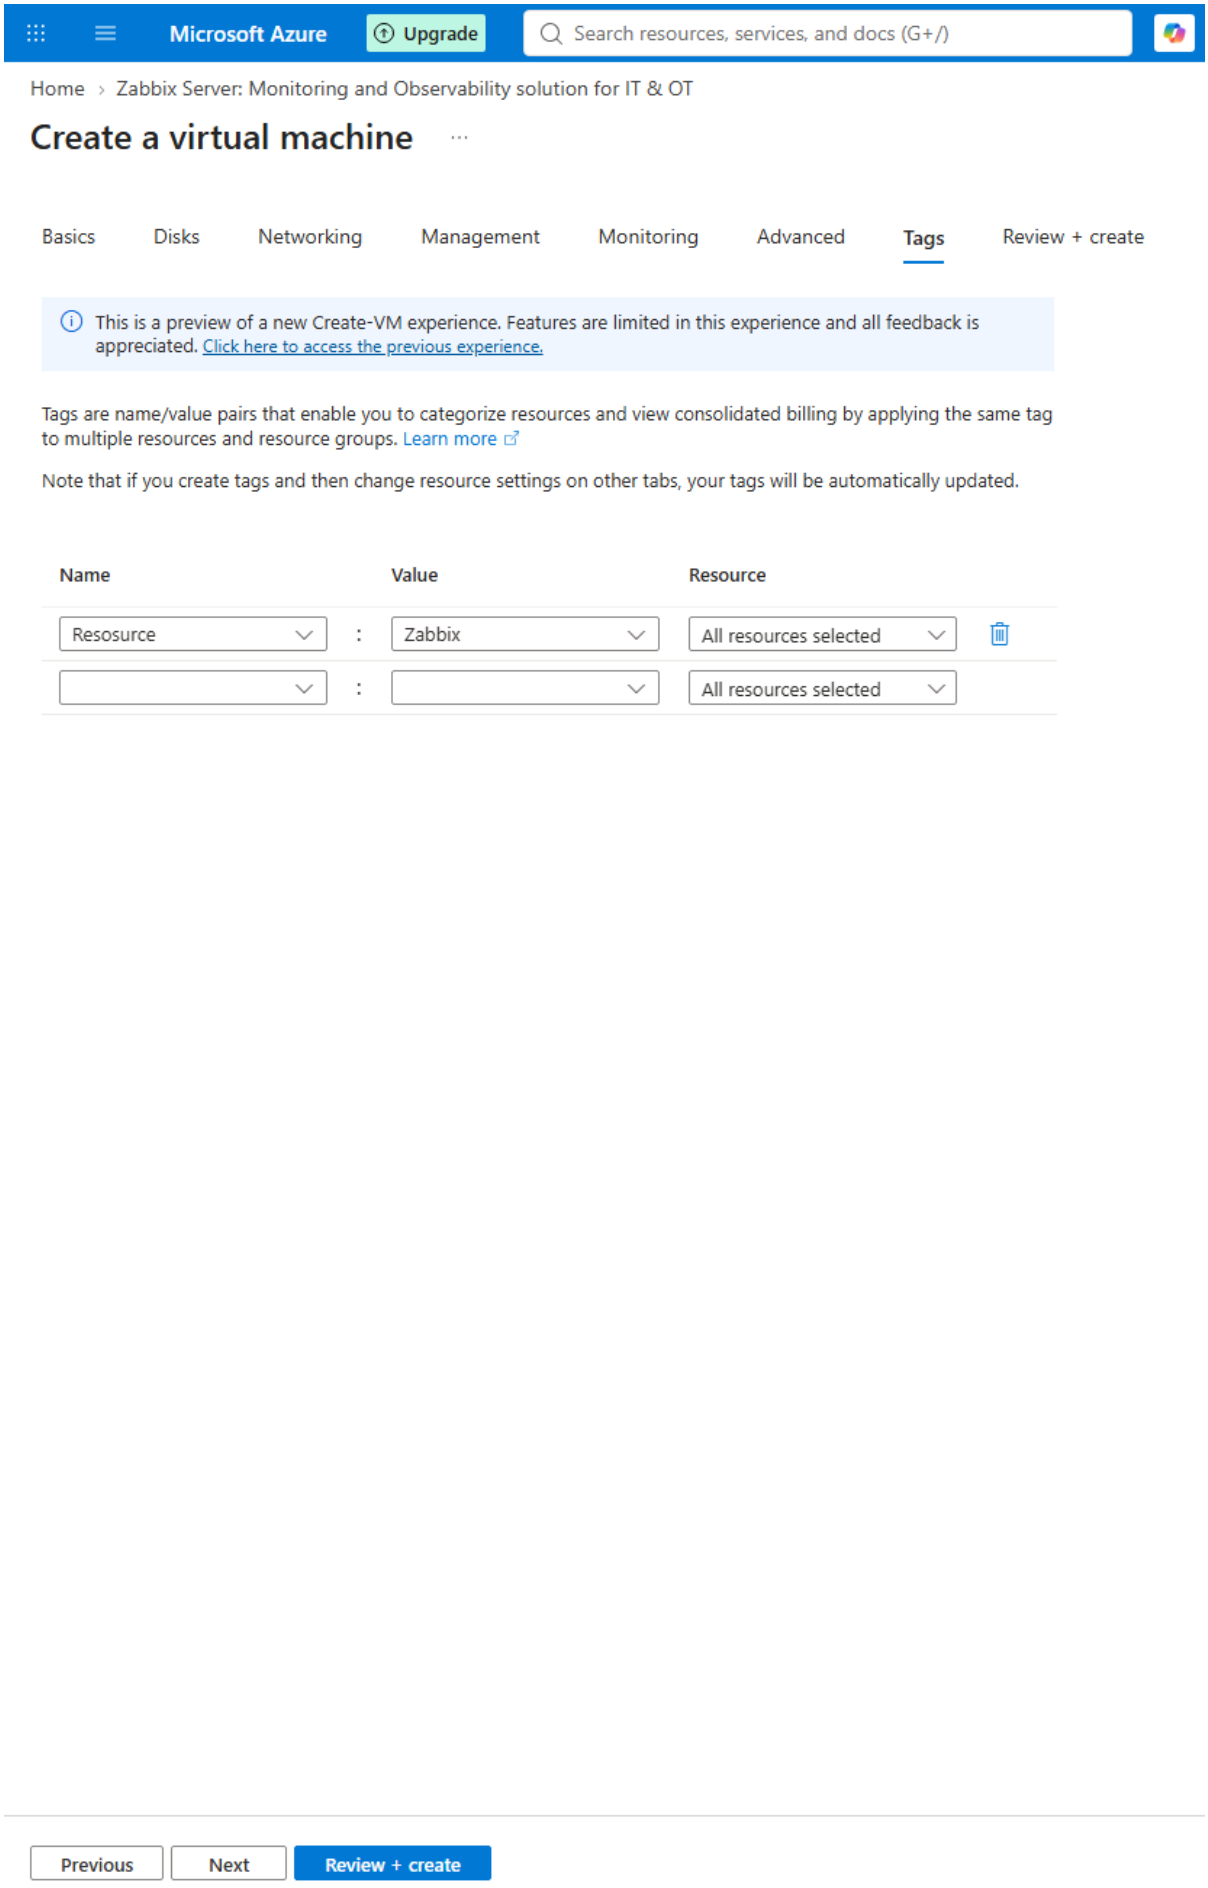
Advanced tab: parameters can be left by default.



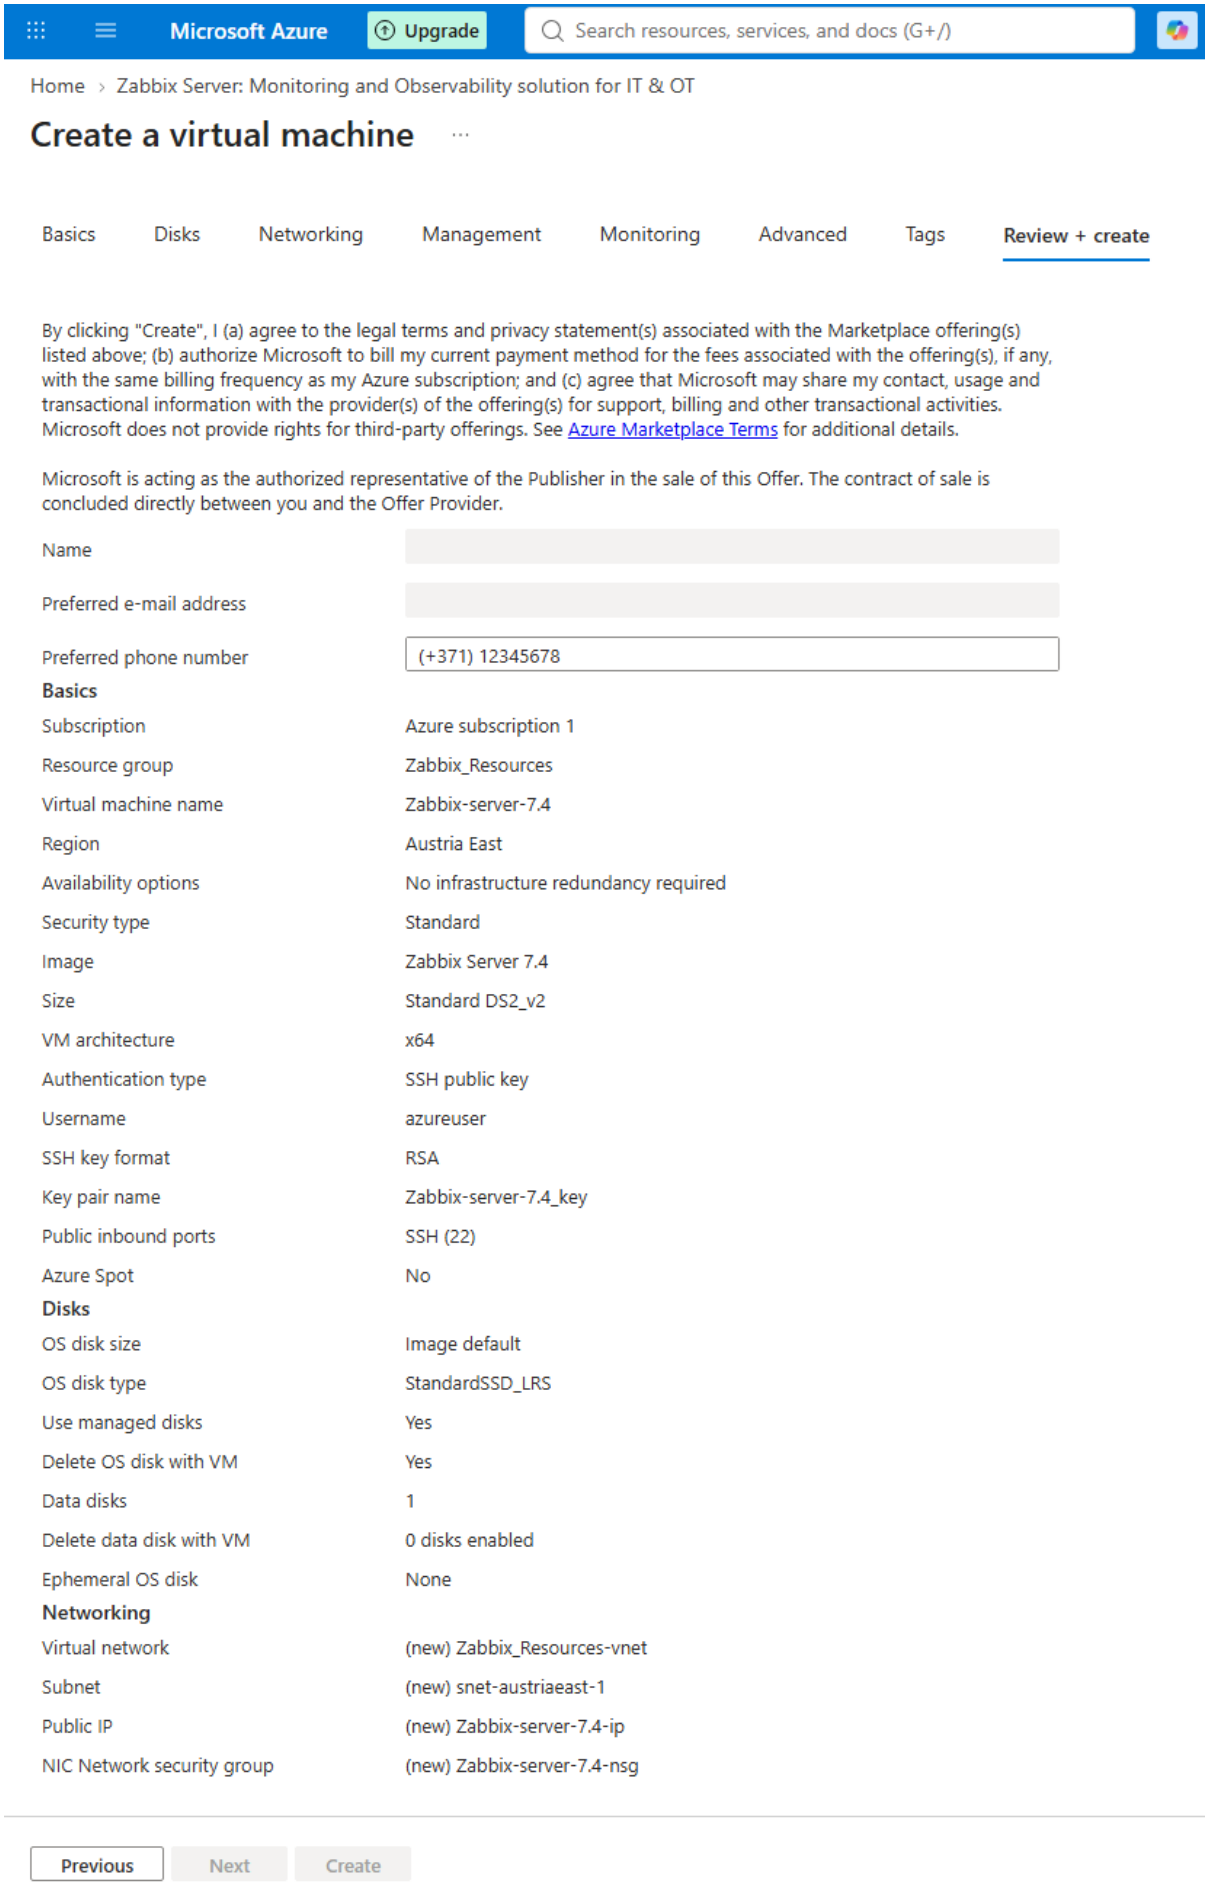
Tags tab:

- *Name*: Resource
- *Value*: Zabbix

- *Resource*: All resource selected



Confirm the VM details and create the VM.



You will be redirected to the VM deployment status page. Wait until the deployment is complete.

Once the deployment is complete, open the resource and note its public IP address. You will now have to connect to Zabbix server using SSH (or another available method) to obtain Zabbix frontend login credentials.

```
********************************************************************************

Zabbix frontend credentials:

Username: Admin

Password: zfpvagcs0s+srUUK+EM

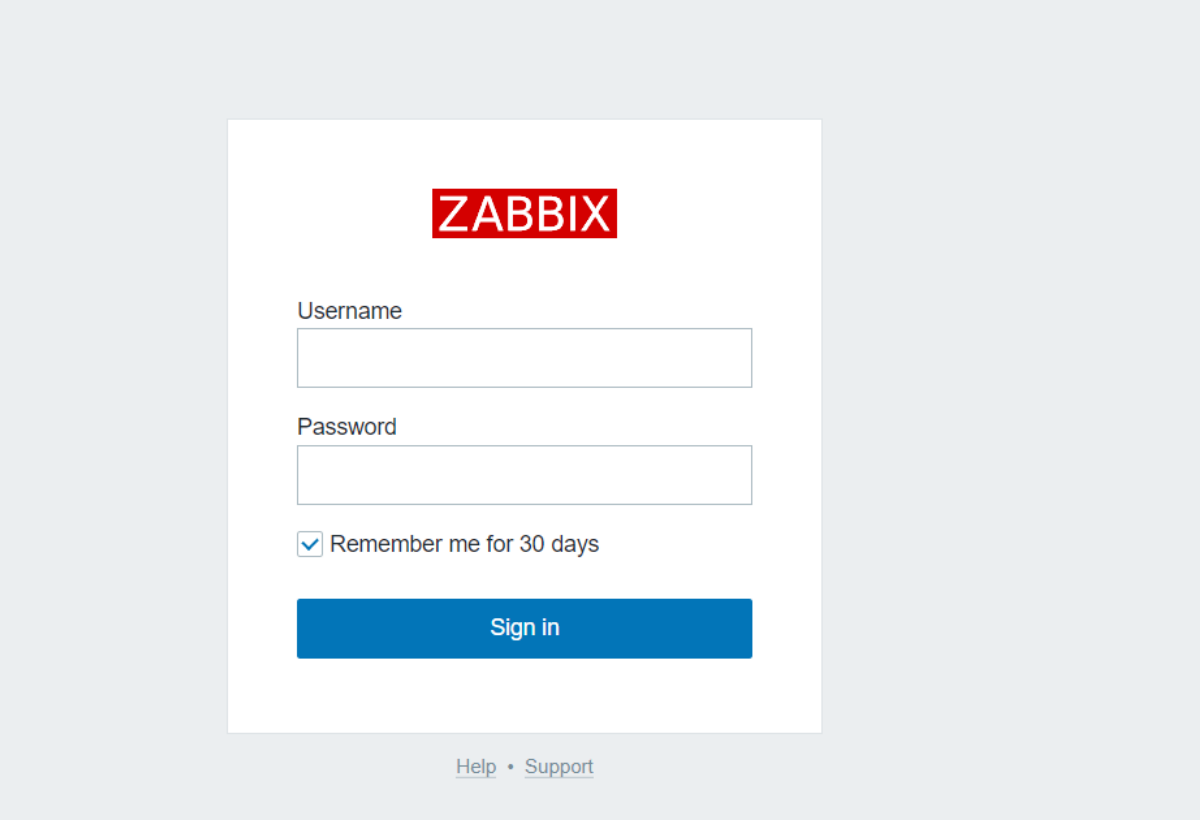

To learn about available professional services, including technical suppport and
 training, please visit https://www.zabbix.com/services

Official Zabbix documentation available at https://www.zabbix.com/documentation/
current/

Note! Do not forget to change timezone PHP variable in /etc/php-fpm.d/zabbix.con
f file.


********************************************************************************
[azureuser@Zabbix-server-6 ~]$
```

Access Zabbix frontend by opening the public IP address of Zabbix appliance in a web browser. To log in, use the username and password obtained in the previous step.



Accessing the Database

You can access Zabbix database using the `root` user without entering a password. The credentials are stored in the `/root/.my.cnf` configuration file.

Azure deployment guide for Zabbix proxy

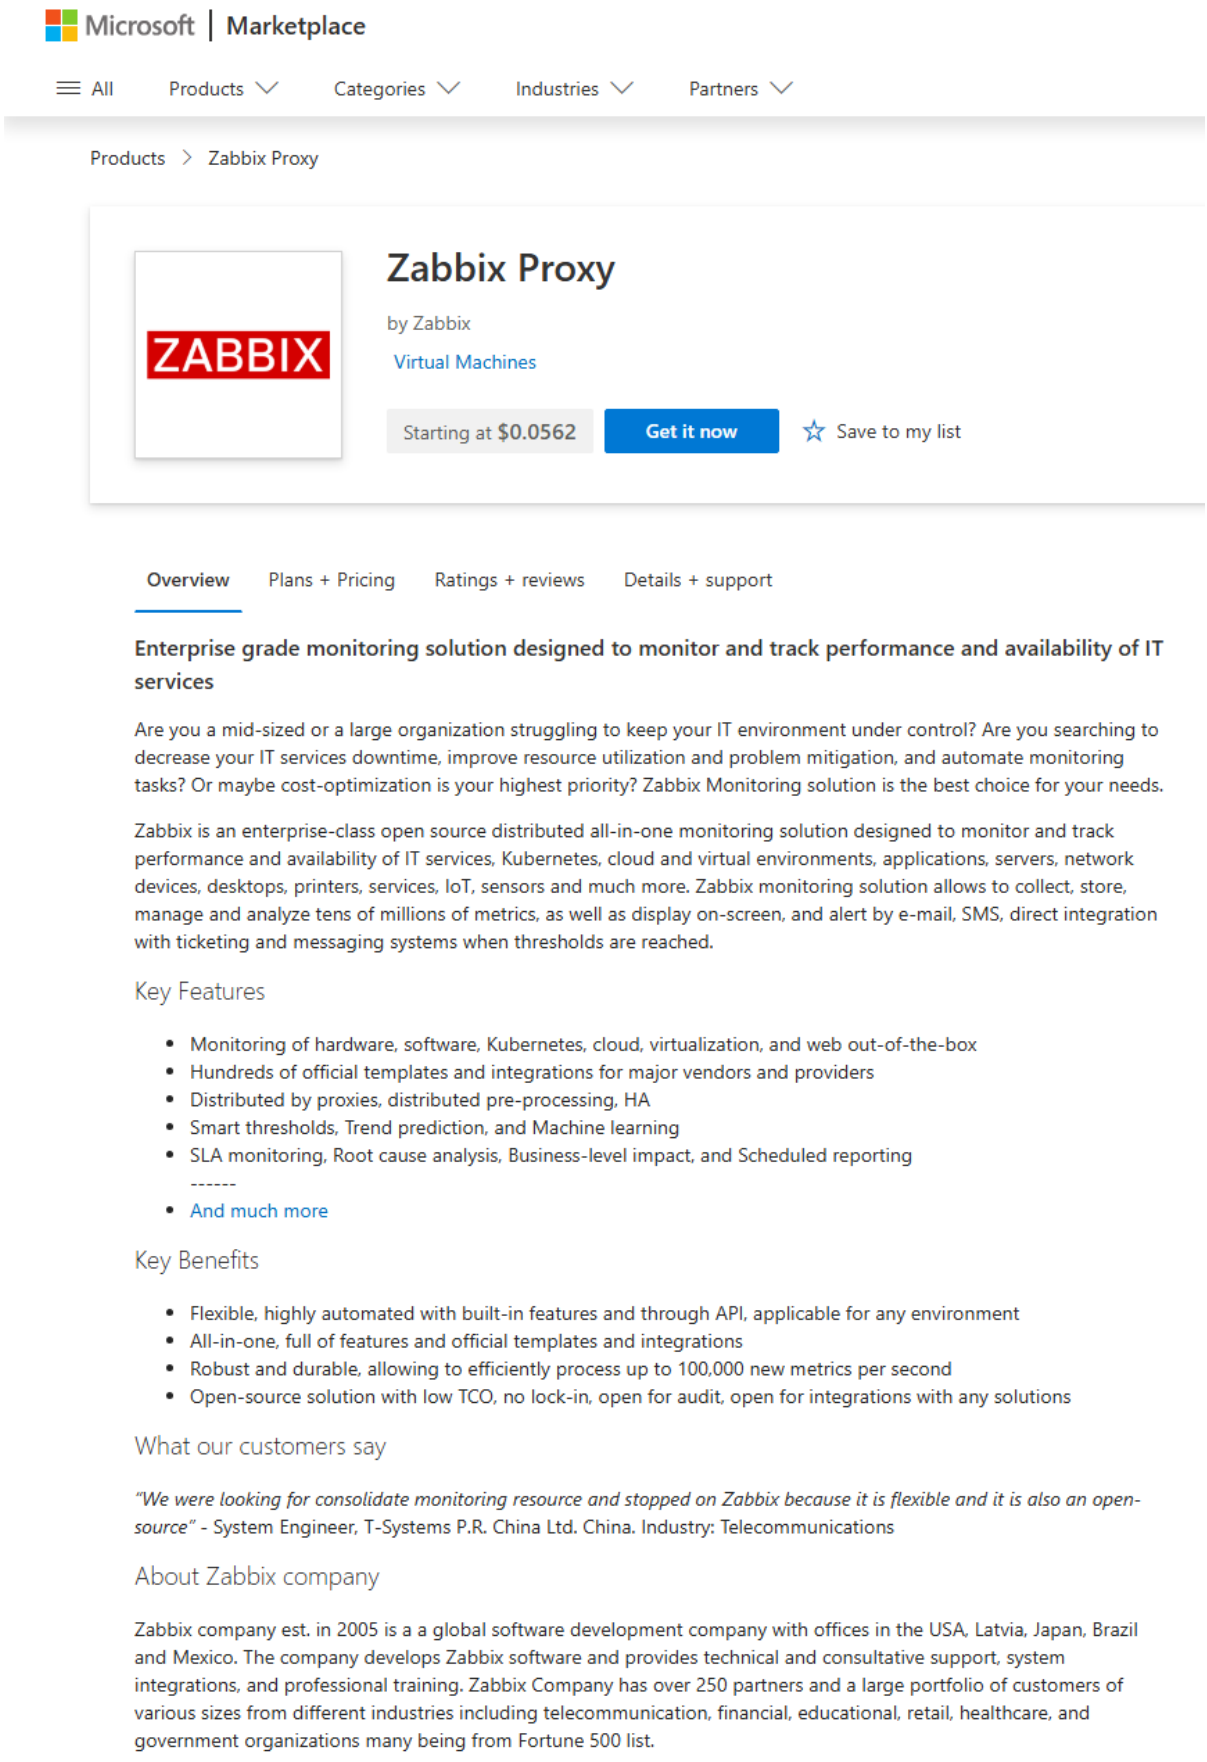
Overview

This page describes how to deploy Zabbix proxy using Azure.

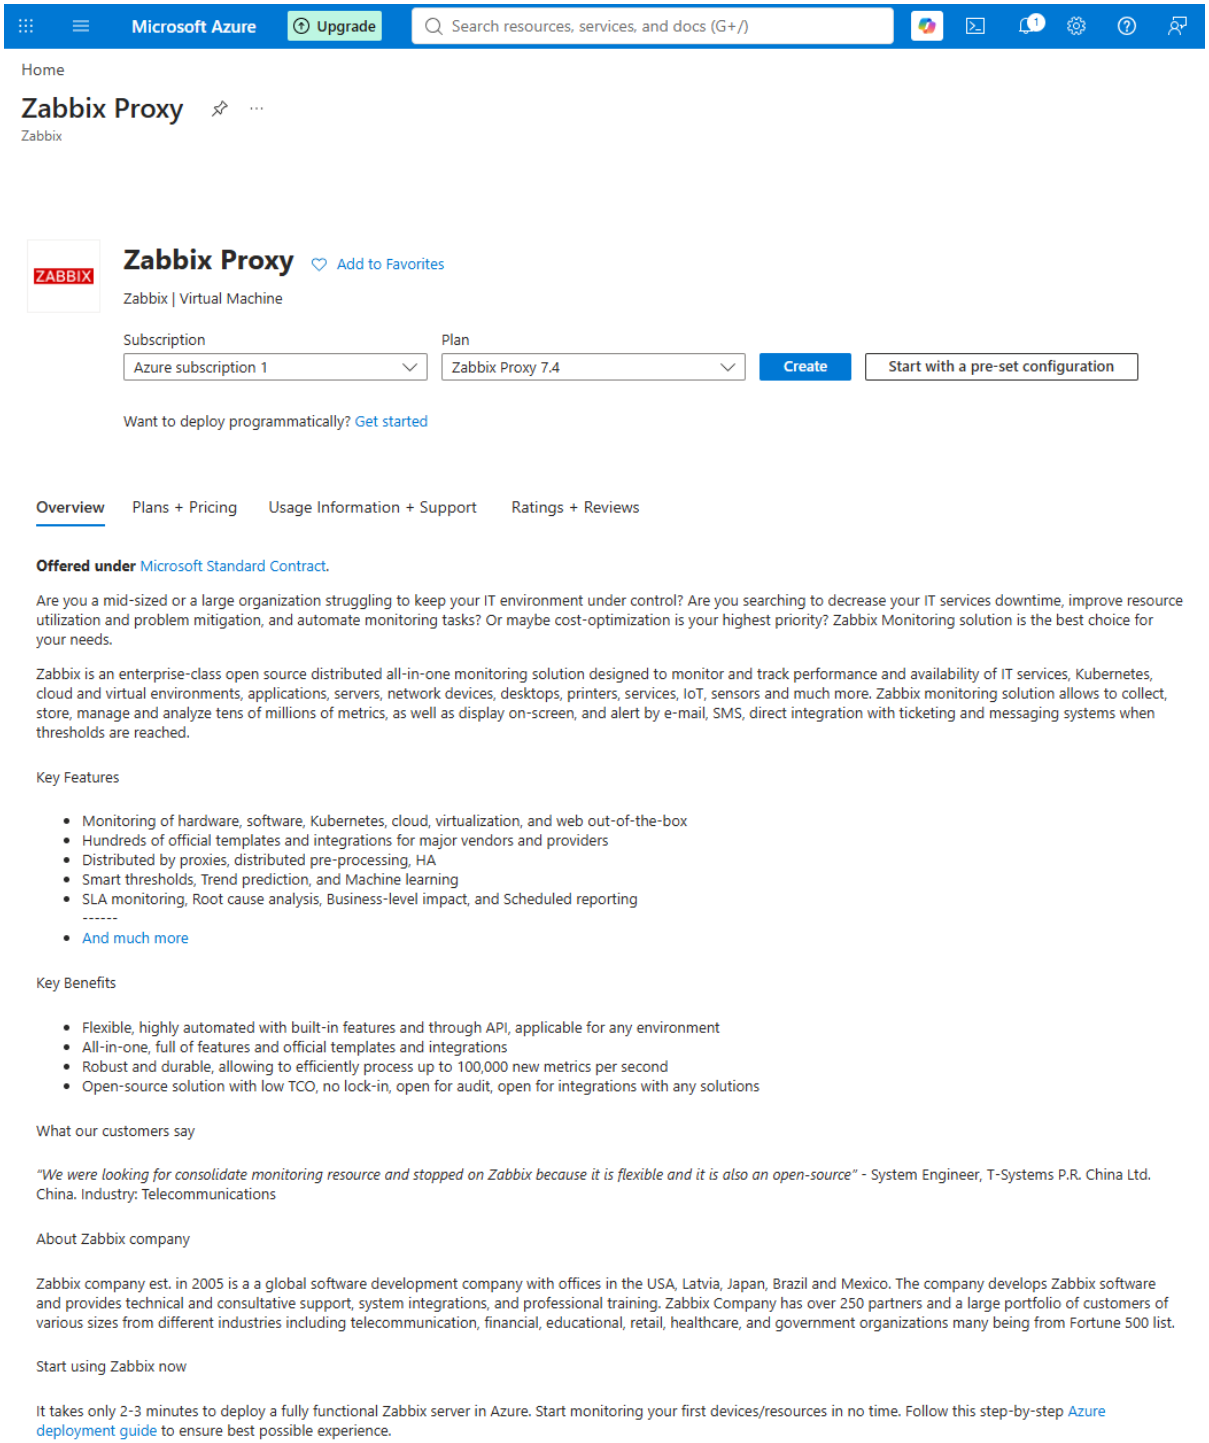
Attention:
Several Zabbix versions are available for deployment on the platform. Make sure to read the documentation for the same Zabbix version that you have deployed. Documentation for different Zabbix versions can be selected from the version dropdown menu.

Deploying Zabbix proxy

To deploy Zabbix proxy, first, navigate to the Zabbix Azure marketplace page. Click *Get it now*.



Here you can read additional information about the current deployment, such as the general overview, pricing plans, and more. Click either *Create* or *Start with a pre-set configuration*.

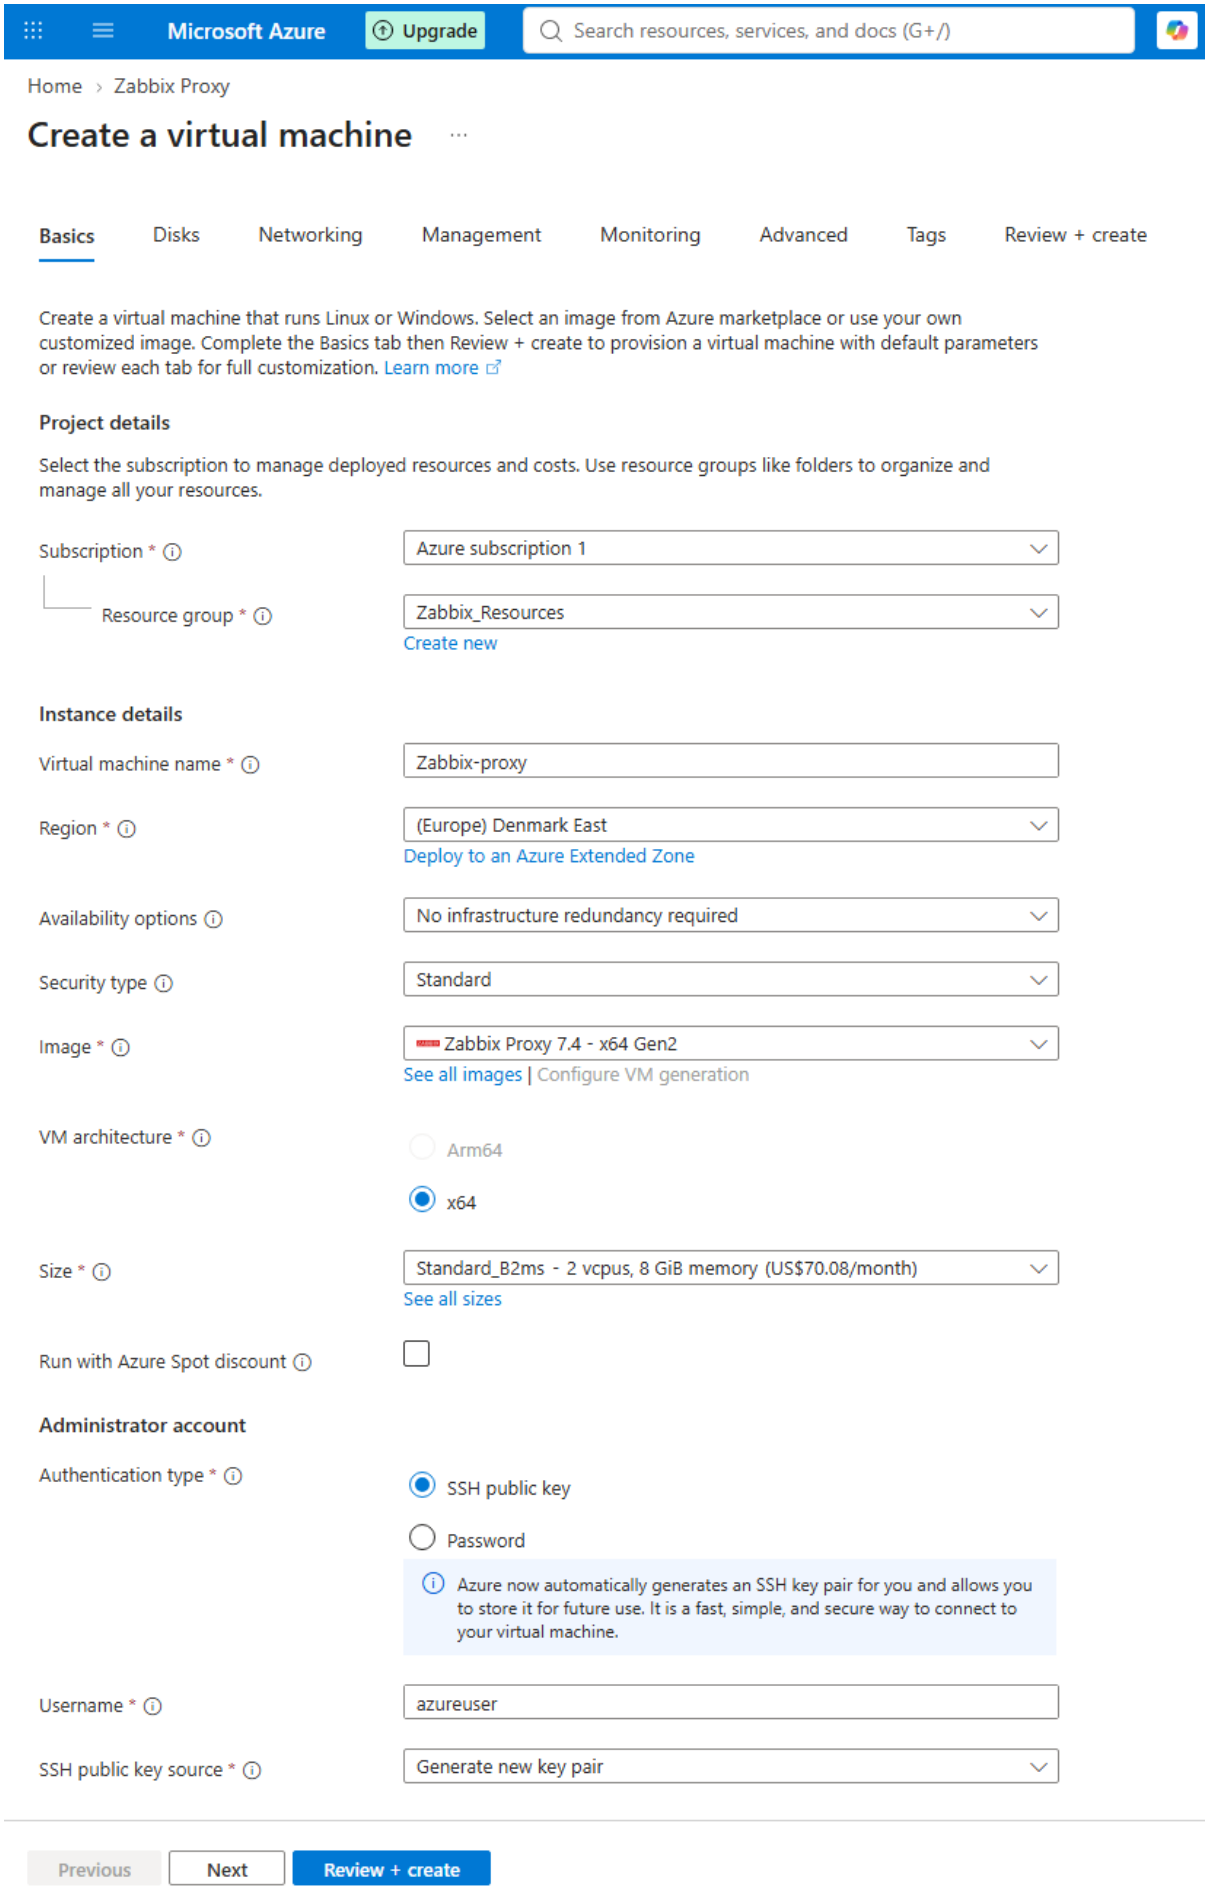


Proceed with the detailed VM configuration: select the VM size, administrator account, disk, network, management, and other settings. These settings may vary depending on your infrastructure requirements.

You can use the following configuration as an example:

Basics tab:

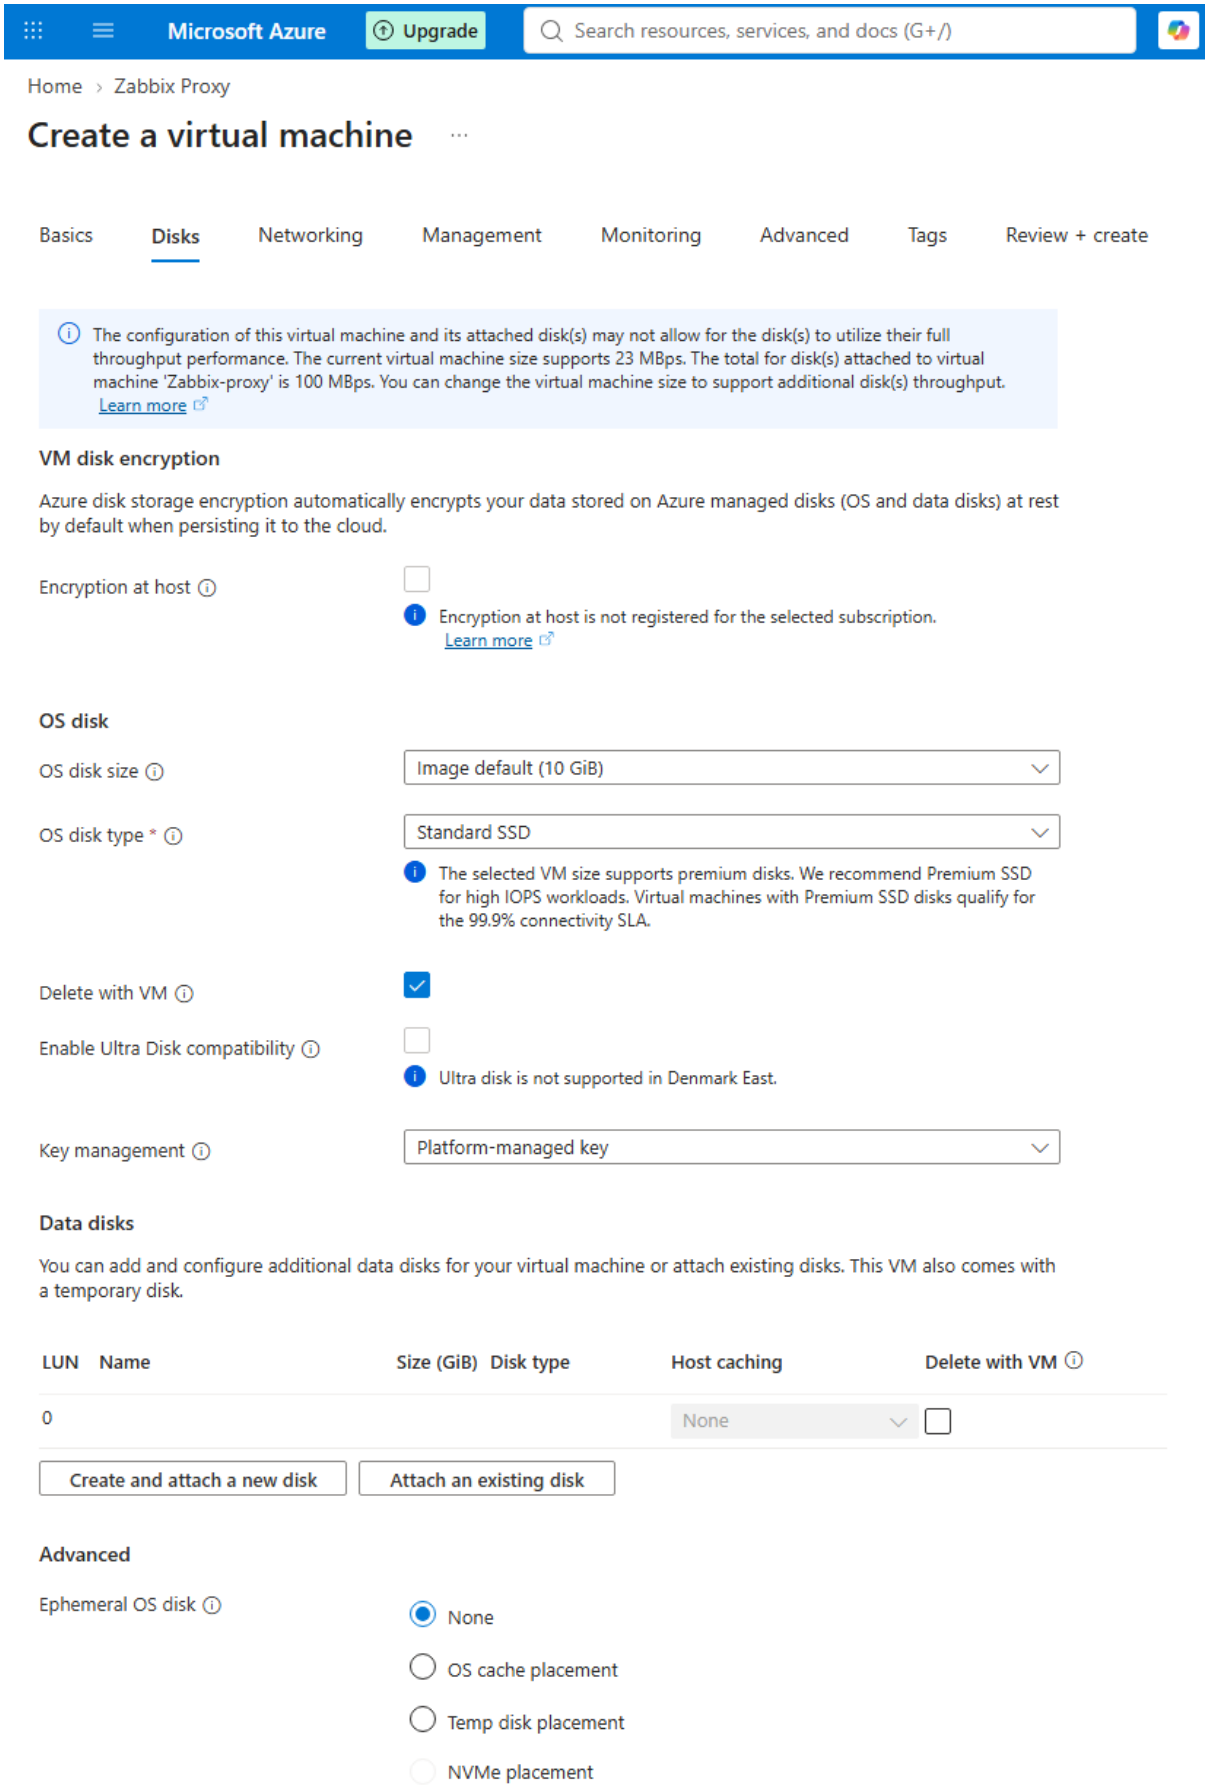
- *Subscription*: Azure subscription 1
- *Resource group*: Zabbix_Resources
- *Virtual machine name*: Zabbix-proxy
- *Availability options*: No infrastructure redundancy required
- *Security type*: Standard
- *Image*: Zabbix Proxy 7.4 - x64 Gen2
- *Size*: Standard_B2ms - 2 vcpus, 8 GiB memory
- *Authentication type*: SSH public key
- *Username*: azureuser
- *SSH public key source*: Generate new key pair



Disks tab:

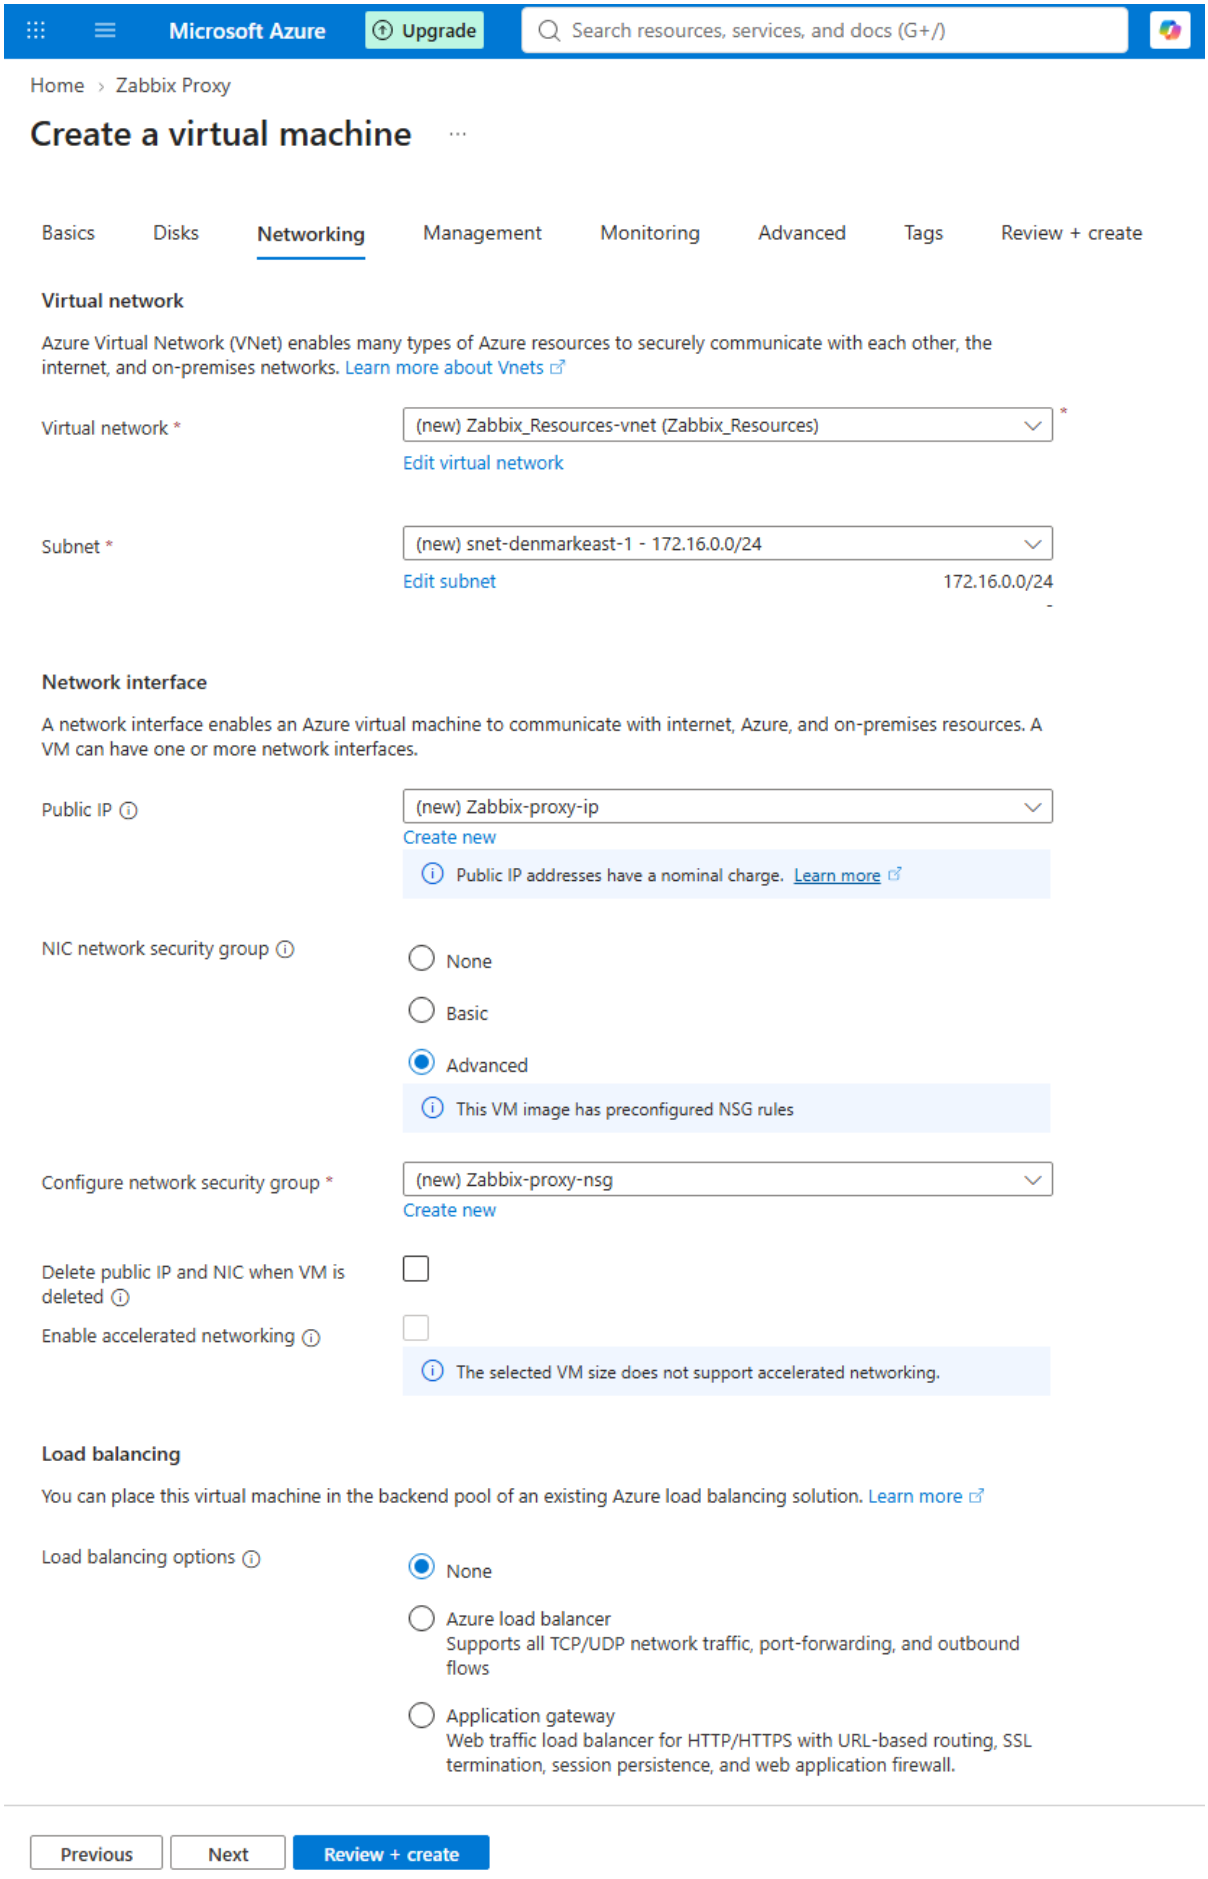
- *Encryption at host*: unchecked

- *OS disk type*: Standard SSD
- *Delete with VM*: checked
- *Enable Ultra Disk compatibility*: unchecked

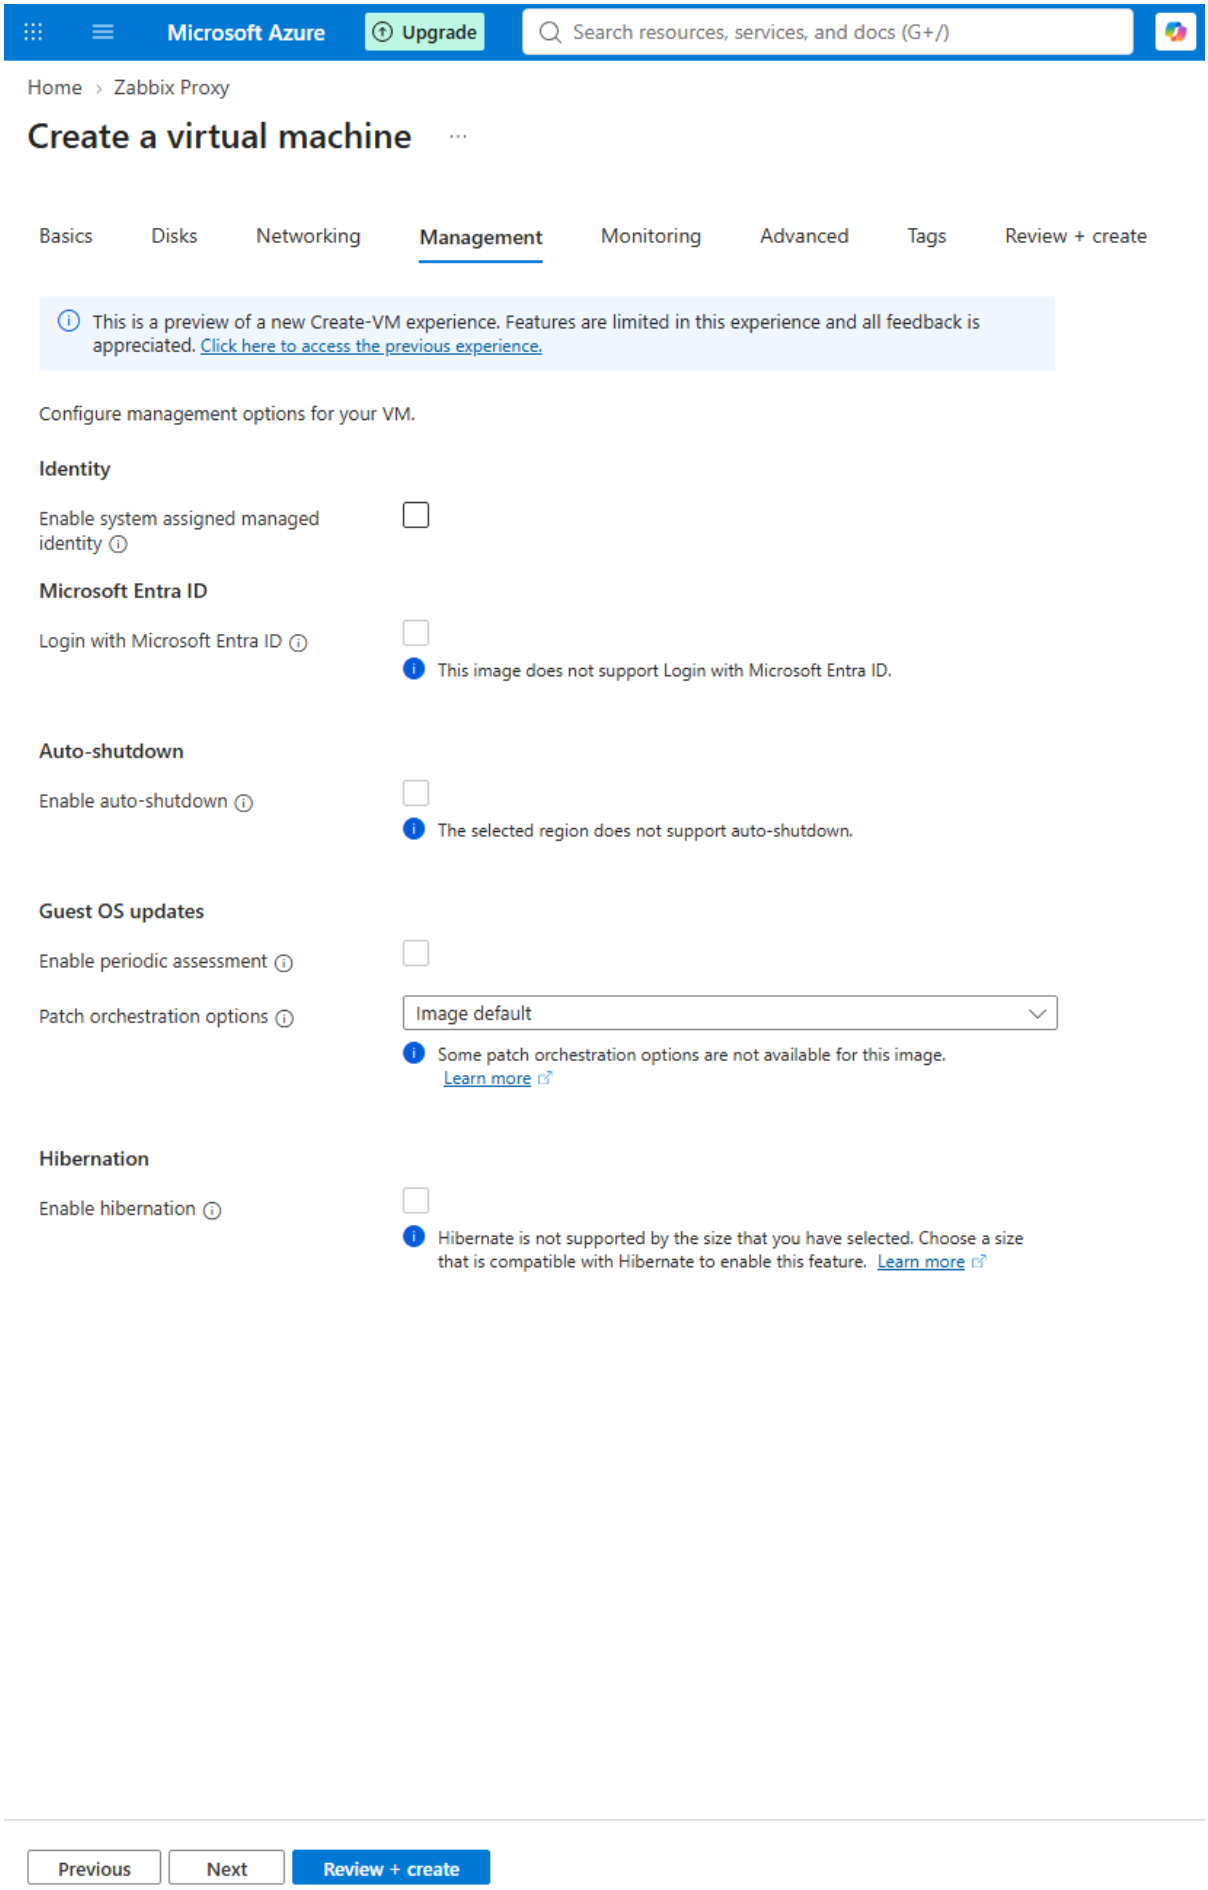


Networking tab:

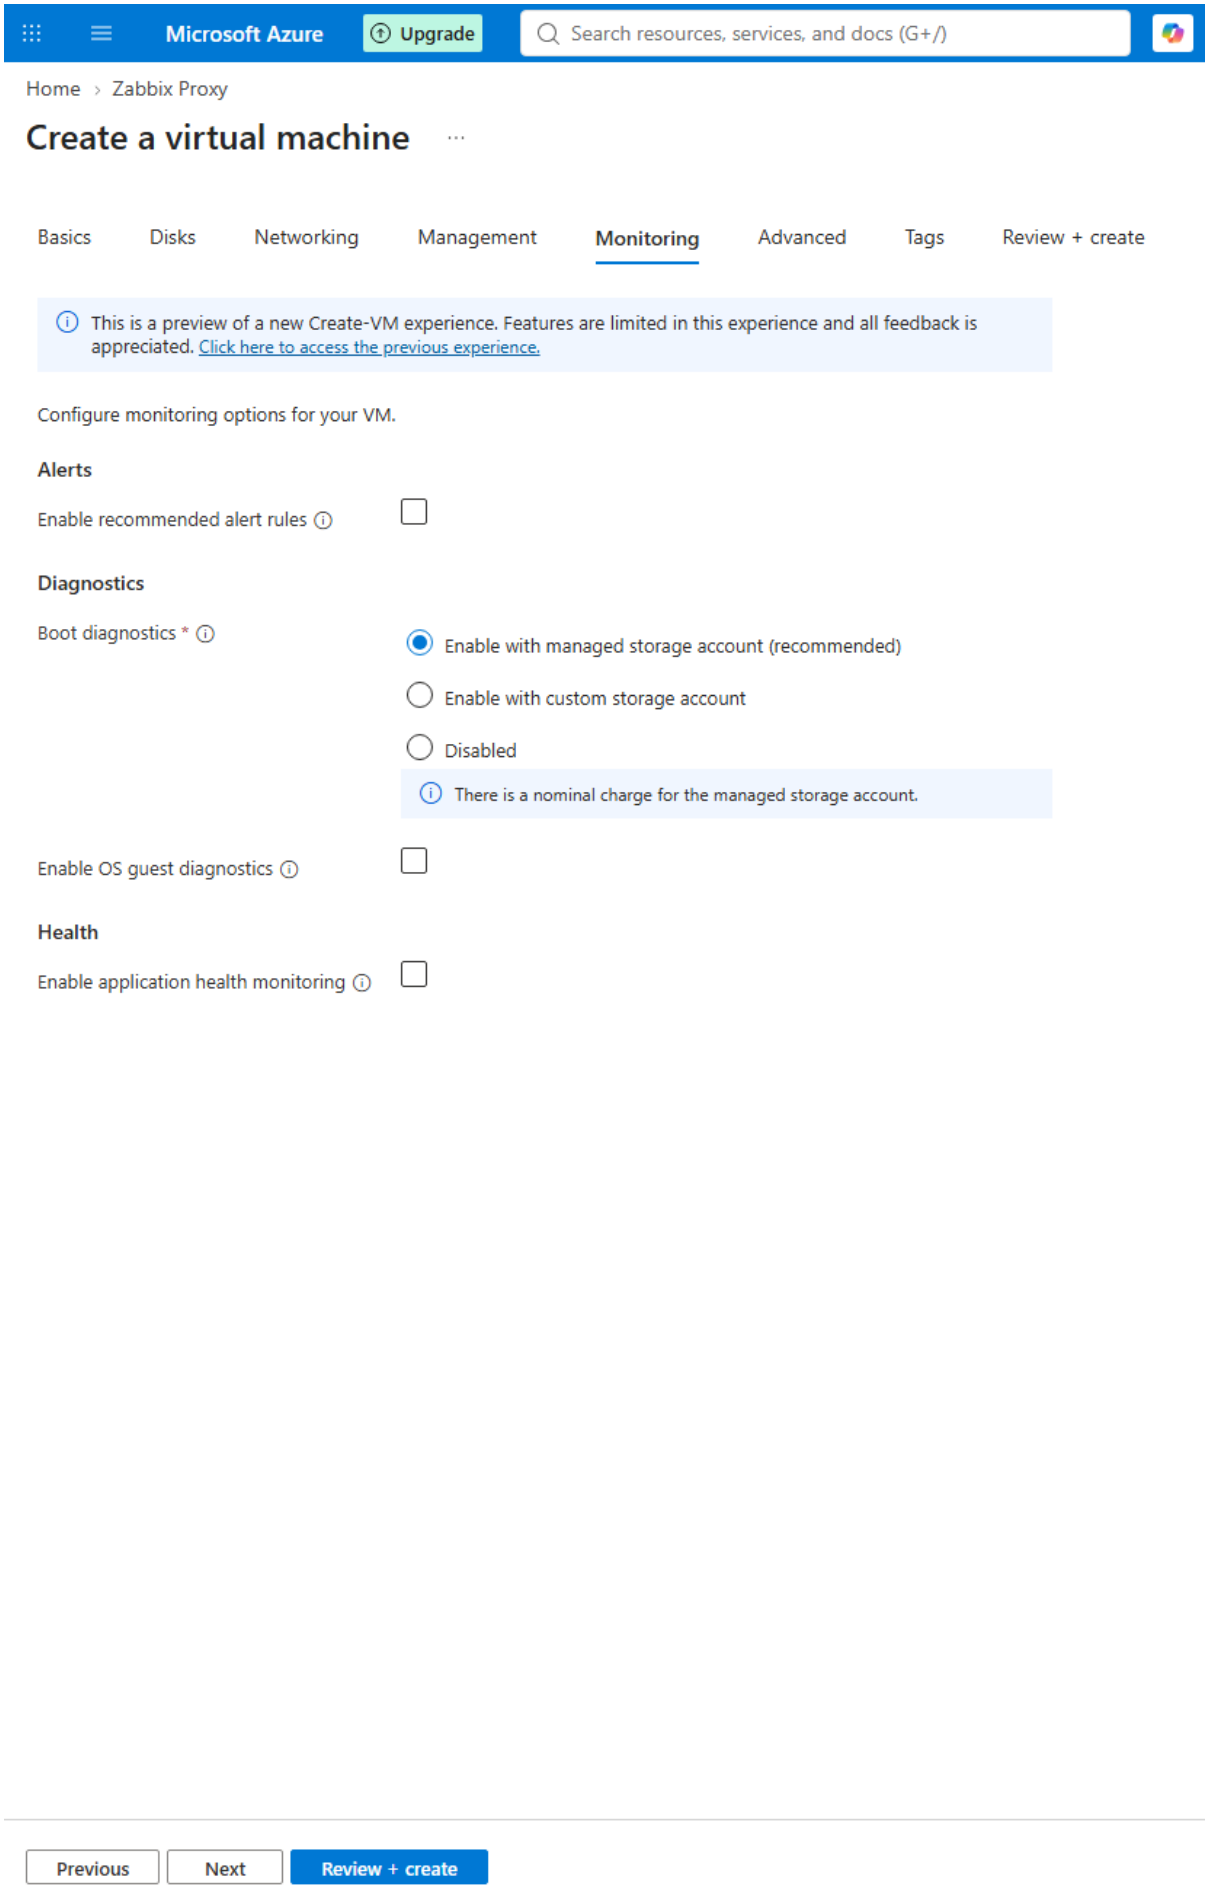
- *Virtual network*: Zabbix_Resources-vnet
- *Public IP*: Zabbix-proxy-ip
- *NIC network security group*: Advanced
- *Configure network security group*: Zabbix-proxy-nsg



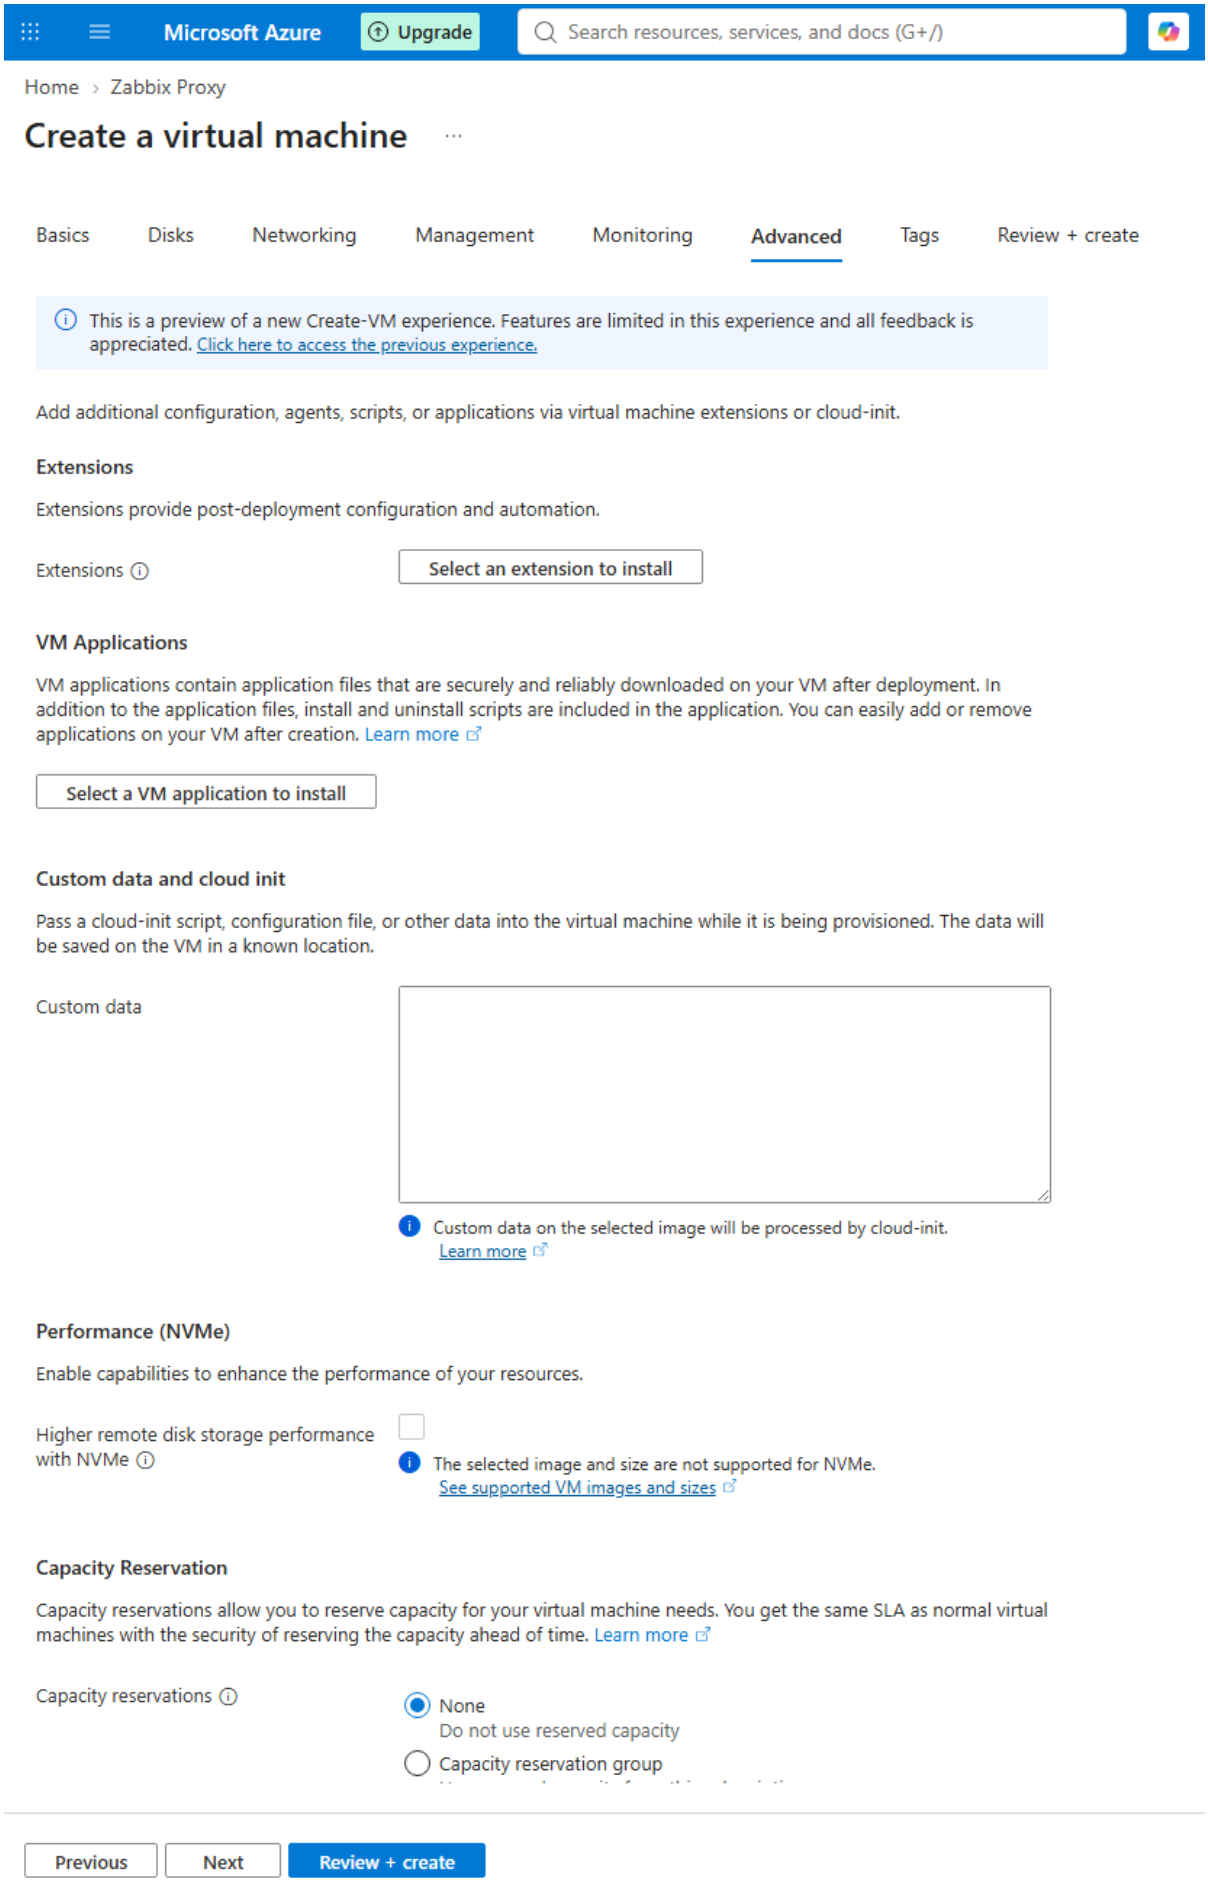
Management tab: parameters can be left by default.



Monitoring tab: parameters can be left by default.



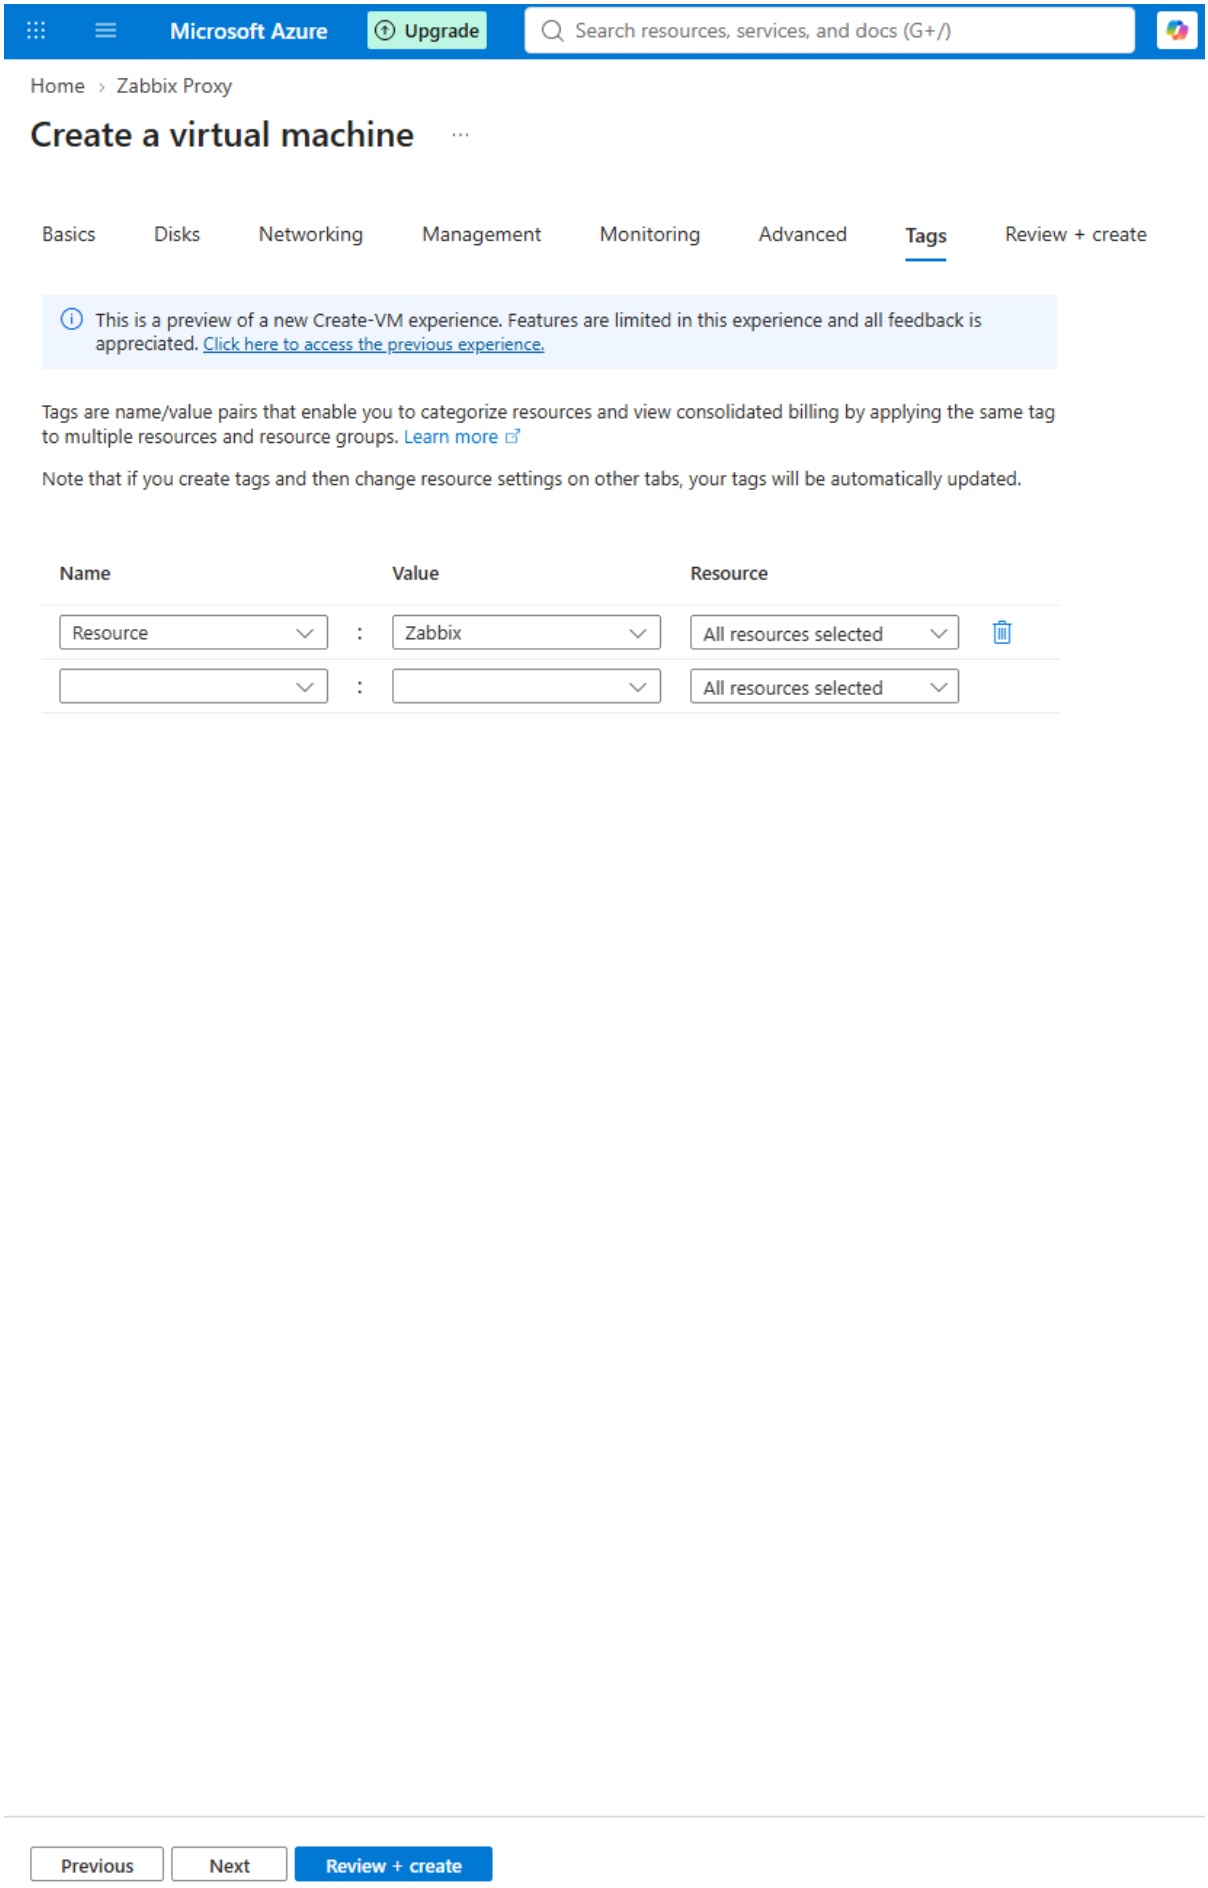
Advanced tab: parameters can be left by default.



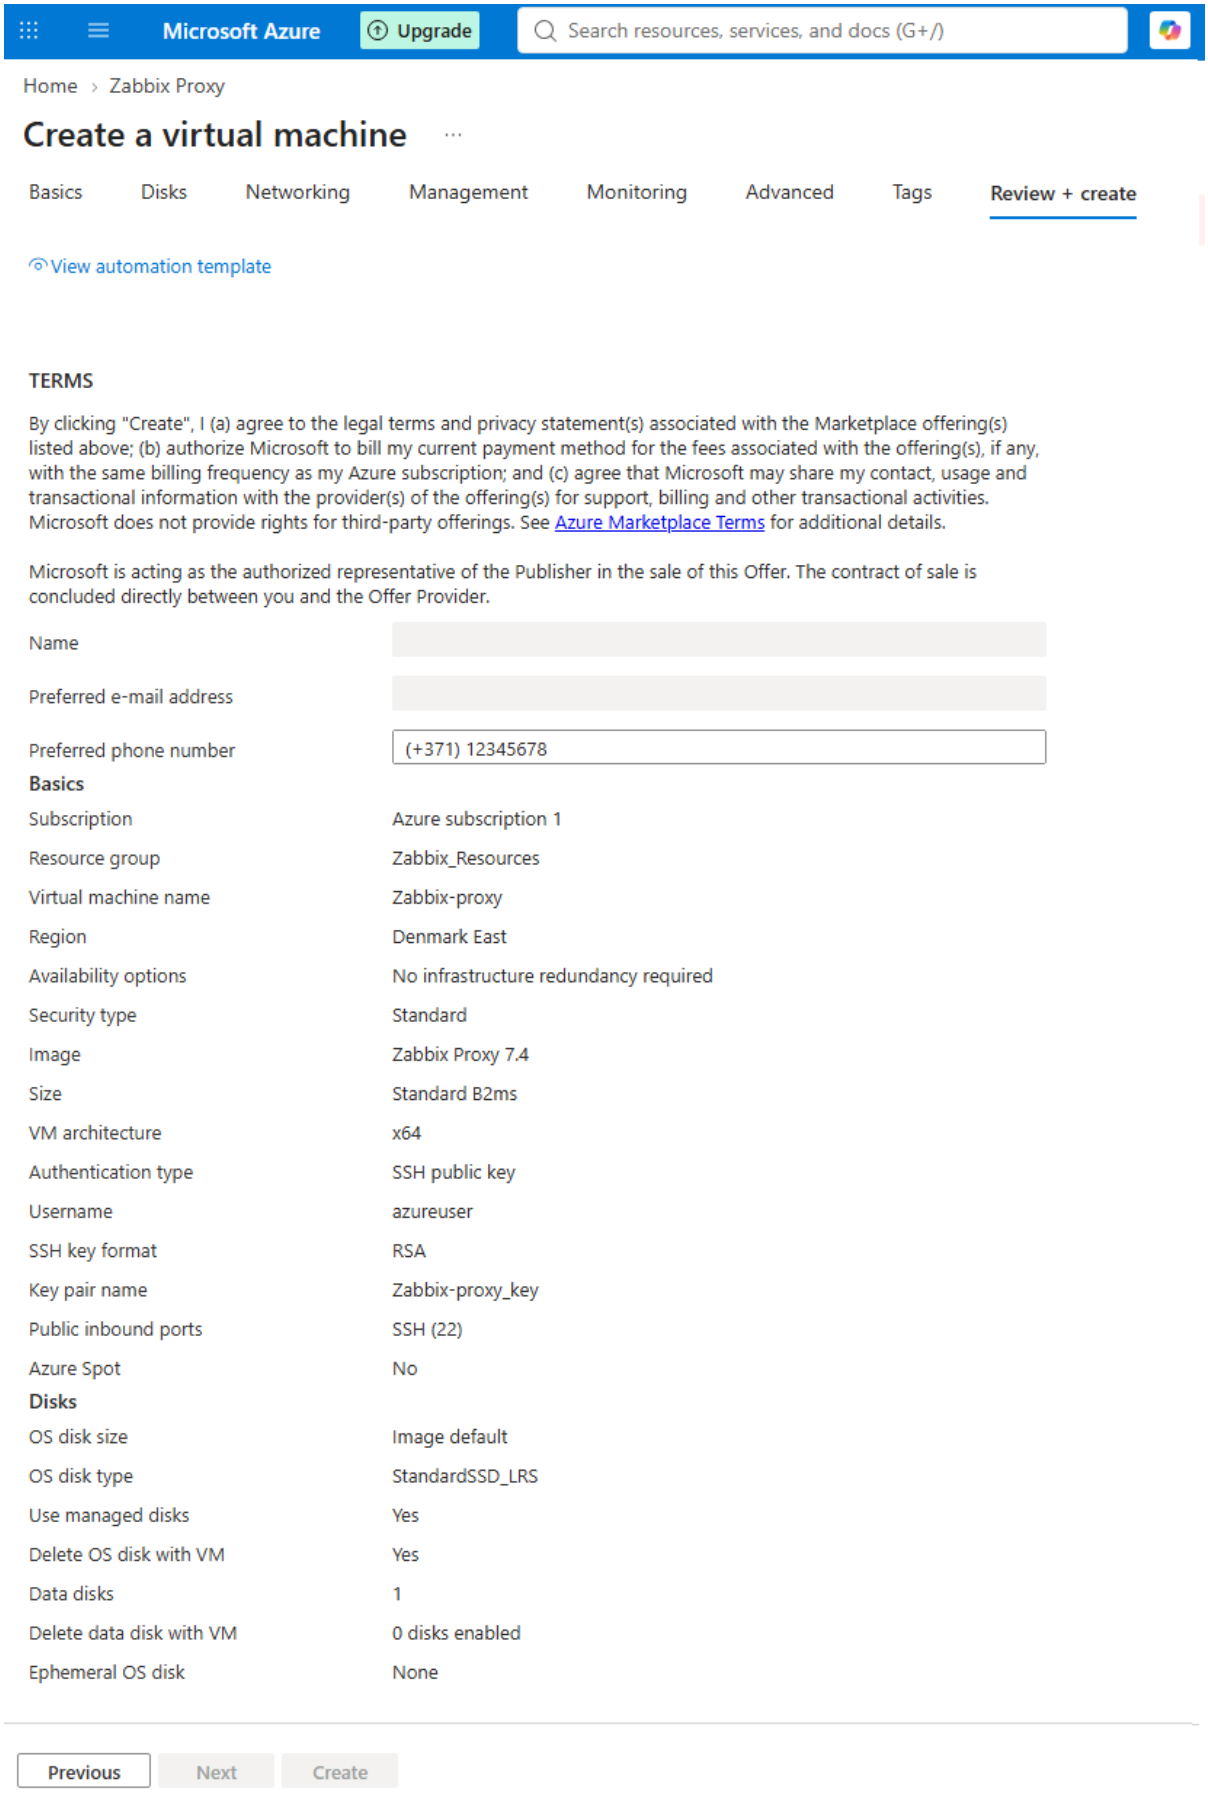
Tags tab:

- *Name*: Resource
- *Value*: Zabbix

- *Resource*: All resource selected



Confirm the VM details and create the VM.



You will be redirected to the VM deployment status page. Wait until the deployment is complete.

Once the deployment is complete, open the resource and note its public IP address. Connect to Zabbix proxy using SSH and the administrator credentials created during VM deployment to configure the proxy.

```
********************************************************************************

To learn about available professional services, including technical suppport and
 training, please visit https://www.zabbix.com/services

Official Zabbix documentation available at https://www.zabbix.com/documentation/
current/

********************************************************************************
[azureuser@Zabbix-proxy ~]$
```

Zabbix proxy configuration file can be found in `/etc/zabbix/zabbix_proxy.conf`.

The initial proxy configuration consists of modifying the `ProxyMode`, `Server`, and `Hostname` parameters in Zabbix proxy configuration file.

The proxy mode can be changed by editing the `ProxyMode` parameter. Set `ProxyMode=0` for active mode or `ProxyMode=1` for passive mode.

The `Server` parameter should contain Zabbix server IP address.

The `Hostname` parameter should contain Zabbix proxy hostname. The hostname defined here should match the proxy name provided during proxy configuration in Zabbix frontend.

```
############ GENERAL PARAMETERS #################

### Option: ProxyMode
#       Proxy operating mode.
#       0 - proxy in the active mode
#       1 - proxy in the passive mode
#
# Mandatory: no
# Default:
# ProxyMode=0

### Option: Server
#       If ProxyMode is set to active mode:
#               IP address or DNS name (address:port) or cluster (address:port;address2:port) of Zabbix server to get
configuration data from and send data to.
#               If port is not specified, default port is used.
#               Cluster nodes need to be separated by semicolon.
#       If ProxyMode is set to passive mode:
#               List of comma delimited IP addresses, optionally in CIDR notation, or DNS names of Zabbix server.
#               Incoming connections will be accepted only from the addresses listed here.
#               If IPv6 support is enabled then '127.0.0.1', '::127.0.0.1', '::ffff:127.0.0.1' are treated equally
#               and '::/0' will allow any IPv4 or IPv6 address.
#               '0.0.0.0/0' can be used to allow any IPv4 address.
#               Example: Server=127.0.0.1,192.168.1.0/24,::1,2001:db8::/32,zabbix.example.com
#
# Mandatory: yes
# Default:
# Server=

Server=127.0.0.1
```

```
### Option: Hostname
#       Unique, case sensitive Proxy name. Make sure the Proxy name is known to the server!
#       Value is acquired from HostnameItem if undefined.
#
# Mandatory: no
# Default:
# Hostname=

Hostname=Zabbix proxy
```

Open Zabbix frontend and navigate to *Administration* > *Proxies*. Click the *Create proxy* button.

Enter the same proxy name as the value of the *Hostname* parameter in the proxy configuration file.

Select the proxy mode and provide the proxy IP address (optional for active proxies).

Click the *Create proxy* button to create a proxy.

Proxy database location

The SQLite proxy database can be found in `/tmp/zabbix_proxy.sqlite3` and is specified by the `DBName` parameter in Zabbix proxy configuration file.

Google Cloud deployment guide

Overview

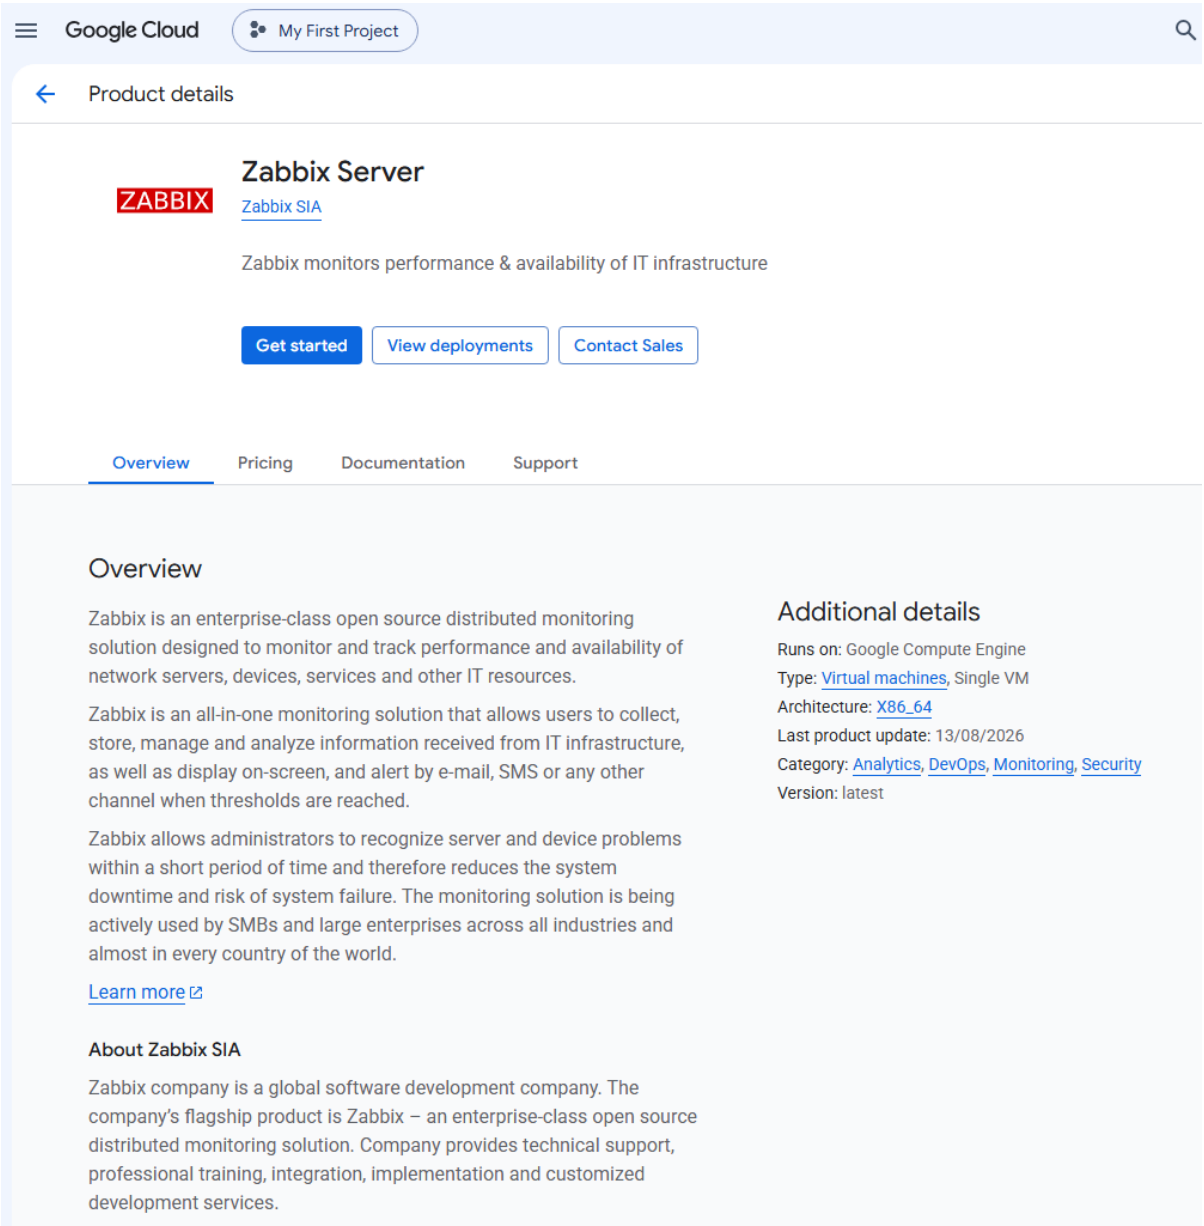
This page describes how to deploy Zabbix using Google Cloud Marketplace.

Attention:
Several Zabbix versions are available for deployment on the platform. Make sure to read the documentation for the same Zabbix version that you have deployed. Documentation for different Zabbix versions can be selected from the version dropdown menu.

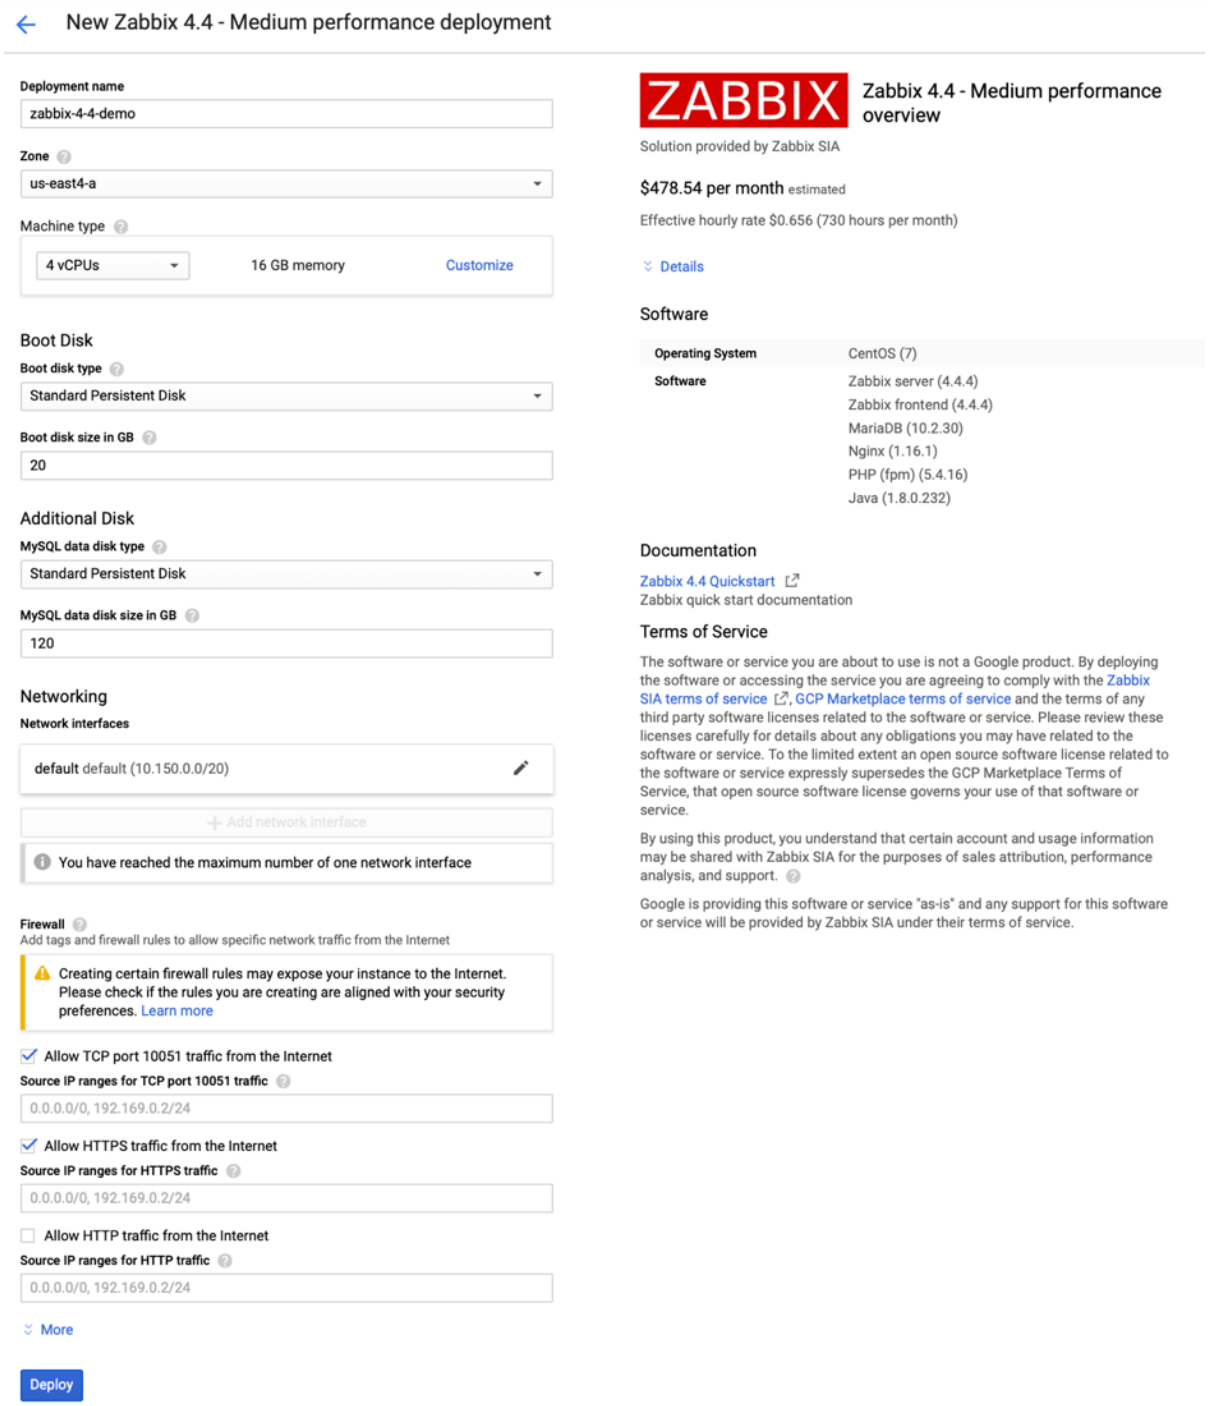
Deploying appliance

First, deploy Zabbix appliance. That's really easy with Google Cloud Marketplace. Simply go to Zabbix Server product on the marketplace.

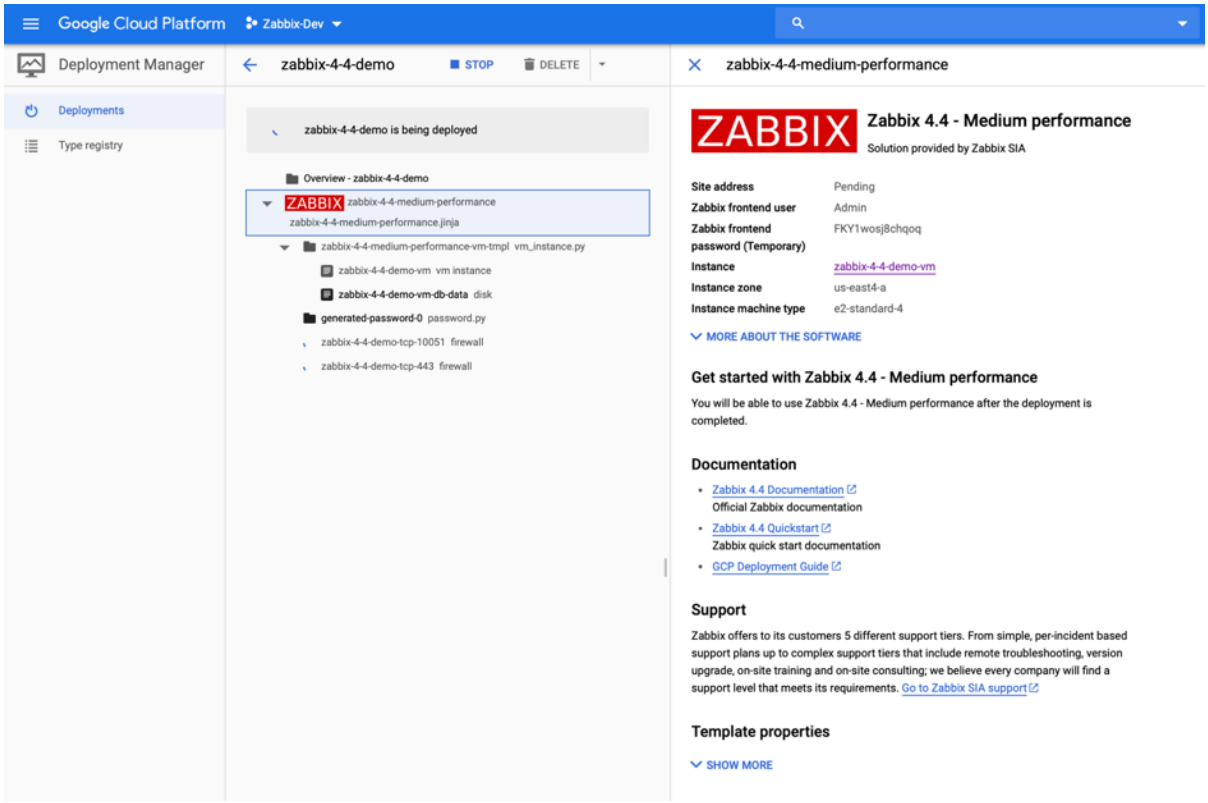
Click *Get started*.



You can take the default settings or customize them. When complete, click *Deploy*.



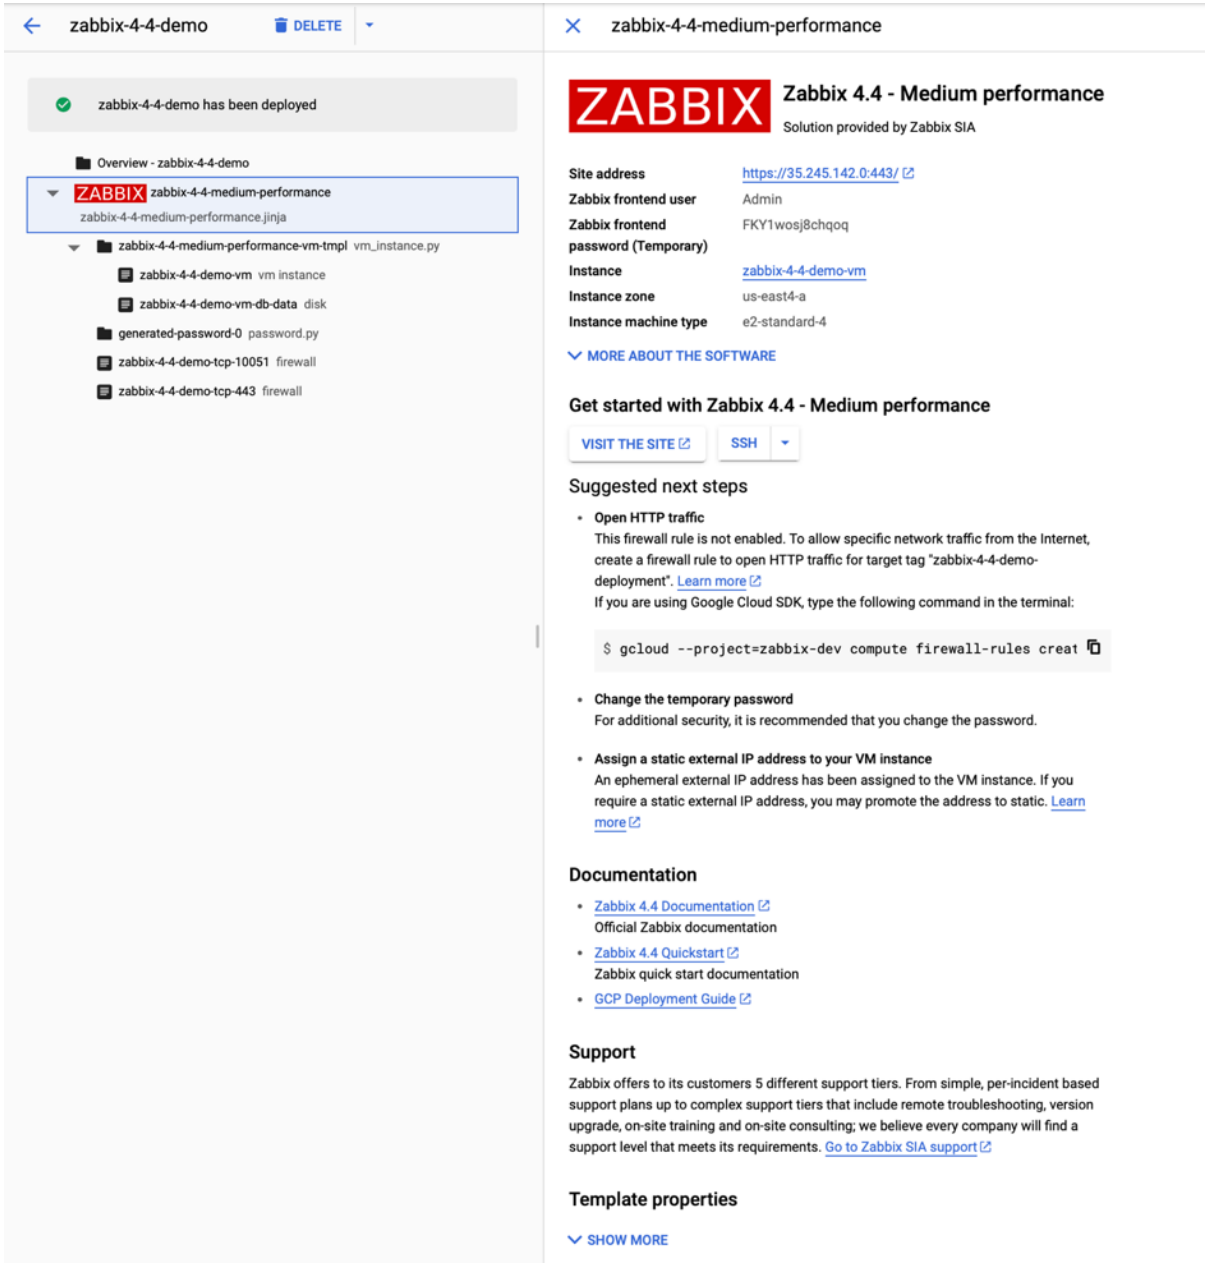
The deployment is now in progress. Wait for it to complete.



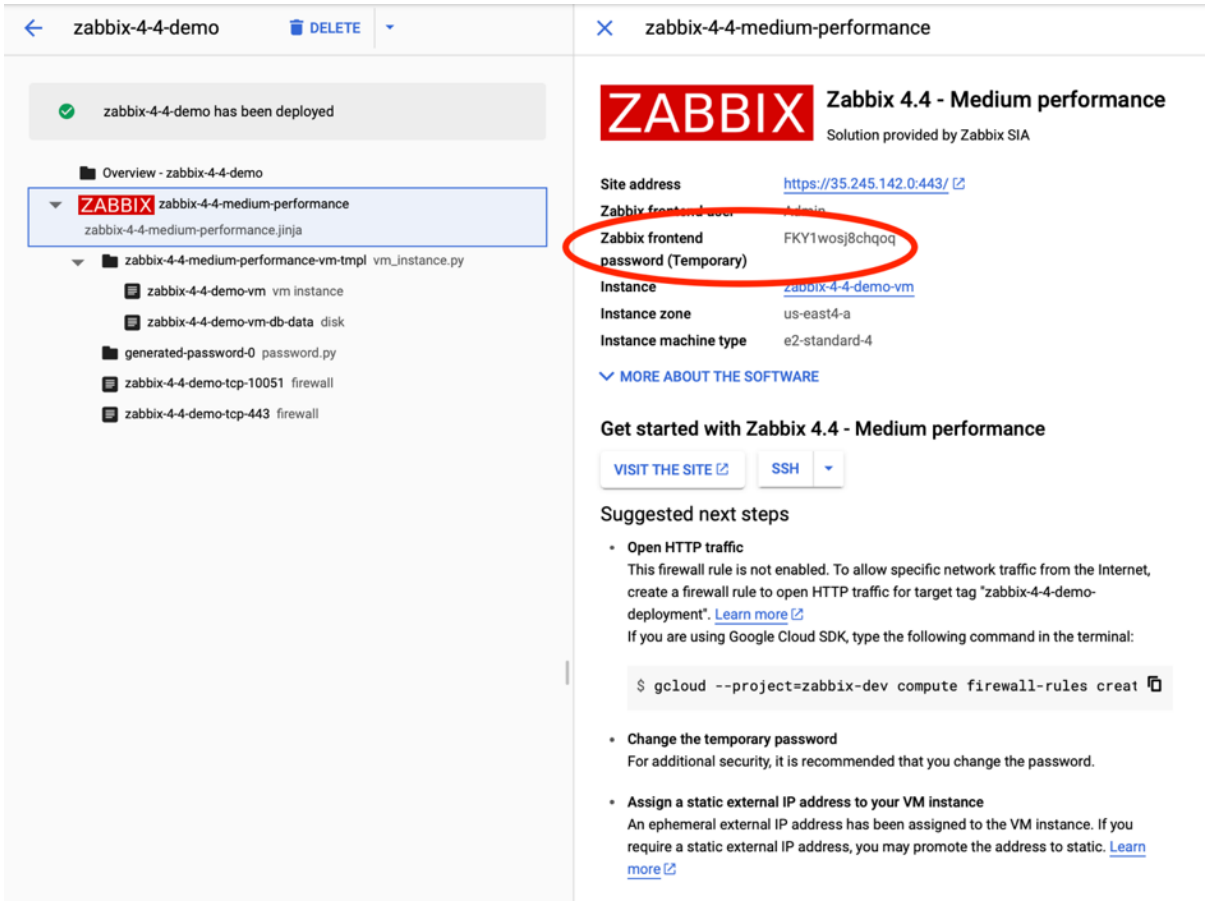
That’s it! The deployment has now started.

Accessing the appliance

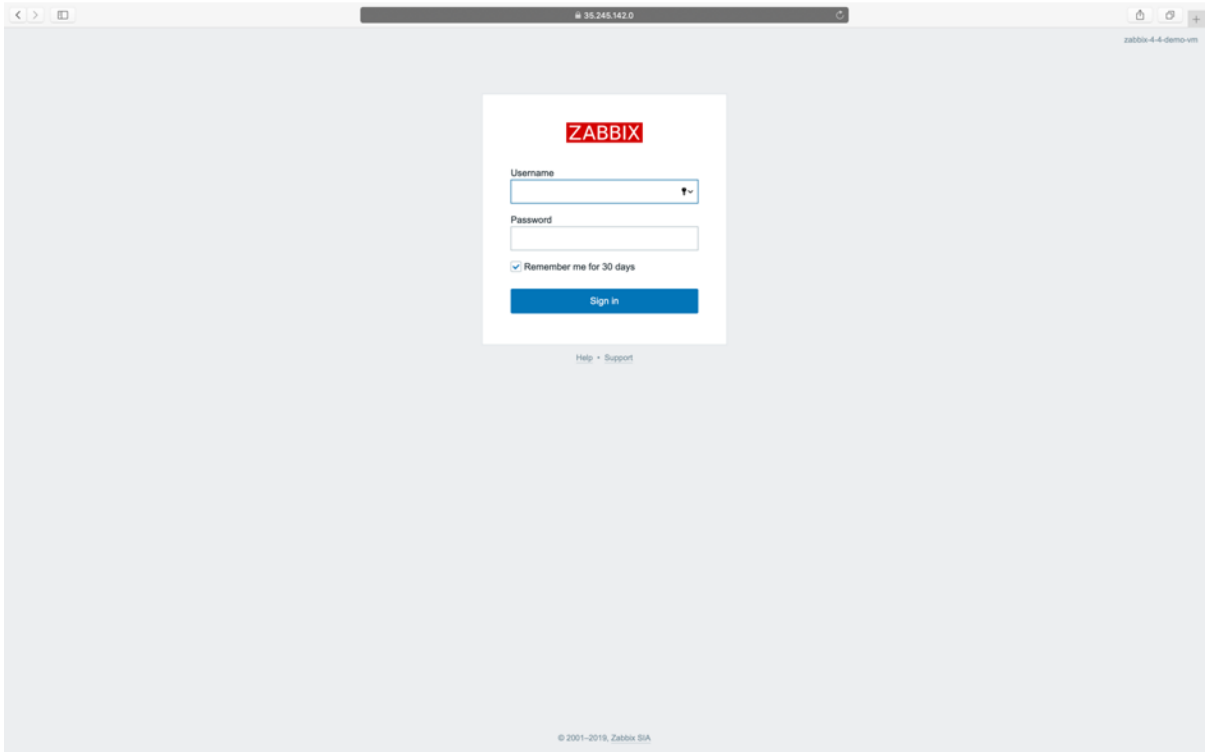
When deployment is complete, you should see a message indicating that the appliance is being deployed.

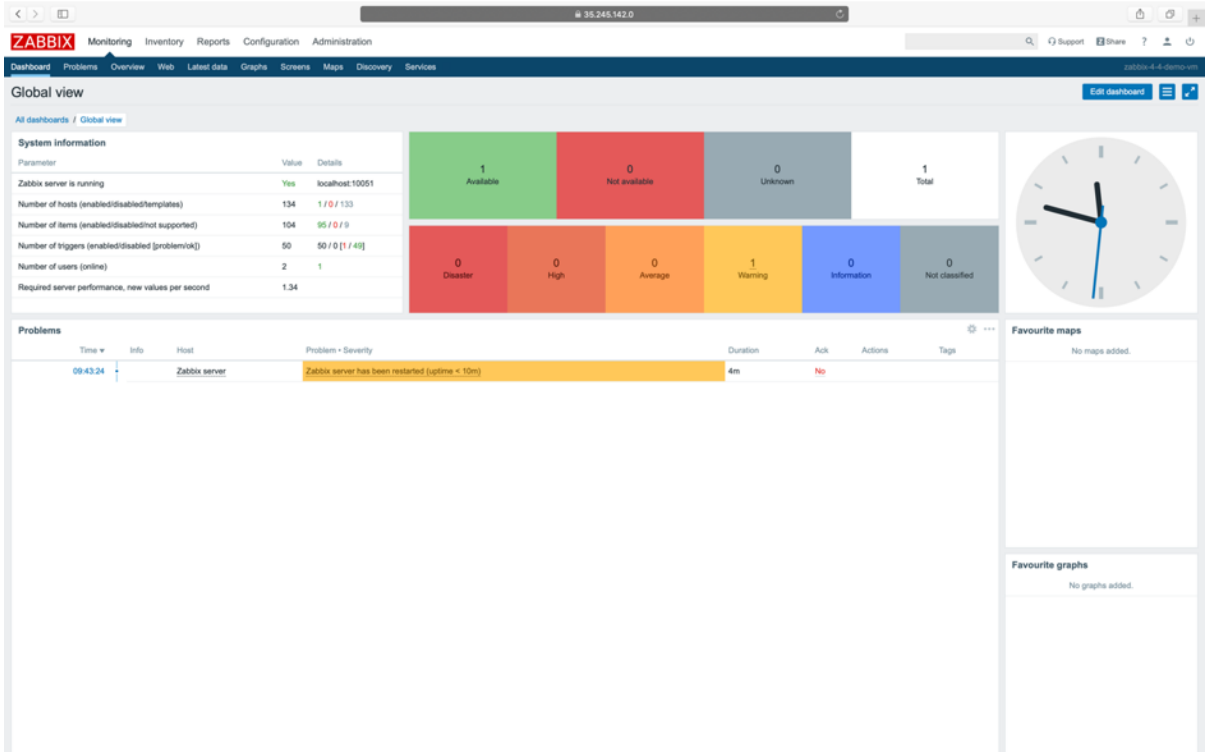


Open a web browser to *Site Address* to view Zabbix frontend. Before that, grab Zabbix *Admin* user's password to log in to your Zabbix appliance.

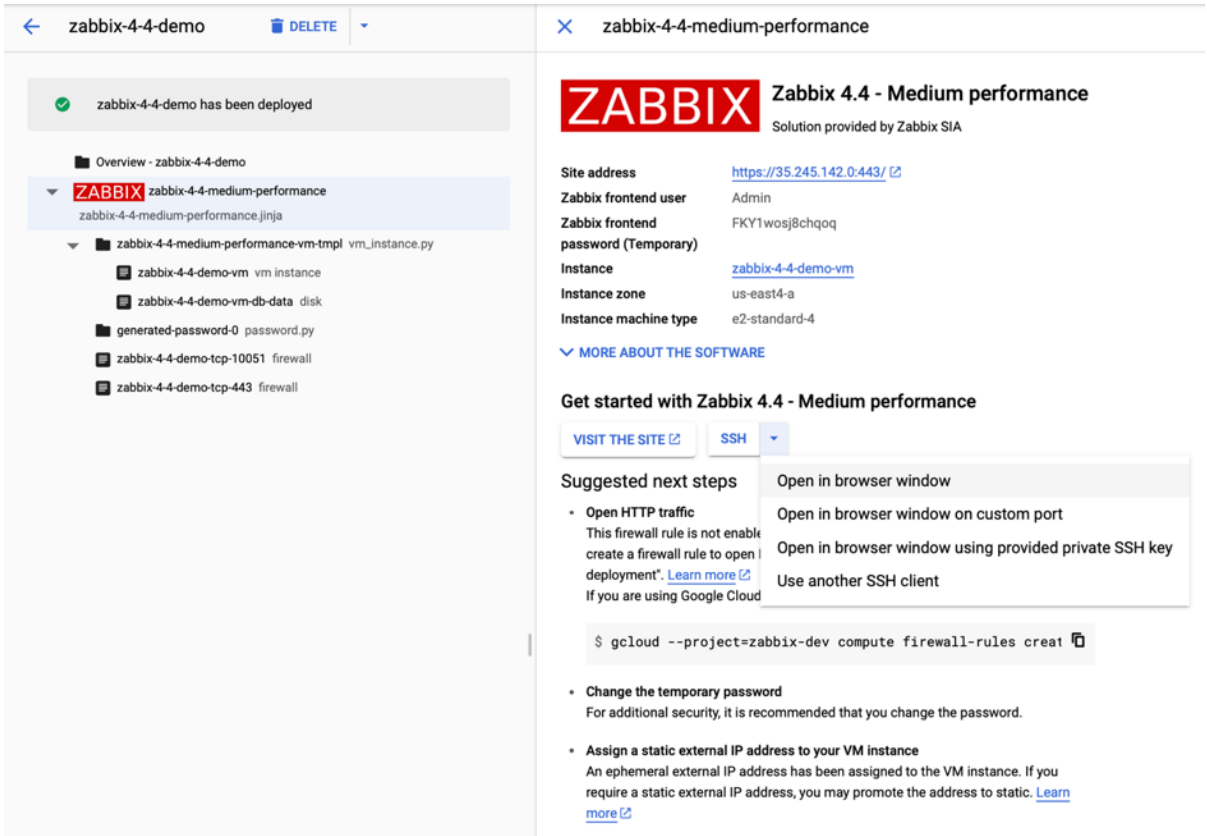


Great! You now have Zabbix appliance.





You can also log in to the instance to interact with the database. To do that go back to the Google console and follow the red arrow as shown below to start an ssh session using the *Open in browser window* option.



MySQL is available without a `root` password. The password is stored in the `/root/.my.cnf` configuration file. By default, Zabbix appliance uses the UTC timezone. If you need to change it, edit `date.timezone` PHP configuration directive in `/etc/php.d/99-zabbix.ini` and restart the php-fpm process: `systemctl restart php-fpm`.

Upgrading

Before upgrading it is highly recommended to consult all the relevant upgrade notes.

It will also be useful to know which components you must upgrade.

Overview

Step-by-step upgrade instructions are available for:

- Red Hat Enterprise Linux (using packages)
- Debian/Ubuntu (using packages)
- Upgrade using containers
- Upgrade using sources

Related instructions:

- For servers in a high-availability (HA) cluster, see Upgrading HA cluster
- For TimescaleDB database, see Upgrading TimescaleDB schema

Components to upgrade

Upgrading Zabbix agents is recommended but not mandatory.

Upgrading Zabbix proxies is highly recommended. Zabbix server fully supports proxies that are of the same major version as the server. Zabbix server also supports proxies that are **no older** than Zabbix server previous LTS release version, but with limited functionality (data collection, execution of remote commands, immediate item value checks). Configuration update is also disabled, and outdated proxies will only work with old configuration.

Attention:
Proxies that are older than Zabbix server previous LTS release version or newer than Zabbix server major version are not supported. Zabbix server will ignore data from unsupported proxies and all communication with Zabbix server will fail with a warning. For more information, see Version compatibility.

To minimize downtime and data loss during the upgrade, it is recommended to stop, upgrade, and start Zabbix server and then stop, upgrade, and start Zabbix proxies one after another. During server downtime, running proxies will continue data collection. Once the server is up and running, outdated proxies will send the data to the newer server (proxy configuration will not be updated though) and will remain partly functional. Any notifications for problems during Zabbix server downtime will be generated only after the upgraded server is started.

If Zabbix proxy is started for the first time and the SQLite database file is missing, proxy creates it automatically.

Note that if Zabbix proxy uses SQLite3 and on startup detects that existing database file version is older than required, it will **delete the database file automatically** and create a new one. Therefore, history data stored in the SQLite database file will be lost. If Zabbix proxy's version is older than the database file version, Zabbix will log an error and exit.

Depending on the database size, the database upgrade to version 8.2 may take a long time.

Upgrade notes

Direct upgrade to Zabbix 8.2.x is supported from versions **2.0**.x. For upgrading from earlier versions consult Zabbix documentation for 2.0 and earlier.

Note:
Please be aware that after upgrading some third-party software integrations in Zabbix might be affected, if the external software is not compatible with the upgraded Zabbix version.

The following upgrade notes are available:

Upgrade from	Read full upgrade notes	Most important changes between versions
8.0.x	For: Zabbix 8.2	
7.4.x	For: Zabbix 8.0 Zabbix 8.2	Pre-1.6 version Zabbix agents are no longer supported. Minimum required database versions raised. Minimum required PHP version raised from 8.0.0 to 8.2.0. The % character has been added to the `UnsafeUserParameters` list for Zabbix agent and Zabbix agent 2.
7.2.x	For: Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	Support of PCRE library dropped.

Upgrade from	Read full upgrade notes	Most important changes between versions
7.0.x	For: Zabbix 7.2 Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	Support of Oracle DB dropped.
6.4.x	For: Zabbix 7.0 Zabbix 7.2 Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	Minimum required PHP version raised from 7.4.0 to 8.0.0. Asynchronous pollers for agent, HTTP agent, SNMP walk[OID], and SNMP get[OID] checks. Separate database table for proxies. Default location for Windows agent configuration file changed. Oracle DB deprecated. Old numeric (float) value type dropped.
6.2.x	For: Zabbix 6.4 Zabbix 7.0 Zabbix 7.2 Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	Minimum required MySQL version raised from 8.0.0 to 8.0.30. 'libevent_pthreads' library is required for Zabbix server/proxy. Upon the first launch after an upgrade, Zabbix proxy with SQLite3 automatically drops the old version of the database (with all the history) and creates a new one.
6.0.x LTS	For: Zabbix 6.2 Zabbix 6.4 Zabbix 7.0 Zabbix 7.2 Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	Minimum required PHP version raised from 7.2.5 to 7.4.0. Service monitoring reworked significantly. Deterministic triggers need to be created during the upgrade. If binary logging is enabled for MySQL/MariaDB, this requires superuser privileges or setting the variable/configuration parameter *log_bin_trust_function_creators = 1*. See Database creation scripts for instructions how to set the variable.
5.4.x	For: Zabbix 6.0 Zabbix 6.2 Zabbix 6.4 Zabbix 7.0 Zabbix 7.2 Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	Minimum required database versions raised. Server/proxy will not start if outdated database. Audit log records lost because of database structure change.
5.2.x	For: Zabbix 5.4 Zabbix 6.0 Zabbix 6.2 Zabbix 6.4 Zabbix 7.0 Zabbix 7.2 Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	Minimum required database versions raised. Aggregate items removed as a separate type.
5.0.x LTS	For: Zabbix 5.2 Zabbix 5.4 Zabbix 6.0 Zabbix 6.2 Zabbix 6.4 Zabbix 7.0 Zabbix 7.2 Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	Minimum required PHP version raised from 7.2.0 to 7.2.5. Password hashing algorithm changed from MD5 to bcrypt.

Upgrade from	Read full upgrade notes	Most important changes between versions
4.4.x	For:	Support of IBM DB2 dropped.
	Zabbix 5.0	Minimum required PHP version raised from 5.4.0 to 7.2.0.
	Zabbix 5.2	Minimum required database versions raised.
	Zabbix 5.4	Changed Zabbix PHP file directory.
	Zabbix 6.0	
	Zabbix 6.2	
	Zabbix 6.4	
	Zabbix 7.0	
	Zabbix 7.2	
	Zabbix 7.4	
	Zabbix 8.0	
	Zabbix 8.2	
4.2.x	For:	Jabber, Ez Texting media types removed.
	Zabbix 4.4	
	Zabbix 5.0	
	Zabbix 5.2	
	Zabbix 5.4	
	Zabbix 6.0	
	Zabbix 6.2	
	Zabbix 6.4	
	Zabbix 7.0	
	Zabbix 7.2	
	Zabbix 7.4	
	Zabbix 8.0	
	Zabbix 8.2	
4.0.x LTS	For:	Older proxies no longer can report data to an upgraded server.
	Zabbix 4.2	Newer agents no longer will be able to work with an older Zabbix
	Zabbix 4.4	server.
	Zabbix 5.0	
	Zabbix 5.2	
	Zabbix 5.4	
	Zabbix 6.0	
	Zabbix 6.2	
	Zabbix 6.4	
	Zabbix 7.0	
	Zabbix 7.2	
	Zabbix 7.4	
	Zabbix 8.0	
	Zabbix 8.2	
3.4.x	For:	'libpthread' and 'zlib' libraries now mandatory.
	Zabbix 4.0	Support for plain text protocol dropped and header is mandatory.
	Zabbix 4.2	Pre-1.4 version Zabbix agents are no longer supported.
	Zabbix 4.4	The Server parameter in passive proxy configuration now mandatory.
	Zabbix 5.0	
	Zabbix 5.2	
	Zabbix 5.4	
	Zabbix 6.0	
	Zabbix 6.2	
	Zabbix 6.4	
	Zabbix 7.0	
	Zabbix 7.2	
	Zabbix 7.4	
	Zabbix 8.0	
	Zabbix 8.2	

Upgrade from	Read full upgrade notes	Most important changes between versions
3.2.x	For: Zabbix 3.4 Zabbix 4.0 Zabbix 4.2 Zabbix 4.4 Zabbix 5.0 Zabbix 5.2 Zabbix 5.4 Zabbix 6.0 Zabbix 6.2 Zabbix 6.4 Zabbix 7.0 Zabbix 7.2 Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	SQLite support as backend database dropped for Zabbix server/frontend. Perl Compatible Regular Expressions (PCRE) supported instead of POSIX extended. 'libpcre' and 'libevent' libraries mandatory for Zabbix server. Exit code checks added for user parameters, remote commands and system.run[] items without the 'nowait' flag as well as Zabbix server executed scripts. Zabbix Java gateway has to be upgraded to support new functionality.
3.0.x LTS	For: Zabbix 3.2 Zabbix 3.4 Zabbix 4.0 Zabbix 4.2 Zabbix 4.4 Zabbix 5.0 Zabbix 5.2 Zabbix 5.4 Zabbix 6.0 Zabbix 6.2 Zabbix 6.4 Zabbix 7.0 Zabbix 7.2 Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	Database upgrade may be slow, depending on the history table size.
2.4.x	For: Zabbix 3.0 Zabbix 3.2 Zabbix 3.4 Zabbix 4.0 Zabbix 4.2 Zabbix 4.4 Zabbix 5.0 Zabbix 5.2 Zabbix 5.4 Zabbix 6.0 Zabbix 6.2 Zabbix 6.4 Zabbix 7.0 Zabbix 7.2 Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	Minimum required PHP version raised from 5.3.0 to 5.4.0. LogFile agent parameter must be specified.

Upgrade from	Read full upgrade notes	Most important changes between versions
2.2.x LTS	For: Zabbix 2.4 Zabbix 3.0 Zabbix 3.2 Zabbix 3.4 Zabbix 4.0 Zabbix 4.2 Zabbix 4.4 Zabbix 5.0 Zabbix 5.2 Zabbix 5.4 Zabbix 6.0 Zabbix 6.2 Zabbix 6.4 Zabbix 7.0 Zabbix 7.2 Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	Node-based distributed monitoring removed.
2.0.x	For: Zabbix 2.2 Zabbix 2.4 Zabbix 3.0 Zabbix 3.2 Zabbix 3.4 Zabbix 4.0 Zabbix 4.2 Zabbix 4.4 Zabbix 5.0 Zabbix 5.2 Zabbix 5.4 Zabbix 6.0 Zabbix 6.2 Zabbix 6.4 Zabbix 7.0 Zabbix 7.2 Zabbix 7.4 Zabbix 8.0 Zabbix 8.2	Minimum required PHP version raised from 5.1.6 to 5.3.0. Case-sensitive MySQL database required for proper server work; character set utf8 and utf8_bin collation is required for Zabbix server to work properly with MySQL database. See database creation scripts. 'mysqli' PHP extension required instead of 'mysql'.

Upgrade from packages

Overview

This section provides the steps required for a successful upgrade using official RPM and DEB packages provided by Zabbix for:

- Red Hat Enterprise Linux
- Debian/Ubuntu

Zabbix packages from OS repositories

Some OS distributions (in particular, Debian-based distributions) provide their own Zabbix packages. These packages **are not supported by Zabbix** and may be outdated or missing the latest features and bug fixes. It is recommended to use only official packages from the Zabbix Official Repository.

If you are upgrading from packages provided by OS distributions (or had them installed at some point), follow this procedure to switch to official Zabbix packages:

1. Uninstall the old packages.
2. Check for and remove any leftover files from the old installation.
3. Install the official Zabbix packages following installation instructions provided by Zabbix.

Do not perform a direct upgrade, as this may result in a broken installation.

Debian/Ubuntu

Overview

This section provides instructions on upgrading from Zabbix **8.0.x** to the latest version of Zabbix **8.2.x** using official Zabbix packages for Debian/Ubuntu.

> **Warning:**
> Before upgrading, please review the relevant upgrade notes and ensure that your system meets the requirements for Zabbix 8.2.

> **Note:**
> Consider running two parallel SSH sessions during the upgrade: one for executing the upgrade steps and another for monitoring server/proxy logs. For example, run `tail -f zabbix_server.log` or `tail -f zabbix_proxy.log` in the second session to view the latest log entries and possible errors in real time. This can be critical for production instances.

For instructions on upgrading between Zabbix 8.2.x minor versions (for example, from 8.2.1 to 8.2.3), see Upgrade between minor versions.

Upgrade procedure

1 Stop Zabbix processes

Stop Zabbix server to make sure that no new data is inserted into database:

```
systemctl stop zabbix-server
```

If upgrading Zabbix proxy, agent, or agent 2, stop these components too:

```
systemctl stop zabbix-proxy
systemctl stop zabbix-agent
systemctl stop zabbix-agent2
```

2 Back up Zabbix database

Back up your existing Zabbix database to safeguard against upgrade failures (for example, disk space issues, power loss, or unexpected problems).

3 Back up Zabbix configuration files, PHP files, and Zabbix binaries

Back up existing Zabbix configuration files, PHP files, and Zabbix binaries.

For configuration files, run:

```
mkdir /opt/zabbix-backup/
cp /etc/zabbix/zabbix_server.conf /opt/zabbix-backup/
##### For Apache:
cp /etc/apache2/conf-enabled/zabbix.conf /opt/zabbix-backup/
##### For Nginx:
cp /etc/nginx/conf.d/zabbix.conf /opt/zabbix-backup/
cp /etc/zabbix/nginx.conf /opt/zabbix-backup/
```

For PHP files and Zabbix binaries, run:

```
cp -R /usr/share/zabbix/ /opt/zabbix-backup/
cp -R /usr/share/zabbix-* /opt/zabbix-backup/
```

4 Update repository configuration package

Before proceeding with the upgrade, uninstall your current Zabbix repository package:

```
rm -Rf /etc/apt/sources.list.d/zabbix.list
```

You may also need to manually remove any old Zabbix packages from your working directory (e.g., `rm zabbix-release_latest+debian12` before downloading the new one to prevent the package manager from reusing an outdated version during the upgrade process.

Then, install the latest repository configuration package to ensure compatibility with the newest packages and to include any recent security patches or bug fixes.

On **Debian 13**, run:

```
wget https://repo.zabbix.com/zabbix/8.2/release/debian/pool/main/z/zabbix-release/zabbix-release_latest+de
dpkg -i zabbix-release_latest+debian13_all.deb
```

On **Debian 12**, run:

```
wget https://repo.zabbix.com/zabbix/8.2/release/debian/pool/main/z/zabbix-release/zabbix-release_latest+de
dpkg -i zabbix-release_latest+debian12_all.deb
```

Note:
For older Debian versions, replace the link above with the correct one from the Zabbix repository. Note, however, that packages for these versions may not include all Zabbix components, and to upgrade those components from packages, consider upgrading your OS. For a list of components included, see Zabbix packages.

On **Ubuntu 26.04**, run:

```
wget https://repo.zabbix.com/zabbix/8.2/release/ubuntu/pool/main/z/zabbix-release/zabbix-release_latest+ub
dpkg -i zabbix-release_latest+ubuntu26.04_all.deb
```

On **Ubuntu 24.04**, run:

```
wget https://repo.zabbix.com/zabbix/8.2/release/ubuntu/pool/main/z/zabbix-release/zabbix-release_latest+ub
dpkg -i zabbix-release_latest+ubuntu24.04_all.deb
```

On **Ubuntu 22.04**, run:

```
wget https://repo.zabbix.com/zabbix/8.2/release/ubuntu/pool/main/z/zabbix-release/zabbix-release_latest+ub
dpkg -i zabbix-release_latest+ubuntu22.04_all.deb
```

Note:
For older Ubuntu versions, replace the link above with the correct one from the Zabbix repository. Note, however, that packages for these versions may not include all Zabbix components, and to upgrade those components from packages, consider upgrading your OS. For a list of components included, see Zabbix packages.

You may see a prompt about the Zabbix repository configuration:

```
Configuration file '/etc/apt/sources.list.d/zabbix.list'
==> Deleted (by you or by a script) since installation.
==> Package distributor has shipped an updated version.
What would you like to do about it ?  Your options are:
Y or I  : install the package maintainer's version
N or O  : keep your currently-installed version
D       : show the differences between the versions
Z       : start a shell to examine the situation
The default action is to keep your current version.
*** zabbix.list (Y/I/N/O/D/Z) [default=N] ?
```

Enter `Y` (or `I`) to install the package maintainer's version of the Zabbix repository configuration.

Then, update the repository information:

```
apt update
```

5 Upgrade Zabbix components

To upgrade Zabbix components, run:

```
apt install --only-upgrade zabbix-server-mysql zabbix-frontend-php zabbix-agent
```

- If using PostgreSQL, replace `mysql` with `pgsql` in the command.
- If upgrading proxy, replace `server` with `proxy` in the command.
- If upgrading Zabbix agent 2, replace `zabbix-agent` with `zabbix-agent2 zabbix-agent2-plugin-*` in the command.

Attention:
Upgrading Zabbix agent 2 with the `apt install zabbix-agent2` command could lead to an error. For more information, see *Known issues*.

You may see a prompt about the Zabbix server (or proxy) configuration:

```
Configuration file '/etc/zabbix/zabbix_server.conf'
==> Modified (by you or by a script) since installation.
==> Package distributor has shipped an updated version.
What would you like to do about it ?  Your options are:
Y or I  : install the package maintainer's version
N or O  : keep your currently-installed version
D       : show the differences between the versions
Z       : start a shell to examine the situation
The default action is to keep your current version.
*** zabbix_server.conf (Y/I/N/O/D/Z) [default=N] ?
```

Enter the option that best fits your situation. For example, enter `D` to compare the current and new configuration, then decide if you want to install the package maintainer's version (`Y` or `I`).

Then, upgrade the Zabbix frontend and restart the web server.

For Apache, run:

```
apt install zabbix-apache-conf
systemctl restart apache2
```

For Nginx, run:

```
apt install zabbix-nginx-conf
systemctl restart nginx
```

For Nginx, you may see a prompt about the `nginx.conf` configuration file. Enter the option that best fits your situation:

```
Configuration file '/etc/zabbix/nginx.conf'
 ==> Modified (by you or by a script) since installation.
 ==> Package distributor has shipped an updated version.
   What would you like to do about it ?  Your options are:
    Y or I  : install the package maintainer's version
    N or O  : keep your currently-installed version
      D     : show the differences between the versions
      Z     : start a shell to examine the situation
 The default action is to keep your current version.
*** nginx.conf (Y/I/N/O/D/Z) [default=N] ?
```

If, after the upgrade, you need to install an additional Zabbix package from the same repository, `apt` may select upgrades for already installed Zabbix packages during dependency resolution. To avoid upgrading an installed Zabbix server or proxy package, temporarily hold that package before installation and unhold it afterwards:

```
apt-mark hold zabbix-server-mysql
apt install zabbix-sender
apt-mark unhold zabbix-server-mysql
```

6 Review component configuration parameters

Review the relevant upgrade notes to check if any changes in the configuration parameters are necessary.

For new optional parameters, see the What's new page.

If using Nginx as the web server and the `/etc/zabbix/nginx.conf` configuration file was reinstalled during the upgrade, review the file and update the `listen` and `server_name` directives to match your Zabbix frontend configuration.

7 Start Zabbix processes

Start the upgraded Zabbix components:

```
systemctl start zabbix-server
systemctl start zabbix-proxy
systemctl start zabbix-agent
systemctl start zabbix-agent2
```

8 Clear web browser cookies and cache

After the upgrade, you may need to clear web browser cookies and web browser cache for the Zabbix web interface to work properly.

Upgrade between minor versions

It is possible to upgrade between Zabbix 8.2.x minor versions (for example, from 8.2.1 to 8.2.3).

First, update the repository information:

```
apt update
```

Then, to upgrade all Zabbix components, run:

```
apt install --only-upgrade 'zabbix*'
```

- To upgrade Zabbix server only, replace `'zabbix*'` with `'zabbix-server*'` in the command.
- To upgrade Zabbix proxy only, replace `'zabbix*'` with `'zabbix-proxy*'` in the command.
- To upgrade Zabbix agent only, replace `'zabbix*'` with `'zabbix-agent*'` in the command.
- To upgrade Zabbix agent 2 only, replace `'zabbix*'` with `'zabbix-agent2*'` in the command.

Red Hat Enterprise Linux

Overview

This section provides instructions on upgrading from Zabbix **8.0.x** to the latest version of Zabbix **8.2.x** using official Zabbix packages for Red Hat Enterprise Linux or its derivatives - AlmaLinux, CentOS Stream, Oracle Linux, and Rocky Linux.

Warning:
Before upgrading, please review the relevant upgrade notes and ensure that your system meets the requirements for Zabbix 8.2.

Note:
Consider running two parallel SSH sessions during the upgrade: one for executing the upgrade steps and another for monitoring server/proxy logs. For example, run `tail -f zabbix_server.log` or `tail -f zabbix_proxy.log` in the second session to view the latest log entries and possible errors in real time. This can be critical for production instances.

For instructions on upgrading between Zabbix 8.2.x minor versions (for example, from 8.2.1 to 8.2.3), see Upgrade between minor versions.

Upgrade procedure

1 Stop Zabbix processes

Stop Zabbix server to make sure that no new data is inserted into database:

```
systemctl stop zabbix-server
```

If upgrading Zabbix proxy, agent, or agent 2, stop these components too:

```
systemctl stop zabbix-proxy
systemctl stop zabbix-agent
systemctl stop zabbix-agent2
```

2 Back up Zabbix database

Back up your existing Zabbix database to safeguard against upgrade failures (for example, disk space issues, power loss, or unexpected problems).

3 Back up Zabbix configuration files, PHP files, and Zabbix binaries

Back up existing Zabbix configuration files, PHP files, and Zabbix binaries.

For configuration files, run:

```
mkdir /opt/zabbix-backup/
cp /etc/zabbix/zabbix_server.conf /opt/zabbix-backup/
cp /etc/httpd/conf.d/zabbix.conf  /opt/zabbix-backup/
```

For PHP files and Zabbix binaries, run:

```
cp -R /usr/share/zabbix/ /opt/zabbix-backup/
cp -R /usr/share/zabbix-* /opt/zabbix-backup/
```

4 Update repository configuration package

Before proceeding with the upgrade, update your current repository package to the latest version to ensure compatibility with the newest packages and to include any recent security patches or bug fixes.

On **RHEL 10**, run:

```
rpm -Uvh https://repo.zabbix.com/zabbix/8.2/release/rhel/10/noarch/zabbix-release-latest.el10.noarch.rpm
```

If you are upgrading Zabbix proxy on RHEL 10 and you want to use APM, also install the EPEL repository and enable the CRB repository, both of which provide dependencies for the Zabbix proxy APM package:

```
dnf -y install epel-release || dnf -y install https://dl.fedoraproject.org/pub/epel/epel-release-latest-$(
crb enable
```

On **RHEL 9**, run:

```
rpm -Uvh https://repo.zabbix.com/zabbix/8.2/release/rhel/9/noarch/zabbix-release-latest.el9.noarch.rpm
```

> **Note:**
> For older RHEL versions or its derivatives, replace the link above with the correct one from the Zabbix repository. Note, however, that packages for these versions may not include all Zabbix components, and to upgrade those components from packages, consider upgrading your OS. For a list of components included, see Zabbix packages.

Then, clean up the `dnf` package manager's cache (including headers, metadata, and package files downloaded during previous installations or updates):

```
dnf clean all
```

On the next `dnf` operation, `dnf` will download fresh metadata from the repositories since the old metadata is cleared.

See also: Known issues for updating the repository configuration package on RHEL.

5 Upgrade Zabbix components

To upgrade Zabbix components, run:

```
dnf install zabbix-server-mysql zabbix-web-mysql zabbix-selinux-policy zabbix-agent
```

- If using PostgreSQL, replace `mysql` with `pgsql` in the command.
- If upgrading the proxy, replace `server` with `proxy` in the command.
- If you have RHEL 10 and want to use APM, add `-apm` to the proxy package name (e.g., `zabbix-proxy-mysql-apm`).
- If upgrading the agent 2, replace `zabbix-agent` with `zabbix-agent2 zabbix-agent2-plugin-*` in the command.
- If you have other Zabbix packages installed, consider upgrading them as well by including them in the command.

> **Attention:**
> Upgrading Zabbix agent 2 with the `dnf install zabbix-agent2` command could lead to an error. For more information, see *Known issues*.

Then, to upgrade Zabbix frontend with Apache and restart Apache, run:

```
dnf install zabbix-apache-conf
systemctl restart httpd
```

6 Review component configuration parameters

Review the relevant upgrade notes to check if any changes in the configuration parameters are necessary.

For new optional parameters, see the What's new page.

7 Start Zabbix processes

Start the upgraded Zabbix components:

```
systemctl start zabbix-server
systemctl start zabbix-proxy
systemctl start zabbix-agent
systemctl start zabbix-agent2
```

8 Clear web browser cookies and cache

After the upgrade, you may need to clear web browser cookies and web browser cache for the Zabbix web interface to work properly.

Upgrade between minor versions

It is possible to upgrade between Zabbix 8.2.x minor versions (for example, from 8.2.1 to 8.2.3).

To upgrade all Zabbix components, run:

```
dnf upgrade 'zabbix-*'
```

- To upgrade Zabbix server only, replace `'zabbix-*'` with `'zabbix-server-*'` in the command.
- To upgrade Zabbix proxy only, replace `'zabbix-*'` with `'zabbix-proxy-*'` in the command.
- To upgrade Zabbix agent only, replace `'zabbix-*'` with `'zabbix-agent-*'` in the command.
- To upgrade Zabbix agent 2 only, replace `'zabbix-*'` with `'zabbix-agent2-*'` in the command.

Upgrade from containers

Overview

This section describes the steps required for a successful upgrade to Zabbix **8.2**.x containers.

Separate sets of instructions are available for:

- Zabbix component image upgrade
- Docker Compose file upgrade

Warning:
Before the upgrade make sure to read the relevant **upgrade notes**!

Attention:
Before starting the upgrade, verify that users have the necessary permissions to the database for upgrading purposes.

 For upgrades from Zabbix 6.0 or older, deterministic triggers will need to be created during the upgrade. If binary logging is enabled for MySQL/MariaDB, this requires superuser privileges or setting the variable/configuration parameter `log_bin_trust_function_creators = 1`. See Database creation scripts for instructions how to set the variable.

 Note that if executing from a console, the variable will only be set temporarily and will be dropped when a Docker is restarted. In this case, keep your SQL service running, only stop `zabbix-server` service by running `docker compose down zabbix-server` and then `docker compose up -d zabbix-server`.

 Alternatively, you can set this variable in the configuration file.

Depending on the database size, the database upgrade to version 8.2 may take a long time.

Zabbix component image upgrade

These instructions upgrade your Docker images used as a base for your Zabbix component containers.

The steps below use a Zabbix proxy as an example. Replace `zabbix-proxy-sqlite3` with your component's container name.

1. Check all containers to identify their names and current image versions:

```
docker ps -a

#### CONTAINER ID   IMAGE                                          (...) NAMES
#### 5a0c2e3d02b4   zabbix/zabbix-proxy-sqlite3:alpine-8.2-latest  (...) zabbix-proxy-sqlite3
```

2. Pull the new image version:

```
docker pull zabbix/zabbix-proxy-sqlite3:alpine-8.2-latest
```

3. Stop the Zabbix component container:

```
docker stop zabbix-proxy-sqlite3
```

4. Remove the Zabbix component container:

Attention:
Make sure you have noted all configuration flags and volume mounts used in the original `docker run` command before proceeding, as they will be needed to recreate the container.

```
docker rm zabbix-proxy-sqlite3
```

5. Start a new container using the updated image:

```
docker run --name zabbix-proxy-sqlite3 \
  -e ZBX_SERVER_HOST=192.0.2.1 \
  -e ZBX_PROXYMODE=0 \
  -e ZBX_HOSTNAME=zabbix-proxy-sqlite3 \
  -v zabbix-proxy-data:/var/lib/zabbix/db_data \
  --init \
  -d \
  zabbix/zabbix-proxy-sqlite3:alpine-8.2-latest
```

Note:
For more information on running Zabbix containers, see Docker (manual deployment).

6. Verify that the upgrade was successful by checking the container logs:

```
docker logs -f zabbix-proxy-sqlite3
```

Docker Compose file upgrade

These instructions upgrade the Docker Compose files used for deploying Zabbix components. Follow these instructions if you installed Zabbix using Docker Compose.

1. Check all containers to identify their names and current image versions:

```
docker ps -a

#### CONTAINER ID   IMAGE                                              (...) NAMES
#### d095f486deef   zabbix/zabbix-server-mysql:alpine-8.0-latest       (...) zabbix-docker-zabbix-server-1
#### 70627d833da0   zabbix/zabbix-web-nginx-mysql:alpine-8.0-latest    (...) zabbix-docker-zabbix-web-ngina
#### aaaa664cc185   zabbix/zabbix-server-mysql:alpine-8.0-latest       (...) zabbix-docker-server-db-init-1
#### 9ab3e77492ff   mysql:8.0-oracle                                   (...) zabbix-docker-mysql-server-1
```

2. Navigate to the Zabbix Docker repository you previously cloned, pull the latest updates, and switch to the 8.2 branch:

Attention:
If you have modified `.env`, `compose.yaml`, or other files in your local repository, a `git pull` may overwrite them. Back up any local changes before running this step.

```
cd zabbix-docker
git pull
git checkout 8.2
```

3. Deploy Zabbix components using the updated Docker Compose file:

```
#### With MySQL as the database:
docker compose -f ./compose.yaml up -d

#### With PostgreSQL as the database:
docker compose -f ./compose_pgsql.yaml up -d
```

Note:
After all containers are recreated and started, it may take 1–3 minutes before data collection resumes. For more information on running Zabbix containers, see Docker Compose.

4. Verify that the upgrade was successful by checking the container logs (replace `zabbix-docker-zabbix-server-1` with the name of your component's container):

```
docker logs -f zabbix-docker-zabbix-server-1
```

Upgrade from sources

Overview

This section provides the steps required for a successful upgrade from Zabbix **8.0**.x to Zabbix **8.2**.x using official Zabbix sources.

Warning:
Before the upgrade make sure to read the relevant **upgrade notes**!

You may also want to review the requirements for 8.2.

Note:
It may be handy to run two parallel SSH sessions during the upgrade, executing the upgrade steps in one and monitoring the server/proxy logs in another. For example, run `tail -f zabbix_server.log` or `tail -f zabbix_proxy.log` in the second SSH session showing you the latest log file entries and possible errors in real time. This can be critical for production instances.

Server upgrade process

1 Stop server

Stop Zabbix server to make sure that no new data is inserted into database.

2 Back up the existing Zabbix database

This is a very important step. Make sure that you have a backup of your database. It will help if the upgrade procedure fails (lack of disk space, power off, any unexpected problem).

3 Back up configuration files, PHP files and Zabbix binaries

Make a backup copy of Zabbix binaries, configuration files and the PHP file directory.

4 Install new server binaries

Use these instructions to compile Zabbix server from sources.

5 Review server configuration parameters

Make sure to review Upgrade notes to check if any changes in the configuration parameters are required.

For new optional parameters, see the What's new page.

6 Start new Zabbix binaries

Start new binaries. Check log files to see if the binaries have started successfully.

Zabbix server will automatically upgrade the database. When starting up, Zabbix server reports the current (mandatory and optional) and required database versions. If the current mandatory version is older than the required version, Zabbix server automatically executes the required database upgrade patches. The start and progress level (percentage) of the database upgrade is written to the Zabbix server log file. When the upgrade is completed, a "database upgrade fully completed" message is written to the log file. If any of the upgrade patches fail, Zabbix server will not start. Zabbix server will also not start if the current mandatory database version is newer than the required one. Zabbix server will only start if the current mandatory database version corresponds to the required mandatory version.

```
8673:20161117:104750.259 current database version (mandatory/optional): 03040000/03040000
8673:20161117:104750.259 required mandatory version: 03040000
```

Before you start the server:

- Make sure the database user has enough permissions (create table, drop table, create index, drop index)
- Make sure you have enough free disk space.

7 Install new Zabbix web interface

The minimum required PHP version is 8.2.0. Update if needed and follow installation instructions.

8 Clear web browser cookies and cache

After the upgrade you may need to clear web browser cookies and web browser cache for the Zabbix web interface to work properly.

Proxy upgrade process

1 Stop proxy

Stop Zabbix proxy.

2 Back up configuration files and Zabbix proxy binaries

Make a backup copy of the Zabbix proxy binary and configuration file.

3 Install new proxy binaries

Use these instructions to compile Zabbix proxy from sources.

4 Review proxy configuration parameters

Make sure to review Upgrade notes to check if any changes in the configuration parameters are required.

For new optional parameters, see the What's new page.

5 Start new Zabbix proxy

Start the new Zabbix proxy. Check log files to see if the proxy has started successfully.

Zabbix proxy will automatically upgrade the database. Database upgrade takes place similarly as when starting Zabbix server.

Agent upgrade process

> **Attention:**
> Upgrading agents is not mandatory. You only need to upgrade agents if it is required to access the new functionality.

The upgrade procedure described in this section may be used for upgrading both the Zabbix agent and the Zabbix agent 2.

1 Stop agent

Stop Zabbix agent.

2 Back up configuration files and Zabbix agent binaries

Make a backup copy of the Zabbix agent binary and configuration file.

3 Install new agent binaries

Use these instructions to compile Zabbix agent from sources.

Alternatively, you may download pre-compiled Zabbix agents from the Zabbix download page.

4 Review agent configuration parameters

Make sure to review Upgrade notes to check if any changes in the configuration parameters are required.

For new optional parameters, see the What's new page.

5 Start new Zabbix agent

Start the new Zabbix agent. Check log files to see if the agent has started successfully.

Upgrade between minor versions

When upgrading between minor versions of 8.2.x (for example, from 8.2.1 to 8.2.3) it is required to execute the same actions for server/proxy/agent as during the upgrade between major versions. The only difference is that when upgrading between minor versions no changes to the database are made.

Upgrade issues

Escaping-related Exceeded character limit for function parameters after upgrade

Proper escaping of backslashes has been added in history function string parameters in Zabbix 7.0.

As additional backslashes are added during upgrade from pre-7.0 Zabbix versions, this leads to longer parameters which may result in broken trigger functions if the parameter length exceeds the maximum data size of 255 characters.

To illustrate what problems can arise during database upgrade, what are their impact, how to recognize and fix them, let's take an example setup with 3 items and 2 triggers.

Setup

1. Zabbix 6.0 version.
2. `/tmp/ONE` file that contains 114 backslashes with 'a' at the end:

```
\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
```

3. A `vfs.file.contents[/tmp/ONE]` item to read /tmp/ONE.
4. Two calculated items:

Calc_228 - its parameter contains 228 backslashes in its parameter, which are escaped, so they **match** the 114 backslashes from the file /tmp/ONE:

```
find(/Zabbix server/vfs.file.contents[/tmp/ONE],1m:now,"regexp","\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
```

Calc_232 - its parameter contains 232 backslashes in its parameter, which are escaped, so they do **not match** the 114 backslashes from the file /tmp/ONE, as they need at least 116 real backslashes:

```
find(/Zabbix server/vfs.file.contents[/tmp/ONE],1m:now,"regexp","\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
```

5. Two triggers:

Trig_228 - similarly to Calc_228, its parameter contains 228 backslashes to **match** 114 real backslashes from the file /tmp/ONE:

```
find(/Zabbix server/vfs.file.contents[/tmp/ONE],1m:now,"regexp","\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
```

Trig_232 - similarly to Calc_232, its parameter contains 232 backslashes to match 116 real backslashes so it does **not match** the 114 backslashes from the file /tmp/ONE:

```
find(/Zabbix server/vfs.file.contents[/tmp/ONE],1h:now,"regexp","\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
```

6. Calc_228 returns 1 and Trig_228 **fires**.

7. Calc_232 returns 0 and Trig_232 **does not** fire.

Upgrade

1. Run the upgrade to 8.2.0, check logs for the warning:

```
2485502:20250228:115442.236 DBpatch_6050165(): cannot save in DB function parameter: resulting size 477 is
functionid:33792 function:'find'
used on host: 'Zabbix server'
  in trigger: 'TRIG_228'.
Current parameter value:
'$,1m:now,"regexp","\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\a"'
Resulting escaped value would be:
'$,1m:now,"regexp","\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
DB upgrade and Zabbix server can continue to run, but to make sure this function works correctly
MANUAL INTERVENTION IS REQUIRED !
Need to manually reduce the size of this parameter with macro as a workaround.

2485502:20250228:115442.237 DBpatch_6050165(): cannot save in DB function parameter: resulting size 485 is
functionid:33795 function:'find'
used on host: 'Zabbix server'
  in trigger: 'TRIG_232'.
Current parameter value:
'$,1m:now,"regexp","\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\a"'
Resulting escaped value would be:
'$,1m:now,"regexp","\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
DB upgrade and Zabbix server can continue to run, but to make sure this function works correctly
MANUAL INTERVENTION IS REQUIRED !
Need to manually reduce the size of this parameter with macro as a workaround.
```

2. The database upgrade **finishes** and Zabbix server continues to run.

3. Calc_228 still returns 1 while Calc_232 returns 0.

4. Triggers were upgraded as follows:

Trig_228:

```
find(/Zabbix server/vfs.file.contents[/tmp/ONE],1m:now,"regexp","\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
```

Trig_232:

```
find(/Zabbix server/vfs.file.contents[/tmp/ONE],1m:now,"regexp","\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
```

Both triggers - Trig_228 and Trig_232 **fire now**! This is unexpected, so manual intervention is required.

Manual intervention

Trigger expressions need to be updated so their backslashes are escaped, however this cannot be done, because the resulting parameters would be too long.

That is why you need to add macros:

{$228_BACKSLASHES_A}:

```
\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
```

{$232_BACKSLASHES_A}:

```
\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\\
```

and **update** trigger expressions to use them:

Trig_228:

```
find(/Zabbix server/vfs.file.contents[/tmp/ONE],1m:now,"regexp",{$228_BACKSLASHES_A})
```

Trig_232:

```
find(/Zabbix server/vfs.file.contents[/tmp/ONE],1m:now,"regexp",{$232_BACKSLASHES_A})
```

Now both triggers work the same way as before - only Trig_228 is **firing**.

There are no issues with the calculated items, but they could also be updated for consistency:

Calc_228:

```
find(/Zabbix server/vfs.file.contents[/tmp/ONE],1m:now,"regexp",{$228_BACKSLASHES_A})
```

Calc_232:

```
find(/Zabbix server/vfs.file.contents[/tmp/ONE],1m:now,"regexp",{$232_BACKSLASHES_A})
```

Known issues

See also: Compilation issues.

Upgrade

SQL mode setting for successful upgrade

The `sql_mode` setting in MySQL/MariaDB must have the "STRICT_TRANS_TABLES" mode set. If it is absent, the Zabbix database upgrade will fail (see also ZBX-19435).

Upgrade with MariaDB 10.2.1 and before

Upgrading Zabbix may fail if database tables were created with MariaDB 10.2.1 and before, because in those versions the default row format is compact. This can be fixed by changing the row format to dynamic (see also ZBX-17690).

MySQL database character set and collation

Zabbix installations using a MySQL database created with the utf8mb3 character set and utf8mb3_bin collation may encounter issues. For example, a known MySQL issue can result in `Row size too large (> 8126)` errors during database schema upgrades or other DDL operations.
 It is strongly recommended to convert databases using the utf8mb3 character set and utf8mb3_bin collation to utf8mb4 character set and utf8mb4_bin collation, as described in Repairing Zabbix database character set and collation.

Memory and file descriptor leak with libcurl 8.21.0

A known issue in libcurl 8.21.0 may cause memory and file descriptor leaks in Zabbix features using libcurl when handling TLS-secured connections.
 A workaround for this issue is to downgrade libcurl to version 8.20.0 until a release containing the fix is available.

Templates

Template compatibility in dual-stack (IPv4/IPv6) environments

In dual-stack environments (systems configured to support both IPv4 and IPv6), the hostname `localhost` typically resolves to both IPv4 and IPv6 addresses. Due to the common prioritization of IPv6 over IPv4 by many operating systems and DNS resolvers, Zabbix templates may fail to work correctly if the service being monitored is configured to listen only on IPv4.

Services that are not configured to listen on IPv6 addresses may become inaccessible, leading to monitoring failures. Users might configure access correctly for IPv4 but still face connectivity issues due to the default behavior of prioritizing IPv6.

A workaround for this is to ensure that the services (Nginx, Apache, PostgreSQL, etc.) are configured to listen on both IPv4 and IPv6 addresses, and Zabbix server/agent is allowed access via IPv6. Additionally, in Zabbix templates and configurations, use `localhost` explicitly instead of `127.0.0.1` to ensure compatibility with both IPv4 and IPv6.

For example, when monitoring PostgreSQL with the PostgreSQL by Zabbix agent 2 template, you may need to edit the `pg_hba.conf` file to allow connections for the `zbx_monitor` user. If the dual-stack environment prioritizes IPv6 (system resolves localhost to `::1`) and you configure `localhost` but only add an IPv4 entry (`127.0.0.1/32`), the connection will fail because there is no matching IPv6 entry.

The following `pg_hba.conf` file example ensures that the `zbx_monitor` user can connect to any database from the local machine using both IPv4 and IPv6 addresses with different authentication methods:

```
### TYPE     DATABASE     USER            ADDRESS          METHOD
  host     all          zbx_monitor     localhost        trust
  host     all          zbx_monitor     127.0.0.1/32     md5
  host     all          zbx_monitor     ::1/128          scram-sha-256
```

If necessary, you can also use the IPv4 address (`127.0.0.1`) directly when configuring the PostgreSQL by Zabbix agent 2 template macro for the connection string.

Accidental installation of EPEL Zabbix packages

If the EPEL repository is installed and enabled, installing Zabbix packages may pull EPEL versions instead of the official Zabbix packages. To resolve this:

1. Remove any Zabbix packages installed from EPEL:

```
dnf remove zabbix-server-mysql
```

2. Exclude Zabbix packages from EPEL by adding the following line to `/etc/yum.repos.d/epel.repo` file:

```
[epel]
...
excludepkgs=zabbix*
```

3. Reinstall the official Zabbix server package:

```
dnf install zabbix-server-mysql
```

During installation, official Zabbix packages include the word `release` in their version string (e.g., `7.0.0-release1.el8`), distinguishing them from EPEL packages.

Zabbix packages for RHEL on Red Hat UBI environments

When installing Zabbix from Red Hat Enterprise Linux packages on Red Hat Universal Base Image environments, ensure access to required repositories and dependencies. Zabbix packages depend on `libOpenIPMI.so` and `libOpenIPMIposix.so` libraries, which are not provided by any package in the default package manager repositories enabled on UBI systems and will result in installation failures.

The `libOpenIPMI.so` and `libOpenIPMIposix.so` libraries are available in the `OpenIPMI-libs` package, which is provided by the `redhat-#-for-<arch>-appstream-rpms` repository. Access to this repository is curated by subscriptions, which, in the case of UBI environments, get propagated by mounting repository configuration and secrets directories of the RHEL host into the container file-system namespace.

For more information, see ZBX-24291.

Expired signing key for RHEL packages

When upgrading Zabbix on Red Hat Enterprise Linux or its derivatives, you may encounter an expired signing key issue for packages on Zabbix repository. When a signing key expires, attempts to verify package signatures will result in an error indicating that the certificate or key is no longer valid. For example:

```
error: Verifying a signature using certificate D9AA84C2B617479C6E4FCF4D19F2475308EFA7DD (Zabbix LLC (Jul 2
  1. Certificate 19F2475308EFA7DD invalid: certificate is not alive
      because: The primary key is not live
      because: Expired on 2024-07-04T11:41:23Z
```

```
  2. Key 19F2475308EFA7DD invalid: key is not alive
      because: The primary key is not live
      because: Expired on 2024-07-04T11:41:23Z
```

To resolve such issues, manually reinstall the latest `zabbix-release` package for your specific variant of RHEL (replace the link below with the correct one from Zabbix repository).

For example, on **RHEL 10**, run:

```
rpm -Uvh https://repo.zabbix.com/zabbix/8.2/release/rhel/10/noarch/zabbix-release-latest.el10.noarch.rpm
```

Then, update the repository information:

```
dnf update
```

For more information, see ZBX-24761.

Database TLS connection with MariaDB

Database TLS connection is not supported with the `verify_ca` option for the DBTLSConnect parameter if MariaDB is used.

Possible deadlocks with MySQL/MariaDB

When running under high load, and with more than one LLD worker involved, it is possible to run into a deadlock caused by an InnoDB error related to the row-locking strategy (see upstream bug). The error has been fixed in MySQL since 8.0.29, but not in MariaDB. For more details, see ZBX-21506.

Global event correlation

Events may not get correlated correctly if the time interval between the first and second event is very small, i.e. half a second and less.

NetBSD 8.0 and newer

Various Zabbix processes may randomly crash on startup on the NetBSD versions 8.X and 9.X. That is due to the too small default stack size (4MB), which must be increased by running:

```
ulimit -s 10240
```

For more information, please see the related problem report: ZBX-18275.

Regular expression limitations in Zabbix agent 2

Zabbix agent 2 does not support lookaheads and lookbehinds in regular expressions due to the standard Go regexp library limitations.

IPMI checks

IPMI checks will not work with the standard OpenIPMI library package on Debian prior to 9 (stretch) and Ubuntu prior to 16.04 (xenial). To fix that, recompile OpenIPMI library with OpenSSL enabled as discussed in ZBX-6139.

IPMI — untrusted hosts may crash OpenIPMI

There is a bug in the OpenIPMI library used by Zabbix for polling IPMI data that can be triggered by specially crafted responses from an untrusted device. An untrusted IPMI device may send crafted data that causes the OpenIPMI library to crash, which in turn can cause Zabbix server process that performs IPMI polling to terminate.

SSH checks

- Some Linux distributions like Debian, Ubuntu do not support encrypted private keys (with passphrase) if the libssh2 library is installed from packages. Please see ZBX-4850 for more details.
- When using libssh 0.9.x on some Linux distributions with OpenSSH 8, SSH checks may occasionally report "Cannot read data from SSH server". This is caused by a libssh issue (more detailed report). The error is expected to have been fixed by a stable libssh 0.9.5 release. See also ZBX-17756 for details.
- Using the pipe "|" in the SSH script may lead to a "Cannot read data from SSH server" error. In this case it is recommended to upgrade the libssh library version. See also ZBX-21337 for details.

ODBC checks

- MySQL unixODBC driver should not be used with Zabbix server or Zabbix proxy compiled against MariaDB connector library and vice versa, if possible it is also better to avoid using the same connector as the driver due to an upstream bug. Recommended setup:

  PostgreSQL, SQLite or Oracle connector → MariaDB or MySQL unixODBC driver
MariaDB connector → MariaDB unixODBC driver
MySQL connector → MySQL unixODBC driver

See ZBX-7665 for more information and available workarounds.

- XML data queried from Microsoft SQL Server may get truncated in various ways on Linux and UNIX systems.
- It has been observed that using ODBC checks for monitoring Oracle databases using various versions of Oracle Instant Client for Linux causes Zabbix server to crash.
 See also: ZBX-18402, ZBX-20803.
- If using FreeTDS UnixODBC driver, you need to prepend a 'SET NOCOUNT ON' statement to an SQL query (for example, `SET NOCOUNT ON DECLARE @strsql NVARCHAR(max) SET @strsql = ....`). Otherwise, database monitor item in Zabbix will fail to retrieve the information with an error "SQL query returned empty result".
 See ZBX-19917 for more information.

Incorrect request method parameter in items

The request method parameter, used only in HTTP checks, may be incorrectly set to '1', a non-default value for all items as a result of upgrade from a pre-4.0 Zabbix version. For details on how to fix this situation, see ZBX-19308.

Web monitoring and HTTP agent

Zabbix server leaks memory on some Linux distributions due to an upstream bug when "SSL verify peer" is enabled in web scenarios or HTTP agent. Please see ZBX-10486 for more information and available workarounds.

Simple checks

There is a bug in **fping** versions earlier than v3.10 that mishandles duplicate echo replay packets. This may cause unexpected results for `icmpping`, `icmppingloss`, `icmppingsec` items. It is recommended to use the latest version of **fping**. Please see ZBX-11726 for more details.

Errors with fping execution in rootless containers

When containers are running in rootless mode or in a specific-restrictions environment, you may face errors related to fping execution when performing ICMP checks, such as `fping: Operation not permitted` or all packets to all resources lost.

To fix this problem add `--cap-add=net_raw` to "docker run" or "podman run" commands.

Additionally fping execution in non-root environments may require sysctl modification, i.e.:

```
sudo sysctl -w "net.ipv4.ping_group_range=0 1995"
```

where "1995" is the zabbix GID. For more details, see ZBX-22833.

SNMP checks

If the OpenBSD operating system is used, a use-after-free bug in the Net-SNMP library up to the 5.7.3 version can cause a crash of Zabbix server if the SourceIP parameter is set in the Zabbix server configuration file. As a workaround, please do not set the SourceIP parameter. The same problem applies also for Linux, but it does not cause Zabbix server to stop working. A local patch for the net-snmp package on OpenBSD was applied and will be released with OpenBSD 6.3.

SNMP data spikes

Spikes in SNMP data have been observed that may be related to certain physical factors like voltage spikes in the mains. See ZBX-14318 more details.

SNMP traps

The "net-snmp-perl" package, needed for SNMP traps, has been removed in RHEL 8.0-8.2; re-added in RHEL 8.3.

So if you are using RHEL 8.0-8.2, the best solution is to upgrade to RHEL 8.3.

Please also see ZBX-17192 for more information.

Alerter process crash in RHEL 7

Instances of a Zabbix server alerter process crash have been encountered in RHEL 7. Please see ZBX-10461 for details.

Upgrading Zabbix agent 2 (6.0.5 or older)

When upgrading Zabbix agent 2 (version 6.0.5 or older) from packages, a plugin-related file conflict error may occur. To fix the error, back up your agent 2 configuration (if necessary), uninstall agent 2 and install it anew.

On RHEL-based systems, run:

```
dnf remove zabbix-agent2
dnf install zabbix-agent2
```

On Debian-based systems, run:

```
apt remove zabbix-agent2
apt install zabbix-agent2
```

For more information, see ZBX-23250.

Flipping frontend locales

It has been observed that frontend locales may flip without apparent logic, i.e. some pages (or parts of pages) are displayed in one language while other pages (or parts of pages) in a different language. Typically the problem may appear when there are several users, some of whom use one locale, while others use another.

A known workaround to this is to disable multithreading in PHP and Apache.

The problem is related to how setting the locale works in PHP: locale information is maintained per process, not per thread. So in a multi-thread environment, when there are several projects run by same Apache process, it is possible that the locale gets changed in another thread and that changes how data can be processed in the Zabbix thread.

For more information, please see related problem reports:

- ZBX-10911 (Problem with flipping frontend locales)
- ZBX-16297 (Problem with number processing in graphs using the `bcdiv` function of BC Math functions)

Graphs

Graphs (classic) issues

If you encounter issues with classic graphs, it is recommended to update GD library (libgd) to version 2.3.3-13 or higher, and PHP to version 8.0.19, 8.1.33, 8.2.29, 8.3.25, 8.4.12, or higher.

Daylight Saving Time

Changes to Daylight Saving Time (DST) result in irregularities when displaying X axis labels (date duplication, date missing, etc).

Sum aggregation

When using sum aggregation in a graph for period that is less than one hour, graphs display incorrect (multiplied) values when data come from trends.

Text overlapping

For some frontend languages (e.g., Japanese), local fonts can cause text overlapping in graph legend. To avoid this, use version 2.3.0 (or later) of PHP GD extension.

Log file monitoring

`log[]` and `logrt[]` items repeatedly reread log file from the beginning if file system is 100% full and the log file is being appended (see ZBX-10884 for more information).

Slow MySQL queries

Zabbix server generates slow `SELECT` queries in case of non-existing values for items. This issue is known to occur in MySQL 5.6/5.7 versions (for an extended discussion, see ZBX-10652), and, in specific cases, may also occur in later MySQL versions. A workaround to this is disabling the `index_condition_pushdown` or `prefer_ordering_index` optimizer in MySQL. Note, however, that this workaround may not fix all issues related to slow queries.

Persistent filter settings from links

When opening a link to Zabbix frontend page that contains filter settings, including the time selector, the filter is automatically saved in the database for the user, replacing the previously saved filter and/or time selector settings for that page. These settings remain active until the user manually updates or resets them.

IPv6 address issue in SNMPv3 traps

Due to a net-snmp bug, IPv6 address may not be correctly displayed when using SNMPv3 in SNMP traps. For more details and a possible workaround, see ZBX-14541.

Trimmed long IPv6 IP address in failed login information

A failed login attempt message will display only the first 39 characters of a stored IP address as that's the character limit in the database field. That means that IPv6 IP addresses longer than 39 characters will be shown incompletely.

Zabbix agent checks on Windows

Non-existing DNS entries in a `Server` parameter of Zabbix agent configuration file (`zabbix_agentd.conf`) may increase Zabbix agent response time on Windows. This happens because Windows DNS caching daemon doesn't cache negative responses for

IPv4 addresses. However, for IPv6 addresses negative responses are cached, so a possible workaround to this is disabling IPv4 on the host.

YAML export/import

There are some known issues with YAML export/import:

- Error messages are not translatable;
- Valid JSON with a .yaml file extension sometimes cannot be imported;
- Unquoted human-readable dates are automatically converted to Unix timestamps.

Setup wizard on SUSE with NGINX and php-fpm

Frontend setup wizard cannot save configuration file on SUSE with NGINX + php-fpm. This is caused by a setting in /usr/lib/systemd/system/php-fpm.service unit, which prevents Zabbix from writing to /etc. (introduced in PHP 7.4).

There are two workaround options available:

- Set the ProtectSystem option to 'true' instead of 'full' in the php-fpm systemd unit.
- Manually save `/etc/zabbix/web/zabbix.conf.php` file.

Authorization header forwarding

In some cases, Apache or NGINX may prevent the Authorization header in API requests from reaching Zabbix. This can cause authentication issues when using Zabbix API or single sign-on (SSO) services, such as SAML with Okta.

To address this, update your web server's configuration.

For **Apache**, if you are using it as a reverse proxy (non-CGI setup), add the following directive to `/etc/httpd/conf/httpd.conf` (on RHEL-based systems) or `/etc/apache2/apache2.conf` (on Debian/Ubuntu):

```
SetEnvIfNoCase ^Authorization$ "(.+)" HTTP_AUTHORIZATION=$1
```

If Apache directly executes scripts to handle requests (e.g., by using mod_cgi), add the following directive instead:

```
CGIPassAuth On
```

In contrast, **NGINX** handles the Authorization header automatically. However, if NGINX is acting as a reverse proxy, you may explicitly forward the Authorization header by adding the following directives to `/etc/nginx/nginx.conf` (for your Zabbix frontend location):

```
...
location / {
...
    proxy_set_header Authorization $http_authorization;
    proxy_pass http://backend_server;
...
}
```

After updating the configuration, restart you web server.

For more details, see:

- ZBX-22952
- Apache 2.4 + PHP-FPM and Authorization headers
- SetEnvIfNoCase and CGIPassAuth directives
- NGINX Reverse Proxy

Chromium for Zabbix web service on Ubuntu 20

Though in most cases, Zabbix web service can run with Chromium, on Ubuntu 20.04 using Chromium causes the following error:

```
Cannot fetch data: chrome failed to start:cmd_run.go:994:
WARNING: cannot create user data directory: cannot create
"/var/lib/zabbix/snap/chromium/1564": mkdir /var/lib/zabbix: permission denied
Sorry, home directories outside of /home are not currently supported. See https://forum.snapcraft.io/t/112
```

This error occurs because `/var/lib/zabbix` is used as a home directory of user 'zabbix'.

MySQL custom error codes

When Zabbix detects that the backend database is inaccessible, it sends a notification and continues attempting to connect. For certain database engines, specific error codes are recognized. In MySQL, these recognized error codes include:

- CR_CONN_HOST_ERROR
- CR_SERVER_GONE_ERROR

- CR_CONNECTION_ERROR
- CR_SERVER_LOST
- CR_UNKNOWN_HOST
- ER_SERVER_SHUTDOWN
- ER_ACCESS_DENIED_ERROR
- ER_ILLEGAL_GRANT_FOR_TABLE
- ER_TABLEACCESS_DENIED_ERROR
- ER_UNKNOWN_ERROR

Additionally, when using Zabbix with a MySQL installation on Azure, the generic error message *[9002] Some errors occurred* may appear in Zabbix logs. This message is sent by the database to the Zabbix server or proxy. To determine the cause of the error, please consult the Azure logs.

Invalid regular expressions after switching to PCRE2

In Zabbix 6.0 support for PCRE2 has been added. Even though PCRE is still supported, Zabbix installation packages for RHEL 7 and newer, SLES (all versions), Debian 9 and newer, Ubuntu 16.04 and newer have been updated to use PCRE2. While providing many benefits, switching to PCRE2 may cause certain existing PCRE regexp patterns becoming invalid or behaving differently. In particular, this affects the pattern *[\w-\.]*. In order to make this regexp valid again without affecting semantics, change the expression to *[-\w\.]* . This happens due to the fact that PCRE2 treats the dash sign as a delimiter, creating a range inside a character class.

Geomap widget error

The maps in the Geomap widget may not load correctly, if you have upgraded from an older Zabbix version with NGINX and didn't switch to the new NGINX configuration file during the upgrade.

To fix the issue, you can discard the old configuration file, use the configuration file from the current version package and reconfigure it as described in the download instructions in section *e. Configure PHP for Zabbix frontend*.

Alternatively, you can manually edit an existing NGINX configuration file (typically, `/etc/zabbix/nginx.conf`). To do so, open the file and locate the following block:

```
location ~ /(api\/|conf[^\.]|include|locale|vendor) {
        deny            all;
        return          404;
}
```

Then, replace this block with:

```
location ~ /(api\/|conf[^\.]|include|locale) {
        deny            all;
        return          404;
}

location /vendor {
        deny            all;
        return          404;
}
```

Preprocessing — global variables are unsafe

Preprocessing JavaScript runs per request, but assignments to undeclared identifiers (for example, `secret = value`) create implicit globals that may persist beyond the current execution. Storing sensitive data (tokens, passwords, etc.) in implicit globals increases risk of accidental exposure or reuse by subsequent preprocessing runs or other integrations executing in the same environment.

Do not rely on implicit globals. Always declare variables with `var` or `const`, and avoid attaching secrets to global objects (for example, `globalThis` or `window`). There is no supported way to override built-in global objects from within preprocessing.

Secure example:

```
var apiToken = payload.token;
var count = 1;
return JSON.stringify({ token: apiToken, calls: count });
```

Processor groups on Windows

Microsoft documentation states that systems with fewer than 64 logical processors always have a single processor group, Group 0. However, Zabbix users have reported a rare bug ZBX-20260, when there are two processor groups on systems with 64 or less logical processors. This resulted in having the "\Processor(n)" performance counters for only one processor group out of two. The

actual root cause of this bug is not known. However, a similar case was described at stackoverflow.com, and the root cause there was in interoperation between BIOS and Windows.

Limits of filtering with utf8mb4 collations

Filters (e.g., in *Data collection* > *Maintenance*) may not function correctly when applied to entities containing certain Unicode characters (e.g., 𝓀, 𝔞). This issue arises due to how the default utf8mb4_bin collation for MySQL or MariaDB databases handles sorting and comparison of Unicode characters.

To address this limitation, users can change the collation of database columns to alternatives such as utf8mb4_0900_bin, utf8mb4_0900_ai_ci, or utf8mb4_unicode_520_ci. Note, however, that changing the collation may cause unexpected behavior in the handling of empty spaces, as well as sorting and filtering for other characters.

For more information on changing collations, see MySQL documentation or MariaDB documentation. For details on collation differences, see Unicode Character Sets in MySQL documentation.

Profiling excessive memory usage with tcmalloc

If you suspect your Zabbix installation is using too much memory, you can use tcmalloc's memory profiling feature to investigate Zabbix server/proxy memory consumption.

1. When installing Zabbix from sources, configure additional flags:

```
export CFLAGS="-std=gnu99 -g -O0"
```

The `-std=gnu99` flag is required for building Zabbix server, Zabbix proxy, or Zabbix agent. The `-g` flag adds extra debugging information, while `-O0` disables optimizations, which can interfere with tcmalloc's profiling.

2. Set the following environment variables before starting the Zabbix server. These variables tell tcmalloc how to track and report memory usage:

```
LD_PRELOAD="/usr/lib/aarch64-linux-gnu/libtcmalloc.so" \
HEAPPROFILE=./heap_profile \
HEAP_PROFILE_ALLOCATION_INTERVAL=0 \
HEAP_PROFILE_INUSE_INTERVAL=4294967296 \
HEAPPROFILESIGNAL=5 \
MALLOCSTATS=1 \
./sbin/zabbix_server -f -c /etc/zabbix/zabbix_server.conf
```

3. Trigger a profile dump by sending signal 5 to the target process. Replace 1234 with the actual process ID (PID):

```
kill -5 1234
```

4. Print the generated profile:

```
pprof-symbolize -text ./sbin/zabbix_server ./heap_profile.0001.heap

Using local file ./sbin/zabbix_server.
Using local file ./heap_profile.0001.heap.
Total: 1078.1 MB
  1076.8  99.9%  99.9%   1076.8  99.9% zbx_malloc2
     1.0   0.1% 100.0%      1.0   0.1% __GI___strdup
     0.2   0.0% 100.0%      0.2   0.0% CRYPTO_zalloc@@OPENSSL_3.0.0
     0.1   0.0% 100.0%      0.1   0.0% OPENSSL_LH_insert@@OPENSSL_3.0.0
     0.0   0.0% 100.0%      0.0   0.0% zbx_realloc2
     0.0   0.0% 100.0%      0.1   0.0% PKCS7_decrypt@@OPENSSL_3.0.0
     0.0   0.0% 100.0%      0.0   0.0% find_best_tree_node
     0.0   0.0% 100.0%      0.0   0.0% CRYPTO_strndup@@OPENSSL_3.0.0
     ...
     0.0   0.0% 100.0%      0.0   0.0% preprocessing_flush_value
     0.0   0.0% 100.0%   1074.0  99.6% preprocessor_add_request
```

In this example, zbx_malloc2 is responsible for almost all memory allocations.

See also:

- ZBX-25050 and ZBX-25584 for the related problem reports.
- GCC Option Summary on compiling options (`-std=gnu99`, `-g`, `-O0`, etc.).
- Gperftools Heap Profiler documentation on environment variables for tcmalloc profiling.

MySQL 8.0 Group Replication in multi-primary mode

When using MySQL 8.0 Group Replication in multi-primary mode, you may encounter an error during transaction commits similar to the following:

```
1531697:20250128:064734.697 query [txnlev:1] [update alerts set status=1,retries=0,error='' where alertid=
1531697:20250128:064734.713 query [txnlev:1] [commit;]
1531697:20250128:064734.753 [Z3005] query failed: [3101] Plugin instructed the server to rollback the curr
```

This error appears to be triggered by issues with rollback operations involving foreign key constraints.

See also:

- ZBX-26060 for the related problem report.
- MySQL Bug #96758 "Rollbacks with Foreign Keys on single node" for the upstream issue.

Login and systemd

We recommend creating a *zabbix* user as system user, that is, without ability to log in. Some users ignore this recommendation and use the same account to log in (e. g. using SSH) to host running Zabbix. This might crash Zabbix daemon on log out. In this case you will get something like the following in Zabbix server log:

```
zabbix_server [27730]: [file:'selfmon.c',line:375] lock failed: [22] Invalid argument
zabbix_server [27716]: [file:'dbconfig.c',line:5266] lock failed: [22] Invalid argument
zabbix_server [27706]: [file:'log.c',line:238] lock failed: [22] Invalid argument
```

and in Zabbix agent log:

```
zabbix_agentd [27796]: [file:'log.c',line:238] lock failed: [22] Invalid argument
```

This happens because of default systemd setting `RemoveIPC=yes` configured in `/etc/systemd/logind.conf`. When you log out of the system the semaphores created by Zabbix previously are removed which causes the crash.

A quote from systemd documentation:

```
RemoveIPC=

Controls whether System V and POSIX IPC objects belonging to the user shall be removed when the
user fully logs out. Takes a boolean argument. If enabled, the user may not consume IPC resources
after the last of the user's sessions terminated. This covers System V semaphores, shared memory
and message queues, as well as POSIX shared memory and message queues. Note that IPC objects of the
root user and other system users are excluded from the effect of this setting. Defaults to "yes".
```

There are 2 solutions to this problem:

1. (recommended) Stop using *zabbix* account for anything else than Zabbix processes, create a dedicated account for other things.
2. (not recommended) Set `RemoveIPC=no` in `/etc/systemd/logind.conf` and reboot the system. Note that `RemoveIPC` is a system-wide parameter, changing it will affect the whole system.

Requirements

Hardware

Memory

Zabbix requires both physical and disk memory. The amount of required disk memory obviously depends on the number of hosts and parameters that are being monitored. If you're planning to keep a long history of monitored parameters, you should be thinking of at least a couple of gigabytes to have enough space to store the history in the database. Each Zabbix daemon process requires several connections to a database server. The amount of memory allocated for the connection depends on the configuration of the database engine.

> **Note:**
> The more physical memory you have, the faster the database (and therefore Zabbix) works.

CPU

Zabbix and especially Zabbix database may require significant CPU resources depending on number of monitored parameters and chosen database engine.

Other hardware

A serial communication port and a serial GSM modem are required for using SMS notification support in Zabbix. USB-to-serial converter will also work.

Examples of hardware configuration

The table provides examples of hardware configuration, assuming a **Linux/BSD/Unix** platform.

These are size and hardware configuration examples to start with. Each Zabbix installation is unique. Make sure to benchmark the performance of your Zabbix system in a staging or development environment, so that you can fully understand your requirements before deploying the Zabbix installation to its production environment.

Installation size	Monitored metrics[1]	CPU/vCPU cores	Memory (GiB)	Database	Amazon EC2[2]
Small	1 000	2	8	MySQL Server, Percona Server, MariaDB Server, PostgreSQL	m6i.large/m6g.large
Medium	10 000	4	16	MySQL Server, Percona Server, MariaDB Server, PostgreSQL	m6i.xlarge/m6g.xlarge
Large	100 000	16	64	MySQL Server, Percona Server, MariaDB Server, PostgreSQL	m6i.4xlarge/m6g.4xlarge
Very large	1 000 000	32	96	MySQL Server, Percona Server, MariaDB Server, PostgreSQL	m6i.8xlarge/m6g.8xlarge

[1] For sizing purposes, 1 monitored metric is roughly represented by the formula 1 metric = 1 item + 1 trigger + 1 graph. This is only an estimation rule, not a formula for calculating monitored metrics.
 [2] Example with Amazon general purpose EC2 instances, using ARM64 or x86_64 architecture, a proper instance type like Compute/Memory/Storage optimised should be selected during Zabbix installation evaluation and testing before installing in its production environment.

Note:
Actual configuration depends on the number of active items and refresh rates very much (see database size section of this page for details). It is highly recommended to run the database on a separate server for large installations.

Supported platforms

Due to security requirements and the mission-critical nature of the monitoring server, UNIX is the only operating system that can consistently deliver the necessary performance, fault tolerance, and resilience. Zabbix operates on market-leading versions.

Zabbix components are available and tested for the following platforms:

Platform	Server	Agent	Agent 2	Comments
Linux	x	x	x	
Windows	-	x	x	Zabbix agent is supported on all desktop and server versions since Windows XP (64-bit)/Server 2003. Zabbix agent 2 is supported on all desktop and server versions since Windows 10 (32-bit)/Server 2016, as it is compiled only with a supported Go version to prevent critical security vulnerabilities. Since Go 1.21, the minimum required Windows versions are raised, making Windows 10/Server 2016 the minimum version for Zabbix agent 2.
macOS	x	x	-	
IBM AIX	x	x	-	Zabbix agent does not work on AIX platforms below versions 6.1 TL07 / 7.1 TL01.
FreeBSD	x	x	-	
OpenBSD	x	x	-	

Platform	Server	Agent	Agent 2	Comments
Solaris	x	x	-	
NetBSD	x	x	-	
HP-UX	x	x	-	

Note:
Zabbix server/agent may work on other Unix-like operating systems.

Attention:
Zabbix disables core dumps if compiled with encryption and does not start if the system does not allow disabling of core dumps.

Required software

Zabbix is built around modern web servers, leading database engines, and PHP scripting language.

Third-party external surrounding software

If stated as mandatory, the required software/library is strictly necessary. Optional ones are needed for supporting some specific function.

Software	Mandatory status	Supported versions	Comments
MySQL/Percona	One of	8.4.0-9.7.X	Required if MySQL (or Percona) is used as Zabbix backend database. InnoDB engine is required. We recommend using the C API (libmysqlclient) library for building server/proxy. We strongly recommend using only LTS releases in production environments. Releases labeled as ”rolling release”, ”early access”, or similar are not recommended.
MariaDB		10.11.00-12.3.X	InnoDB engine is required. We recommend using the MariaDB Connector/C library for building server/proxy. See also: Possible deadlocks with MariaDB and Access to UI elements with MariaDB 10.5.1-10.5.9. We strongly recommend using only LTS releases in production environments. Releases labeled as ”rolling release”, ”early access”, or similar are not recommended.
PostgreSQL		15.0-18.X	Required if PostgreSQL is used as Zabbix backend database. Depending on the installation size, it might be required to increase PostgreSQL *work_mem* configuration property (4MB being the default value), so that the amount of memory used by the database for particular operation is sufficient and query execution does not take too much time.
TimescaleDB for PostgreSQL		2.20.X-2.29.X	Required if TimescaleDB is used as a PostgreSQL database extension. Make sure to install TimescaleDB Community Edition, which supports compression. See TimescaleDB documentation for details regarding PostgreSQL and TimescaleDB version compatibility.
Elasticsearch		7.X	Required if Elasticsearch is used as Zabbix backend database. Elasticsearch can be used only for storing history data and is only supported with Zabbix servers. The support of Elasticsearch is currently experimental. See also required software for server/proxy.

Software	Mandatory status	Supported versions	Comments
ClickHouse		26.3.0.0-26.9.X.X	Required if ClickHouse is used as Zabbix backend database for storing history data (supported only with Zabbix servers) or as a telemetry data source. To use ClickHouse as a telemetry data source, ClickHouse 26.8.17.4 or later is required. The system time on Zabbix server and the ClickHouse server must also be synchronized; otherwise, telemetry queries may miss data or return the same data twice.
SQLite	Optional	3.3.5-3.53.X	SQLite is only supported with Zabbix proxies. Required if SQLite is used as Zabbix proxy database.
smartmontools		7.1 or later	Required for Zabbix agent 2.
who			Required for the user count plugin.
dpkg			Required for the system.sw.packages plugin.
pkgtool			Required for the system.sw.packages plugin.
rpm			Required for the system.sw.packages plugin.
pacman			Required for the system.sw.packages plugin.
q applets			`qlist` and `qsize`, as part of q applets, are required for the system.sw.packages plugin on Gentoo Linux.

Note:
Although Zabbix can work with databases available in the operating systems, for the best experience, we recommend using databases installed from the official database developer repositories.

Frontend

The minimum supported screen width for Zabbix frontend is 1200px.

If stated as mandatory, the required software/library is strictly necessary. Optional ones are needed for supporting some specific function.

Software	Mandatory status	Supported versions	Comments
PHP	Yes	8.2.0 - 8.5.X	
Apache	One of	2.4 or later	
Nginx		1.20 or later	
MySQL	One of	See Third-party external surrounding software	
PostgreSQL			
PHP extensions			
mysqli	Yes		Required if MySQL is used as Zabbix backend database.
pgsql			Required if PostgreSQL is used as Zabbix backend database.
bcmath			php-bcmath (*--enable-bcmath*)
mbstring			php-mbstring (*--enable-mbstring*)
sockets			php-net-socket (*--enable-sockets*); required for user script support.
gd		2.0.28 or later	php-gd (if provided as a separate package by the distributor); PHP GD extension must support PNG images (*--with-png-dir*), JPEG images (*--with-jpeg-dir*), and FreeType 2 (*--with-freetype-dir*). Version 2.3.0 or later may be required to avoid possible text overlapping in graphs for some frontend languages.
libxml		2.6.15 or later	php-xml (if provided as a separate package by the distributor)
xmlwriter			php-xmlwriter (if provided as a separate package by the distributor)
xmlreader			php-xmlreader (if provided as a separate package by the distributor)
ctype			php-ctype (*--enable-ctype*)

Software	Mandatory status	Supported versions	Comments
session			php-session (if provided as a separate package by the distributor)
ldap	No		php-ldap; required for LDAP authentication.
openssl			php-openssl; required for SAML authentication.
gettext			php-gettext (*--with-gettext*); required for translations.
cURL		7.19.4 or later	php-curl; required for Duo Universal Prompt MFA, SMTP authentication, and for testing the telemetry data source.

Third-party frontend libraries that are supplied with Zabbix:

Library	Mandatory status	Supplied version	Comments
jQuery JavaScript Library	Yes	3.6.0	JavaScript library that simplifies the process of cross-browser development.
jQuery UI		1.12.1	A set of user interface interactions, effects, widgets, and themes built on top of jQuery.
SAML PHP Toolkit		4.3.1	A PHP toolkit that adds SAML 2.0 authentication support to be able to sign in to Zabbix.
Duo Universal PHP library		1.1.2	A PHP library that adds Duo Universal Prompt multi-factor authentication support for Zabbix.
Symfony Yaml Component		5.1.0	Adds support to export and import Zabbix configuration elements in the YAML format.

Note:
Zabbix may work on previous versions of Apache, MySQL, and PostgreSQL as well.

Attention:
For other fonts than the default DejaVu, PHP function imagerotate might be required. If it is missing, these fonts might be rendered incorrectly when a graph is displayed. This function is only available if PHP is compiled with bundled GD, which is not the case in Debian and other distributions.

Third-party libraries used for writing and debugging Zabbix frontend code:

Library	Mandatory status	Minimum version	Description
Composer	No	2.4.1	An application-level package manager for PHP that provides a standard format for managing dependencies of PHP software and required libraries.
PHPUnit		8.5.29	A PHP unit testing framework for testing Zabbix frontend.
SASS		3.4.22	A preprocessor scripting language that is interpreted and compiled into Cascading Style Sheets (CSS).

Web browser on client side

Cookies and JavaScript must be enabled.

The latest stable versions of Google Chrome, Mozilla Firefox, Microsoft Edge, Apple Safari, and Opera are supported.

Warning:
The same-origin policy for IFrames is implemented, which means that Zabbix cannot be placed in frames on a different domain.

 Still, pages placed into a Zabbix frame will have access to Zabbix frontend (through JavaScript) if the page that is placed in the frame and Zabbix frontend are on the same domain. A page like `http://secure-zabbix.com/cms/page.html`, if placed into dashboards on `http://secure-zabbix.com/zabbix/`, will have full JS access to Zabbix.

Server/proxy

If stated as mandatory, the required software/library is strictly necessary. Optional ones are needed for supporting some specific function.

Requirement	Mandatory status	Description
libpcre2	Yes	PCRE2 library is required for Perl Compatible Regular Expression (PCRE) support. PCRE2 v10.x is supported.
libevent		Required for inter-process communication. Version 2.0.10 or higher.
libevent-pthreads		Required for inter-process communication.
libpthread		Required for mutex and read-write lock support (could be part of libc).
libresolv		Required for DNS resolution (could be part of libc).
libiconv		Required for text encoding/format conversion (could be part of libc). Mandatory for Zabbix server on Linux.
libz		Required for compression support.
libm		Math library. Required by Zabbix server only.
libmysqlclient	One of	Required if MySQL is used.
libmariadb		Required if MariaDB is used.
libpq5		Required if PostgreSQL is used; *libpq5* version must match or be higher than the version of the PostgreSQL database used.
libsqlite3		Required if SQLite is used. Required for Zabbix proxy only.
libOpenIPMI	No	Required for IPMI support. Required for Zabbix server only.
libssh2 or *libssh*		Required for SSH checks. Version 1.8.0 or higher (libssh2); 0.9.0 or higher (libssh).
libcurl		Required for the following features: - Web monitoring, VMware monitoring, and HTTP agent items (for all: version 7.19.1 or higher); - Zabbix agent web.page.* items (version 7.19.1 or higher; see also agent requirements); - SMTP authentication (Basic: version 7.20.0 or higher; OAuth: version 7.33 or higher; see also frontend requirements); - Elasticsearch (version 7.28.0 or higher); - Telemetry query items (version 7.28.0 or higher). Version 7.28.0 or higher for all features is recommended. To use upgraded cURL features for `web.page.*` items, restart Zabbix server/proxy. For SMTP authentication, use the `libcurl-full` package at runtime.
libxml2		Required for VMware monitoring and XML XPath preprocessing.
net-snmp		Required for SNMP support. Version 5.3.0 or higher. Support of strong encryption protocols (AES192/AES192C, AES256/AES256C) is available starting with net-snmp library 5.8; on RHEL 8+ based systems it is recommended to use net-snmp 5.8.15 or later.
libunixodbc		Required for database monitoring.
libgnutls or *libopenssl*		Required when using encryption. Minimum versions: *libgnutls* - 3.1.18, *libopenssl* - 1.0.1
libldap		Required for LDAP support.
fping		Required for ICMP ping items.
c-ares		Required for asynchronous DNS resolution if Zabbix is configured with `--with-ares` option. Otherwise, *libevent* will be used. Minimum version: 1.16.0. The DNS query caching and failover feature (query cache enabled by default) requires c-ares 1.26.0 or later. On systems where c-ares is provided by the OS package, make sure the installed runtime library exposes the resolver features required for DNS caching and failover.
g++		C++ compiler. Required to compile Zabbix proxy with the `--with-apm` option.
gRPC		Required for APM support on Zabbix proxy. Required libraries: *grpc++*, *grpc*, *gpr*. The `grpc_cpp_plugin` utility is also required to compile Zabbix proxy. The required version is available in the standard package repositories of these operating systems: - Ubuntu 23.04 or later; - Debian 12 (Bookworm) or later; - Alpine Linux 3.15 or later; - RHEL, Rocky Linux, CentOS Stream 9 or later (only from the EPEL repository); - Fedora X or later. On older operating system versions, *gRPC* must be built from source.
protobuf		Required for APM support on Zabbix proxy. The `protoc` compiler is also required to compile Zabbix proxy.

Requirement	Mandatory status	Description
Abseil		Required for APM support on Zabbix proxy. Required libraries: *absl_synchronization*, *absl_base*. Usually installed together with *gRPC*.

Agent

Requirement	Mandatory status	Description
libpcre2	Yes	PCRE2 library is required for Perl Compatible Regular Expression (PCRE) support. PCRE2 v10.x is supported. Required for log monitoring. Also required on Windows.
libpthread		Required for mutex and read-write lock support (could be part of libc). Not required on Windows.
libresolv		Required for DNS resolution (could be part of libc). Not required on Windows.
libiconv		Required for text encoding/format conversion to UTF-8 in log items, file content, file regex and regmatch items (could be part of libc). Not required on Windows.
libgnutls or *libopenssl*	No	Required if using encryption. Minimum versions: *libgnutls* - 3.1.18, *libopenssl* - 1.0.1 On Microsoft Windows OpenSSL 1.1.1 or later is required.
libldap		Required if LDAP is used. Not supported on Windows.
libcurl		Required for extended support of Zabbix agent web.page.* items. Without libcurl, basic functionality is available (e.g., web.page.get[http://example.com]). With libcurl, the agent supports additional features, such as HTTP URLs with credentials (e.g., http://user:password@example.com) and HTTPS URLs. Version 7.19.1 or higher is required (7.28.0 or higher is recommended). To use upgraded cURL features, restart Zabbix agent.
libmodbus		Only required if Modbus monitoring is used. Version 3.0 or higher.

Agent 2

Requirement	Mandatory status	Description
Go	Yes	Required to build Zabbix agent 2 and its plugins from source. Go 1.25 or later is supported. See go.dev for installation instructions. Go libraries used by Zabbix agent 2 and its plugins are listed in the Zabbix Git repository (libraries marked as `indirect` in the repository are dependencies of other required libraries): - Zabbix agent 2 - Plugin support - PostgreSQL - MongoDB - MSSQL - Ember+ - NVIDIA GPU - Example plugin
libpcre2	Yes	PCRE2 library is required for Perl Compatible Regular Expression (PCRE) support. PCRE2 v10.x is supported. Required for log monitoring. Also required on Windows.
libopenssl	No	Required when using encryption. OpenSSL 1.0.1 or later is required on UNIX platforms. The OpenSSL library must have PSK support enabled. LibreSSL is not supported. On Microsoft Windows systems OpenSSL 1.1.1 or later is required.

Web service

The latest stable version of Google Chrome is supported for generating scheduled reports using the Zabbix web service.

The required Go version for building the web service matches the one used for Zabbix agent 2.

MCP server

The required Go version for building the MCP server is Go 1.25 or later. See go.dev for installation instructions.

The MCP server is officially supported on Linux only. Connecting MCP clients must support the Streamable HTTP transport; see Zabbix MCP server for details.

Java gateway

If you obtained Zabbix from the source repository or an archive, then the necessary dependencies are already included in the source tree.

If you obtained Zabbix from your distribution's package, then the necessary dependencies are already provided by the packaging system.

In both cases above, the software is ready to be used and no additional downloads are necessary.

If, however, you wish to provide your versions of these dependencies (for instance, if you are preparing a package for some Linux distribution), below is the list of library versions that Java gateway is known to work with. Zabbix may work with other versions of these libraries, too.

The following table lists JAR files that are currently bundled with Java gateway in the original code:

Library	Mandatory status	Bundled version	Comments
android-json	Yes	4.3r1	JSON (JavaScript Object Notation) is a lightweight data-interchange format. This is the `org.json` compatible Android implementation extracted from the Android SDK.
logback-classic		1.5.25	
logback-core		1.5.25	
slf4j-api		2.0.16	

The Java gateway can be built using either Oracle Java or open source OpenJDK (version 8 or newer). Packages provided by Zabbix are compiled using OpenJDK. The following table lists OpenJDK packages used for building Zabbix packages by distribution:

Distribution	OpenJDK package
AlmaLinux 10	`java-21-openjdk-devel` (amd64, arm64: 21.0.12.0.8-1.2)
AlmaLinux 9	`java-17-openjdk-devel` (amd64, arm64: 17.0.15.0.6-2)
AlmaLinux 8	`java-1.8.0-openjdk-devel` (amd64: 1.8.0.332.b09-2; arm64: 1.8.0.382.b05-2)
Amazon Linux 2023	`java-22-amazon-corretto-devel` (amd64, arm64: 22.0.2+9-1)
CentOS Stream 10	`java-21-openjdk-devel` (amd64, arm64: 21.0.11.0.10-2)
CentOS Stream 9	`java-17-openjdk-devel` (amd64, arm64: 17.0.15.0.6-2)
CentOS Stream 8	`java-1.8.0-openjdk-devel` (amd64, arm64: 1.8.0.362.b08-3)
CentOS 7	`java-1.8.0-openjdk-devel` (amd64: 1.8.0.282.b08-1)
Debian 13	`default-jdk-headless` (amd64, arm64: 2:1.21-76)
Debian 12	`default-jdk-headless` (amd64, arm64: 2:1.17-74)
Debian 11	`default-jdk-headless` (amd64: 2:1.11-72)
OpenSUSE Leap 16	`java-21-openjdk-devel` (amd64, arm64: 21.0.11.0-160000.2.1)
OpenSUSE Leap 15	`java-17-openjdk-devel` (amd64: 17.0.5.0-150400.3.9.3; arm64: 17.0.8.0-150400.3.27.1)
Oracle Linux 10	`java-21-openjdk-devel` (amd64, arm64: 21.0.11.0.10-2.0.1)
Oracle Linux 9	`java-17-openjdk-devel` (amd64, arm64: 17.0.15.0.6-2.0.1)
Oracle Linux 8	`java-1.8.0-openjdk-devel` (amd64: 1.8.0.372.b07-4.0.1); `java-11-openjdk-devel` (arm64: 11.0.20.0.8-3.0.1)
Oracle Linux 7	`java-1.8.0-openjdk-devel` (amd64: 1.8.0.282.b08-1)
Raspberry Pi OS 13	`default-jdk-headless` (arm64, armhf: 2:1.21-76)
Raspberry Pi OS 12	`default-jdk-headless` (arm64, armhf: 2:1.17-74)
Raspberry Pi OS 11	`default-jdk-headless` (arm64: 2:1.11-72; armhf: 2:1.11-72+b4)
RHEL 10	`java-21-openjdk-devel` (amd64, arm64: 21.0.12.0.8-1.2)
RHEL 9	`java-17-openjdk-devel` (amd64, arm64: 17.0.15.0.6-2)
RHEL 8	`java-1.8.0-openjdk-devel` (amd64: 1.8.0.372.b07-4; arm64: 1.8.0.382.b05-2)
RHEL 7	`java-1.8.0-openjdk-devel` (amd64: 1.8.0.282.b08-1)
Rocky Linux 10	`java-21-openjdk-devel` (amd64, arm64: 21.0.12.0.8-1.2)
Rocky Linux 9	`java-17-openjdk-devel` (amd64, arm64: 17.0.15.0.6-2)

Distribution	OpenJDK package
Rocky Linux 8	`java-1.8.0-openjdk-devel` (amd64: 1.8.0.372.b07-4; arm64: 1.8.0.382.b05-2)
SLES 16	`java-21-openjdk-devel` (amd64, arm64: 21.0.11.0-160000.1.1)
SLES 15	`java-17-openjdk-devel` (amd64: 17.0.5.0-150400.3.9.3; arm64: 17.0.8.0-150400.3.27.1)
Ubuntu 26.04	`default-jdk-headless` (amd64, arm64: 2:1.25-77)
Ubuntu 24.04	`default-jdk-headless` (amd64, arm64: 2:1.21-75+exp1)
Ubuntu 22.04	`default-jdk-headless` (amd64, arm64: 2:1.11-72build2)
Ubuntu 20.04	`default-jdk-headless` (amd64, arm64: 2:1.11-72)

Bridge adapter

Zabbix bridge adapter has no special requirements.

Default ports

The following list contains ports that Zabbix components may use in a default configuration:

Zabbix component	Protocol and port	Connection initiated by
Zabbix agent	TCP 10050	Zabbix server or Zabbix proxy (passive checks only)
Zabbix agent 2	TCP 10050	Zabbix server or Zabbix proxy (passive checks only)
Zabbix server	TCP 10051	Active Zabbix agents, active Zabbix proxy, Zabbix sender, Zabbix frontend
Zabbix proxy	TCP 10051	Active Zabbix agents, Zabbix sender, Zabbix server (passive proxy only)
Zabbix Java gateway	TCP 10052	Zabbix server or proxy
Zabbix web service	TCP 10053	Zabbix server
Zabbix MCP server	TCP 8443	MCP clients (AI agents)
Zabbix frontend	TCP 80	Web browsers, Zabbix web service (scheduled reports), API clients
	TCP 443	
SNMP traps (snmptrapd)	UDP 162	SNMP devices
Zabbix bridge adapter	TCP 10054	Zabbix server

Warning:
Open only the ports required by the components you actually run, and restrict each firewall rule to the specific hosts that need to connect (for example, allow port 10050 only from the Zabbix server or proxy, not from any address). Do not expose Zabbix ports directly to the internet.
 Before opening ports, harden the installation: enable encryption between components, restrict allowed hosts in component configuration files, and review the security best practices.

Database size

Zabbix configuration data require a fixed amount of disk space and do not grow much.

Zabbix database size mainly depends on these variables, which define the amount of stored historical data:

- Number of processed values per second

This is the average number of new values Zabbix server receives every second. For example, if we have 3000 items for monitoring with a refresh rate of 60 seconds, the number of values per second is calculated as 3000/60 = **50**.

It means that 50 new values are added to Zabbix database every second.

- Housekeeper settings for history

Zabbix keeps values for a fixed period of time, normally several weeks or months. Each new value requires a certain amount of disk space for data and index.

So, if we would like to keep 30 days of history and we receive 50 values per second, the total number of values will be around (**30***24*3600)* **50** = 129.600.000, or about 130M of values.

Depending on the database engine used, type of received values (floats, integers, strings, log files, etc), the disk space for keeping a single value may vary from 40 bytes to hundreds of bytes. Normally it is around 90 bytes per value for numeric items[2]. In our case, it means that 130M of values will require 130M * 90 bytes = **10.9GB** of disk space.

Note:
The size of text/log item values is impossible to predict exactly, but you may expect around 500 bytes per value.

- Housekeeper setting for trends

Zabbix keeps a 1-hour max/min/avg/count set of values for each item in the table **trends**. The data is used for trending and long period graphs. The one hour period cannot be customized.

Zabbix database, depending on the database type, requires about 90 bytes per each total. Suppose we would like to keep trend data for 5 years. Values for 3000 items will require 3000*24*365* **90** = **2.2GB** per year, or **11GB** for 5 years.

- Housekeeper settings for events

Each Zabbix event requires approximately 250 bytes of disk space[1]. It is hard to estimate the number of events generated by Zabbix daily. In the worst-case scenario, we may assume that Zabbix generates one event per second.

For each recovered event, an event_recovery record is created. Normally most of the events will be recovered so we can assume one event_recovery record per event. That means additional 80 bytes per event.

Optionally events can have tags, each tag record requiring approximately 100 bytes of disk space[1]. The number of tags per event (#tags) depends on configuration. So each will need an additional #tags * 100 bytes of disk space.

It means that if we want to keep 3 years of events, this would require 3*365*24*3600* (250+80+#tags*100) = **~30GB**+#tags*100B disk space[2].

Note:
[1] More when having non-ASCII event names, tags and values.
 [2] The size approximations are based on MySQL and might be different for other databases.

The table contains formulas that can be used to calculate the disk space required for Zabbix system:

Parameter	Formula for required disk space (in bytes)
Zabbix configuration	Fixed size. Normally 10MB or less.
History	days*(items/refresh rate)*24*3600*bytes items : number of items days : number of days to keep history refresh rate : average refresh rate of items bytes : number of bytes required to keep single value, depends on database engine, normally ~90 bytes.
Trends	days*(items/3600)*24*3600*bytes items : number of items days : number of days to keep history bytes : number of bytes required to keep single trend, depends on the database engine, normally ~90 bytes.
Events	days*events*24*3600*bytes events : number of event per second. One (1) event per second in worst-case scenario. days : number of days to keep history bytes : number of bytes required to keep single event, depends on the database engine, normally ~330 + average number of tags per event * 100 bytes.

So, the total required disk space can be calculated as:

Configuration + History + Trends + Events

The disk space will NOT be used immediately after Zabbix installation. Database size will grow then it will stop growing at some point, which depends on housekeeper settings.

Time synchronization

It is very important to have precise system time on the server with Zabbix running. ntpd is the most popular daemon that synchronizes the host's time with the time of other machines. It's strongly recommended to maintain synchronized system time on all systems Zabbix components are running on.

What's new and upgrade notes: 8.2

Zabbix 8.2.0 continues from Zabbix 8.0.0, adding new features and enhancements.

Also on this page:

- What's new in templates, including new templates and changes to existing templates.
- Upgrade notes, including breaking changes.

See also: API changes.

Upgrade notes

These notes are for upgrading from Zabbix 8.0.x to Zabbix 8.2.0.

All notes are grouped into:

- **Breaking changes** - changes that may break existing installations and other critical information related to the upgrade process
- **Other** - all remaining information describing the changes in Zabbix functionality

See also:

- Upgrade procedure for all relevant information about upgrading from versions before Zabbix 8.2.0;
- Upgrading HA cluster for instructions on upgrading servers in a **high-availability** (HA) cluster.

Breaking changes

Other

What's new and upgrade notes: 8.2.x

This page introduces all new features included in minor releases of the major Zabbix version. Template changes and upgrade notes are also included.

See also What's new of the major version.

What's new in Zabbix 8.2.1 This version is not released yet.

See upgrade notes for this version.

Upgrade notes

Zabbix appliance

Overview Zabbix appliance provides a way to instantly deploy Zabbix server and frontend, instead of setting them up manually or reusing an existing server for Zabbix.

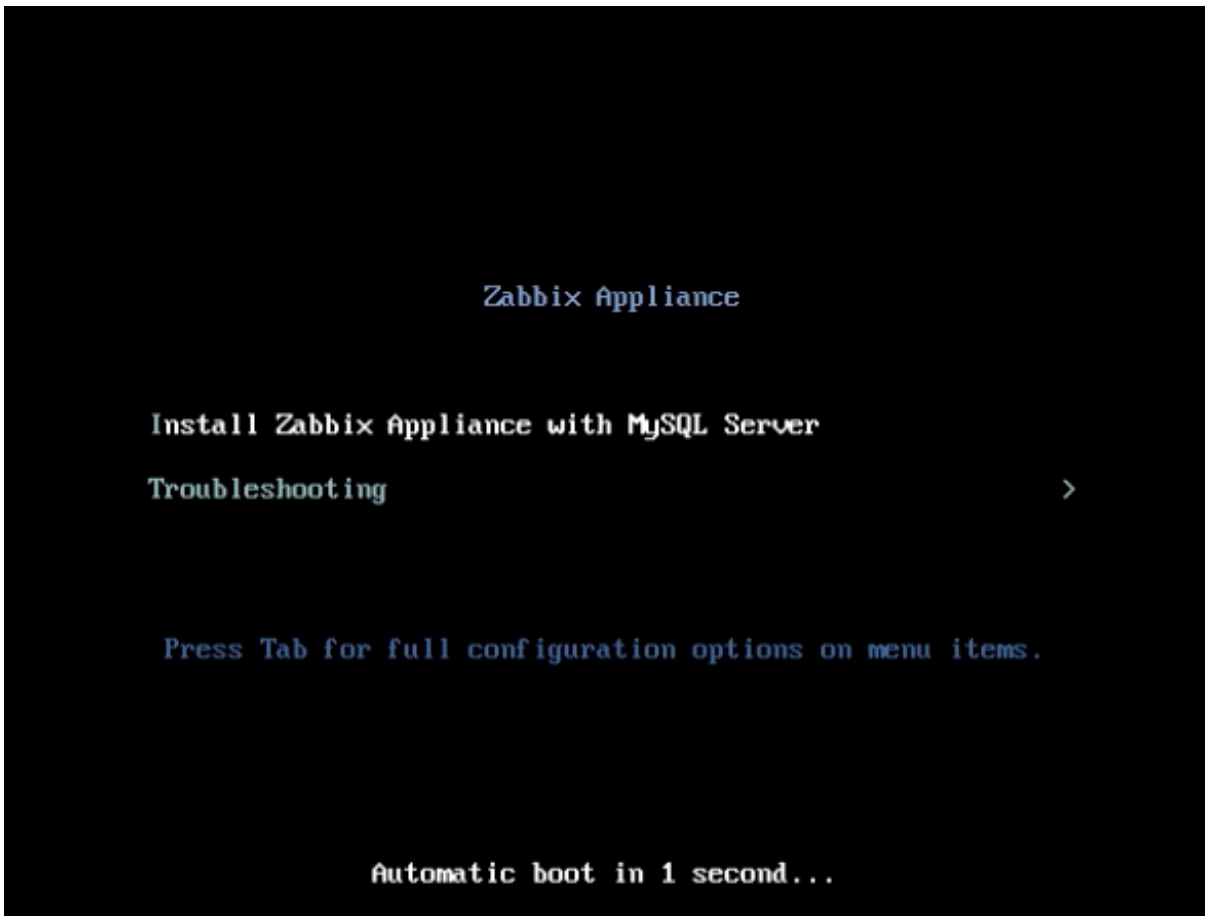
The appliance for Zabbix 8.2 is based on Red Hat Enterprise Linux 10 derivatives such as AlmaLinux or Rocky Linux and contains a pre-configured Zabbix server running on MySQL and the frontend running on an NGINX web server.

> **Attention:**
> This appliance is designed for evaluating Zabbix. Using it in serious production environments is discouraged.

Appliance images are available for download in the following formats:

- Installation CD/DVD (.iso)
- VMware (.vmx) - see notes
- Open virtualization format (.ovf)
- Microsoft Hyper-V (.vhd/.vhdx) - see notes
- KVM, Parallels, QEMU, USB stick, VirtualBox, Xen (.raw) - see notes
- KVM, QEMU (.qcow2)

Zabbix installation CD/DVD boot menu:



Quick start Pre-requisites

1. Make sure the host machine has sufficient resources to meet the virtual machine's system requirements:

- *RAM*: 4 GB
- *Disk space*: at least 8 GB should be allocated for the virtual machine
- *CPU*: 2 cores minimum

2. If not yet installed, install the virtualization software for booting the appliance image (for example, VirtualBox).
3. Download the appliance in the format supported by your virtualization software.
4. Check network settings to ensure that DHCP is enabled on the host machine.

Installation

1. Boot the appliance virtual machine from the downloaded image. For instructions, please refer to the documentation of your virtualization software, for example, VirtualBox documentation.
2. Configure virtual machine network settings to allow access from a browser on the host machine. This can be achieved by enabling *Bridged mode*.
3. Log in to the virtual machine using default system credentials.
4. To retrieve the IP address, run the following command on the virtual machine:

```
ip addr show
```

5. Open a browser on the host machine and point it to the IP address the appliance has received over DHCP.
6. Log in to Zabbix using default credentials and start monitoring.

Configuration This section describes frequently required default configuration settings along with available customization options.

Credentials

System

- Username: `root`

- Password: `zabbix`

Zabbix frontend

- Username: `Admin`
- Password: `zabbix`

After logging in, you can change default password in the user profile settings or create new users and delete the default one.

Database

Passwords for all database users are randomly generated during the installation process. The following users are defined for the database:

Root:

- Username: `root`
- Password: the password is stored in the `/root/.my.cnf` file It is not required to input a password under the root account.

Zabbix server:

- Username: `zabbix_srv`
- Password: the password is stored in the server configuration file

Zabbix frontend:

- Username: `zabbix_web`
- Password: the password is stored in the frontend configuration file

To change a database user password, modify it in MySQL and the corresponding configuration file.

Frontend access

Zabbix frontend can be accessed at *http://<virtual machine's IP>*.

By default, access is allowed from anywhere. To limit the access, modify `/etc/nginx/conf.d/zabbix.conf`. After saving the edited file, restart NGINX by logging in via SSH as root user and executing:

```
systemctl restart nginx
```

Static IP address

By default, the appliance uses DHCP to obtain the IP address. To set a static IP address:

1. Log in as root user.
2. Run the following commands, replacing values with your custom IP address:

```
nmcli con mod eth0 ipv4.addresses 192.168.1.10/24    # Appliance IP address/CIDR prefix
nmcli con mod eth0 ipv4.gateway 192.168.1.1     # Gateway IP address
nmcli con mod eth0 ipv4.dns 8.8.8.8    # DNS server IP address
nmcli con mod eth0 ipv4.method manual
nmcli con up eth0
```

Alternatively, you can open the file `/etc/NetworkManager/system-connections/eth0.nmconnection` and add the following lines:

```
[ipv4]
address1=192.168.1.10/24
dns=8.8.8.8;
gateway=192.168.1.1
method=manual
```

After saving the edited file, run the commands:

```
nmcli con reload
nmcli con up eth0
```

To verify that the network configuration changes have been applied, run the following command:

```
nmcli dev show eth0
```

Firewall configuration

To manage firewall settings, the appliance uses iptables with predefined rules:

- Open SSH port (22 TCP)

- Open Zabbix agent (10050 TCP) and Zabbix trapper (10051 TCP) ports
- Open HTTP (80 TCP) and HTTPS (443 TCP) ports
- Open SNMP trap port (162 UDP)
- Open outgoing connections to NTP port (123 UDP)
- Limit ICMP packets to 5 packets per second
- Drop all other incoming connections

To open additional ports, modify `/etc/sysconfig/iptables` file and reload firewall rules:

```
systemctl reload iptables
```

Repositories

Zabbix appliance uses *zabbix-release* package from the Zabbix repository. Repositories are configured in `/etc/yum.repos.d/*` directory.

Time zone

By default, the appliance uses UTC for the system clock. To change the time zone, copy the appropriate file from `/usr/share/zoneinfo` to `/etc/localtime`, for example:

```
cp /usr/share/zoneinfo/Europe/Riga /etc/localtime
```

> **Note:**
> Zabbix frontend time zone is set separately and can be changed in the frontend settings. The default time zone for Zabbix frontend is Europe/Riga.

File locations

- The Zabbix server and agent configuration files are located in `/etc/zabbix`.
- The Zabbix server and agent log files are located in `/var/log/zabbix`.
- The Zabbix frontend configuration file is located in `/etc/zabbix/web`.
- The Zabbix frontend is located in `/usr/share/zabbix`.
- The home directory for *zabbix* user is `/var/lib/zabbix`.

System services

Systemd services are available. To see the list of Zabbix services, run the following command on the virtual machine:

```
systemctl list-units zabbix*
```

Image format-specific notes VMware

The images in *vmdk* format are usable directly in VMware Player, Server, and Workstation products. For use in ESX, ESXi, and vSphere they must be converted using VMware vCenter Converter (authentication required for download). If you use VMware vCenter Converter, you may encounter issues with the hybrid network adapter. In that case, you can try specifying the E1000 adapter during the conversion process. Alternatively, after the conversion is complete, you can delete the existing adapter and add an E1000 adapter.

HDD/flash image (raw)

To boot the image, run:

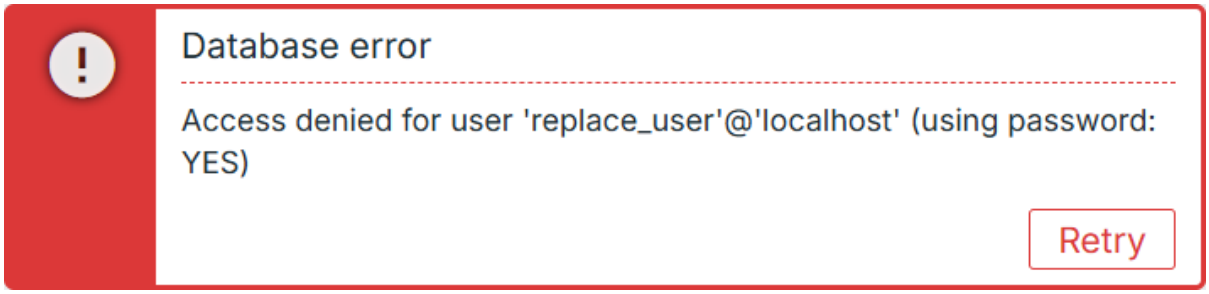
```
dd if=./zabbix_appliance_8.0.0.raw of=/dev/sdc bs=4k conv=fdatasync
```

Replace `/dev/sdc` with your Flash/HDD disk device path.

Hyper-V

If the appliance fails to start up in Hyper-V, try pressing **Ctrl+Alt+F2** to switch a TTY session.

Troubleshooting If you encounter the error message *Access denied for user 'replace_user'@'localhost' (using password: YES)* while trying to log in to the frontend, it may indicate that the installation is still in progress.



If the error persists after waiting a few minutes, or if you observe any other unexpected behavior, it likely means that the installation process has not completed successfully. In this case, we recommend deleting the current appliance and redeploying it by following the same installation instructions. This step typically resolves the issue.

Note that attempting to manually fix a broken installation is not recommended, as it may lead to further complications.

Processes

Overview Zabbix consists of several processes that work together to monitor systems and send notifications. Each process has a specific role in the monitoring infrastructure.

Core processes

- Server - the central process that performs polling and trapping of data, calculates triggers, and sends notifications to users.
- Agent - collects data from monitored hosts.
- Agent 2 - a new generation agent with plugin architecture and enhanced capabilities.
- Proxy - collects monitoring data from monitored devices on behalf of Zabbix server and sends the information to a centralized Zabbix server in a distributed setup.
- Java gateway - enables monitoring of Java applications.
- MCP server - allows AI agents to work with Zabbix using the Model Context Protocol.

For easy access to Zabbix from anywhere and from any platform, the web-based interface is provided. The interface is part of Zabbix server, and usually (but not necessarily) runs on the same physical machine as the one running the server.

All configuration information as well as the data gathered by Zabbix is stored in a database.

Header The header is present in all request and response messages between Zabbix components. It is required to determine the message length, if it is compressed or not, if it is a large packet or not.

Zabbix communications protocol has 1GB packet size limit per connection. The limit of 1GB is applied to both the received packet data length and the uncompressed data length.

When sending configuration to Zabbix proxy, the packet size limit is increased to 4GB to allow syncing large configurations. When data length before compression exceeds 4GB, Zabbix server automatically starts using the large packet format (0x04 flag) which increases the packet size limit to 16GB.

Note that while a large packet format can be used for sending any data, currently only the Zabbix proxy configuration syncer can handle packets that are larger than 1GB.

Header structure The header consists of four fields. All numbers in the header are formatted as little-endian.

Field	Size	Size (large packet)	Description
<PROTOCOL>	4	4	"ZBXD" or 5A 42 58 44
<FLAGS>	1	1	Protocol flags: 0x01 - Zabbix communications protocol 0x02 - compression 0x04 - large packet
<DATALEN>	4	8	Data length.
<RESERVED>	4	8	When compression is used (0x02 flag) - the length of uncompressed data When compression is not used - 00 00 00 00

Header examples Here are some code snippets showing how to add Zabbix protocol header to the data you want to send in order to obtain the packet you should send to Zabbix so that it is interpreted correctly. These code snippets assume that the data is not larger than 1GB, thus the large packet format is not used.

Python

```
packet = b"ZBXD\1" + struct.pack("<II", len(data), 0) + data
```

or

```
def zbx_create_header(plain_data_size, compressed_data_size=None):
    protocol = b"ZBXD"
    flags = 0x01
    if compressed_data_size is None:
        datalen = plain_data_size
        reserved = 0
    else:
        flags |= 0x02
        datalen = compressed_data_size
        reserved = plain_data_size
    return protocol + struct.pack("<BII", flags, datalen, reserved)

packet = zbx_create_header(len(data)) + data
```

Perl

```
my $packet = "ZBXD\1" . pack("(II)<", length($data), 0) . $data;
```

or

```
sub zbx_create_header($;$)
{
    my $plain_data_size = shift;
    my $compressed_data_size = shift;

    my $protocol = "ZBXD";
    my $flags = 0x01;
    my $datalen;
    my $reserved;

    if (!defined($compressed_data_size))
    {
        $datalen = $plain_data_size;
        $reserved = 0;
    }
    else
    {
        $flags |= 0x02;
        $datalen = $compressed_data_size;
        $reserved = $plain_data_size;
    }

    return $protocol . chr($flags) . pack("(II)<", $datalen, $reserved);
}

my $packet = zbx_create_header(length($data)) . $data;
```

PHP

```
$packet = "ZBXD\1" . pack("VV", strlen($data), 0) . $data;
```

or

```
function zbx_create_header($plain_data_size, $compressed_data_size = null)
{
    $protocol = "ZBXD";
    $flags = 0x01;
    if (is_null($compressed_data_size))
```

```
    {
        $datalen = $plain_data_size;
        $reserved = 0;
    }
    else
    {
        $flags |= 0x02;
        $datalen = $compressed_data_size;
        $reserved = $plain_data_size;
    }
    return $protocol . chr($flags) . pack("VV", $datalen, $reserved);
}

$packet = zbx_create_header(strlen($data)) . $data;
```

Bash

```
datalen=$(printf "%08x" ${#data})
datalen="\\x${datalen:6:2}\\x${datalen:4:2}\\x${datalen:2:2}\\x${datalen:0:2}"
printf "ZBXD\1${datalen}\0\0\0\0%s" "$data"
```

Server

Overview Zabbix server is the central process of Zabbix software.

The server performs the polling and trapping of data, it calculates triggers, sends notifications to users. It is the central component to which Zabbix agents and proxies report data on availability and integrity of systems. The server can itself remotely check networked services (such as web servers and mail servers) using simple service checks.

The server is the central repository in which all configuration, statistical and operational data is stored, and it is the entity in Zabbix that will actively alert administrators when problems arise in any of the monitored systems.

The functioning of a basic Zabbix server is broken into three distinct components; they are: Zabbix server, web frontend and database storage.

All of the configuration information for Zabbix is stored in the database, which both the server and the web frontend interact with. For example, when you create a new item using the web frontend (or API) it is added to the items table in the database. Then, about once a minute Zabbix server will query the items table for a list of the items which are active that is then stored in a cache within the Zabbix server. This is why it can take up to two minutes for any changes made in Zabbix frontend to show up in the latest data section.

See also:

- High availability setup for Zabbix servers
- Configuration parameters: Zabbix server

Running server If installed as package

Zabbix server runs as a daemon process. The server can be started by executing:

```
systemctl start zabbix-server
```

This will work on most of GNU/Linux systems. On other systems you may need to run:

```
/etc/init.d/zabbix-server start
```

Similarly, for stopping/restarting/viewing status, use the following commands:

```
systemctl stop zabbix-server
systemctl restart zabbix-server
systemctl status zabbix-server
```

Start up manually

If the above does not work you have to start it manually. Find the path to the `zabbix_server` binary and execute:

```
zabbix_server
```

You can use the following command-line parameters with Zabbix server:

Parameter	Description
`-c --config <config-file>`	Path to the configuration file.
`-f --foreground`	Run Zabbix server in foreground.
`-R --runtime-control <option>`	Perform administrative functions.
`-T --test-config`	Validate configuration file and exit.
`-h --help`	Print help information and exit.
`-V --version`	Display version number.

Examples of running Zabbix server with command-line parameters:

```
zabbix_server -c /usr/local/etc/zabbix_server.conf
zabbix_server --help
zabbix_server -V
```

Runtime control

Runtime control options:

Option	Description	Target
`config_cache_reload`	Reload configuration cache. Ignored if cache is being currently loaded.	
`history_cache_clear=<target>`	Clear history cache for the item specified by its ID. Affects all values of the item, except the first and last value.	**target** - ID of the item
`diaginfo[=<section>]`	Gather diagnostic information in the server log file.	`historycache` - history cache statistics; `valuecache` - value cache statistics; `preprocessing` - preprocessing manager statistics; `alerting` - alert manager statistics; `lld` - LLD manager statistics; `locks` - list of mutexes (is empty on *BSD* systems); `connector` - statistics for connectors with the largest queue.
`ha_status`	Log high availability (HA) cluster status.	
`ha_remove_node=<target>`	Remove the high availability (HA) node specified by its name or ID. Note that active/standby nodes cannot be removed.	**target** - name or ID of the node (can be obtained by running `ha_status`).
`ha_set_failover_delay=<delay>`	Set high availability (HA) failover delay. Time suffixes are supported, e.g., 10s, 1m.	
`proxy_config_cache_reload[=<target>]`	Reload proxy configuration cache.	**target** - comma-delimited list of proxy names. If no target is specified, reload configuration for all proxies.
`secrets_reload`	Reload secrets from Vault.	
`service_cache_reload`	Reload the service manager cache.	
`snmp_cache_reload`	Reload SNMP cache—clear SNMP engine properties (engine time, engine boots, engine id, credentials) for all hosts. Use to force a global cache clear when troubleshooting SNMP issues.	
`housekeeper_execute`	Start the housekeeping procedure. Ignored if the housekeeping procedure is currently in progress.	
`trigger_housekeeper_execute`	Start the trigger housekeeping procedure. Ignored if the trigger housekeeping procedure is currently in progress.	

Option	Description	Target
`log_level_increase[=<target>]`	Increase log level, affects all processes if target is not specified. Not supported on *BSD* systems.	**process type** - all processes of specified type (e.g., `poller`). See all server process types. **process type,N** - process type and number (e.g., `poller,3`). **pid** - process identifier (1 to 65535). For larger values specify target as 'process type,N'.
`log_level_decrease[=<target>]`	Decrease log level, affects all processes if target is not specified. Not supported on *BSD* systems.	
`prof_enable[=<target>]`	Enable profiling. Affects all processes if target is not specified. Enabled profiling provides details of all rwlocks/mutexes by function name.	**process type** - all processes of specified type (e.g., `history syncer`) Supported process types as profiling targets: `alerter`, `alert manager`, `availability manager`, `configuration syncer`, `discovery manager`, `escalator`, `history poller`, `history syncer`, `housekeeper`, `http poller`, `icmp pinger`, `ipmi manager`, `ipmi poller`, `java poller`, `lld manager`, `lld worker`, `odbc poller`, `poller`, `preprocessing manager`, `proxy poller`, `self-monitoring`, `service manager`, `snmp trapper`, `task manager`, `timer`, `trapper`, `unreachable poller`, `vmware collector`. **process type,N** - Process type and number (e.g., `history syncer,1`). **pid** - Process identifier (1 to 65535). For larger values specify target as 'process type,N'. **scope** - `rwlock`, `mutex`, `processing` can be used with the process type and number (e.g., `history syncer,1,processing`) or all processes of type (e.g., `history syncer,rwlock`).
`prof_disable[=<target>]`	Disable profiling. Affects all processes if target is not specified.	**process type** - all processes of specified type (e.g., `history syncer`). Supported process types as profiling targets: see `prof_enable`. **process type,N** - Process type and number (e.g., `history syncer,1`). **pid** - Process identifier (1 to 65535). For larger values specify target as 'process type,N'.

When a process type is specified as a target for runtime control, the command applies only to running processes of that type. If no process of the specified type is currently running, the command returns an error:

```
zabbix_server -R log_level_increase='discovery manager'
Cannot redirect signal: "discovery manager" process does not exist
```

Example of using runtime control to reload the server configuration cache:

```
zabbix_server -c /usr/local/etc/zabbix_server.conf -R config_cache_reload
```

Examples of using runtime control to reload the proxy configuration:

```
### Reload configuration of all proxies:
zabbix_server -R proxy_config_cache_reload

### Reload configuration of Proxy1 and Proxy2:
zabbix_server -R proxy_config_cache_reload=Proxy1,Proxy2
```

Example of using runtime control to clear the history cache for an item:

```
zabbix_server -c /usr/local/etc/zabbix_server.conf -R history_cache_clear=42243
```

Examples of using runtime control to gather diagnostic information:

```
### Gather all available diagnostic information in the server log file:
zabbix_server -R diaginfo

### Gather history cache statistics in the server log file:
zabbix_server -R diaginfo=historycache
```

Example of using runtime control to reload the SNMP cache:

```
zabbix_server -R snmp_cache_reload
```

Attention:
When an SNMPv3 interface is updated via the Zabbix UI, Zabbix will automatically reload the new SNMPv3 credentials for that interface in most cases; use `-R snmp_cache_reload` only if polling still fails after credential changes (for example, due to engineBoots/engineID inconsistencies or non-RFC devices), or when you need to force a global SNMP cache clear for troubleshooting.

Example of using runtime control to trigger execution of housekeeper:

```
zabbix_server -c /usr/local/etc/zabbix_server.conf -R housekeeper_execute
```

Examples of using runtime control to change log level:

```
### Increase log level of all processes:
zabbix_server -c /usr/local/etc/zabbix_server.conf -R log_level_increase

### Increase log level of second poller process:
zabbix_server -c /usr/local/etc/zabbix_server.conf -R log_level_increase=poller,2

### Increase log level of process with PID 1234:
zabbix_server -c /usr/local/etc/zabbix_server.conf -R log_level_increase=1234

### Decrease log level of all http poller processes:
zabbix_server -c /usr/local/etc/zabbix_server.conf -R log_level_decrease="http poller"
```

Example of setting the HA failover delay to the minimum of 10 seconds:

```
zabbix_server -R ha_set_failover_delay=10s
```

Process user

Zabbix server is designed to run as a non-root user. It will run as whatever non-root user it is started as. So you can run server as any non-root user without any issues.

If you will try to run it as 'root', it will switch to a hardcoded 'zabbix' user, which must be present on your system. You can only run server as 'root' if you modify the 'AllowRoot' parameter in the server configuration file accordingly.

If Zabbix server and agent are run on the same machine it is recommended to use a different user for running the server than for running the agent. Otherwise, if both are run as the same user, the agent can access the server configuration file and any Admin level user in Zabbix can quite easily retrieve, for example, the database password.

Configuration file

The default path to the server configuration file depends on the operating system and Zabbix installation method:

- For installation from sources: `/usr/local/etc/zabbix_server.conf`, or as set by compile-time variables `--sysconfdir` or `--prefix`.
- For installation from packages and appliance: `/etc/zabbix/zabbix_server.conf`.

See the configuration file options for details on configuring `zabbix_server`.

Start-up scripts

The scripts are used to automatically start/stop Zabbix processes during system's start-up/shutdown. The scripts are located under directory misc/init.d.

Server process types and threads

- `agent poller` - asynchronous poller process for passive checks with a worker thread;
- `alert manager` - alert queue manager;
- `alert syncer` - alert DB writer;
- `alerter` - process for sending notifications;
- `availability manager` - process for host availability updates;
- `browser poller` - poller for browser item checks;
- `CEP manager` - manager process for complex event processing;
- `CEP worker` - process for handling requests from the CEP manager;
- `configuration syncer` - process for managing in-memory cache of configuration data;
- `configuration syncer worker` - process for resolving and synchronizing user macro values in item names;
- `connector manager` - manager process for connectors;
- `connector worker` - process for handling requests from the connector manager;
- `discovery manager` - manager process for discovery of devices;
- `discovery worker` - process for handling discovery tasks from the discovery manager;
- `escalator` - process for escalation of actions;
- `ha manager` - process for managing high availability;
- `history poller` - process for handling calculated checks requiring a database connection;
- `history syncer` - history DB writer;
- `housekeeper` - process for removing outdated data (item history and trends, user sessions, events, etc.), as well as data left behind by deleted objects;
- `http agent poller` - asynchronous poller process for HTTP checks with a worker thread;
- `http poller` - web monitoring poller;
- `icmp pinger` - poller for icmpping checks;
- `internal poller` - poller for internal checks;
- `ipmi manager` - IPMI poller manager;
- `ipmi poller` - poller for IPMI checks;
- `java poller` - poller for Java checks;
- `lld manager` - manager process of low-level discovery tasks;
- `lld worker` - worker process of low-level discovery tasks;
- `odbc poller` - poller for ODBC checks;
- `poller` - normal poller for passive checks;
- `preprocessing manager` - manager of preprocessing tasks with preprocessing worker threads;
- `preprocessing worker` - thread for data preprocessing;
- `proxy poller` - poller for passive proxies;
- `proxy group manager` - manager of proxy load balancing and high availability;
- `report manager`- manager of scheduled report generation tasks;
- `report writer` - process for generating scheduled reports;
- `self-monitoring` - process for collecting internal server statistics;
- `service manager` - process for managing services by receiving information about problems, problem tags, and problem recovery from history syncer, task manager, and alert manager;
- `snmp poller` - asynchronous poller process for SNMP checks with a worker thread (`walk[OID]` and `get[OID]` items only);
- `snmp trapper` - trapper for SNMP traps;
- `supervisor` - coordinates the startup and shutdown of server processes and threads, and monitors their readiness;
- `task manager` - process for remote execution of tasks requested by other components (e.g., close problem, acknowledge problem, check item value now, remote command functionality);
- `telemetry query poller` - poller for telemetry query item checks;
- `timer` - timer for processing maintenances;
- `trapper` - trapper for active checks, traps, proxy communication;
- `trigger housekeeper` - process for removing problems and events generated by triggers that have since been deleted;
- `unreachable poller` - poller for unreachable devices;
- `vmware collector` - VMware data collector responsible for data gathering from VMware services.

The server log file can be used to observe these process types.

The server log file is created with read-write permissions for the file owner only. Additionally the file is readable by the owner group. All other permissions are denied.

Various types of Zabbix server processes can be monitored using the `zabbix[process,<type>,<mode>,<state>]` internal item.

History syncer transaction statistics

The history syncer process title displays detailed statistics about history syncer transactions:

```
205182 ?        S      0:00  zabbix_server: history syncer #2 [processed 0 values, 0+0 triggers in 0.00002
205183 ?        S      0:00  zabbix_server: history syncer #3 [processed 18 values, 7+0 triggers in 0.0026
205184 ?        S      0:00  zabbix_server: history syncer #4 [processed 0 values, 0+0 triggers in 0.00002
```

In "A+B triggers":

- A - triggers processed because of history values;
- B - triggers processed because of timers.

The timings, in `processed...in N (<timings>) sec`, are:

- Time spent writing item values into database;
- Time spent updating item data (state, errors, host inventory, etc);
- Time spent flushing trends to database;
- Time spent calculating triggers;
- Time spent processing events and actions.

Housekeeping procedure

The `housekeeper` process periodically removes outdated data (item history and trends, user sessions, events, etc.), as well as data left behind by deleted objects. It runs in cycles, with frequency and per-cycle deletion limit determined by `HousekeepingFrequency` and `MaxHousekeeperDelete`. Any data not removed in one cycle is carried over to the next. Automatic housekeeping can be enabled and configured in *Administration* > *Housekeeping*.

When the `housekeeper` process removes data left behind by deleted objects, it relies on tasks from the `housekeeper` table, which is populated whenever an object is deleted. For example, when you delete a host, Zabbix deletes its items, but not their history, trends, or problems. Instead, database triggers populate the `housekeeper` table with tasks consisting of these fields:

- `housekeeperid` - task ID
- `object` - object type (0 - item; 1 - trigger; 2 - service; 3 - discovered host; 4 - discovered service)
- `objectid` - object ID (helps the housekeeper find object-related data)

For example, deleting a host with two items and one trigger populates the `housekeeper` table as follows:

```
+---------------+--------+----------+
| housekeeperid | object | objectid |
+---------------+--------+----------+
|             1 |      1 |    28724 |
|             2 |      0 |    59396 |
|             3 |      0 |    59397 |
+---------------+--------+----------+
```

Database triggers populate the `housekeeper` table with tasks, without checking for object-related data; that check is done by the `housekeeper` process when it removes data:

- For deleted items (including LLD rules) - removes their data from all history and trends tables (`history`, `history_str`, `history_log`, `history_uint`, `history_text`, `history_bin`, `history_json`, `trends`, `trends_uint`). Also removes their internal events from the `problem` table.
- For deleted triggers - removes their events from event-related tables (`problem`, `event_symptom`, `events`; the database itself also removes events from the `event_recovery` table). Also notifies the `service manager` process about the removed events.

Note:
A separate `trigger housekeeper` process handles a narrower task—removing problems and events that have no known source trigger. Its run frequency is controlled by `ProblemHousekeepingFrequency`.
Until the trigger housekeeping procedure is started, problems created by triggers that have since been deleted might still generate service problems and assign them to services. If your setup involves many service status calculation rules based on frequently discovered/undiscovered triggers, consider increasing the frequency of the housekeeping procedure by adjusting the `ProblemHousekeepingFrequency` server configuration parameter.

- For deleted services - removes their problems from the `problem` table; their events from the `events` table are removed only after the configured data storage period.
- For deleted discovered hosts/services (network discovery) - removes their events from the `events` table.

The `housekeeper` process removes only those events that are not associated with problems. For example, an outdated problem/recovery event will not be removed if it's associated to an open problem. When `housekeeper` removes outdated entities, it first removes problems, then events.

Note:
Tables using `partition` mode (TimescaleDB partitioned tables), are skipped; only tables using `regular` mode are processed.

Supported platforms Due to the security requirements and mission-critical nature of server operation, UNIX is the only operating system that can consistently deliver the necessary performance, fault tolerance and resilience. Zabbix operates on market leading versions.

Zabbix server is tested on the following platforms:

- Linux
- Solaris
- AIX
- HP-UX
- Mac OS X
- FreeBSD
- OpenBSD
- NetBSD
- SCO Open Server

Note:
Zabbix may work on other Unix-like operating systems as well.

Locale Note that the server requires a UTF-8 locale so that some textual items can be interpreted correctly. Most modern Unix-like systems have a UTF-8 locale as default, however, there are some systems where that may need to be set specifically.

High availability setup for Zabbix servers

Overview

High availability (HA) is typically required in critical infrastructures that can afford virtually no downtime. So for any service that may fail there must be a failover option in place to take over should the current service fail.

Zabbix offers a **native** high-availability solution that is easy to set up and does not require any previous HA expertise. Native Zabbix HA may be useful for an extra layer of protection against software/hardware failures of Zabbix server or to have less downtime due to maintenance.

In the Zabbix high availability mode multiple Zabbix servers are run as nodes in a cluster. While one Zabbix server in the cluster is active, others are on standby, ready to take over if necessary.



Switching to Zabbix HA is non-committal. You may switch back to standalone operation at any point.

See also: Implementation details

Enabling high availability

Starting Zabbix server as cluster node

Two parameters are required in the server configuration to start a Zabbix server as cluster node:

- `HANodeName` parameter must be specified for each Zabbix server that will be an HA cluster node.

This is a unique node identifier (e.g., `zabbix-node-01`) that the server will be referred to in agent and proxy configurations. If you do not specify `HANodeName`, then the server will be started in standalone mode.

- `NodeAddress` parameter must be specified for each node.

The `NodeAddress` parameter (`address:port`) will be used by Zabbix frontend to connect to the active server node. `NodeAddress` must match the IP or FQDN name of the respective Zabbix server.

Restart all Zabbix servers after making changes to the configuration files. They will now be started as cluster nodes. The new status of the servers can be seen in *Reports* > *System information* and also by running:

```
zabbix_server -R ha_status
```

This runtime command will log the current HA cluster status into the Zabbix server log (and to stdout):

```
Failover delay: 60 seconds
Cluster status:
   #  ID                        Name                 Address                 Status      Last Access
   1. ckzxxqg7u0001lsropenyzh3m zabbix-node-01       64.227.66.193:10051     standby     0s
   2. ckzxyqo1k00013frpq539e1jp zabbix-node-02       64.227.74.25:10051      active      3s
```

Preparing frontend

Make sure that Zabbix server `address:port` is not defined in the Zabbix frontend configuration file (`zabbix.conf.php`).

```
// Uncomment and set to desired values to override Zabbix hostname/IP and port.
// $ZBX_SERVER                  = '';
// $ZBX_SERVER_PORT             = '';
```

Zabbix frontend will autodetect the active node by reading settings from the nodes table in Zabbix database. Node address of the active node will be used as the Zabbix server address.

Proxy configuration

HA cluster nodes (servers) must be listed in the configuration of either passive or active Zabbix proxy.

For a passive proxy, the node names must be listed in the Server parameter of the proxy, separated by a **comma**.

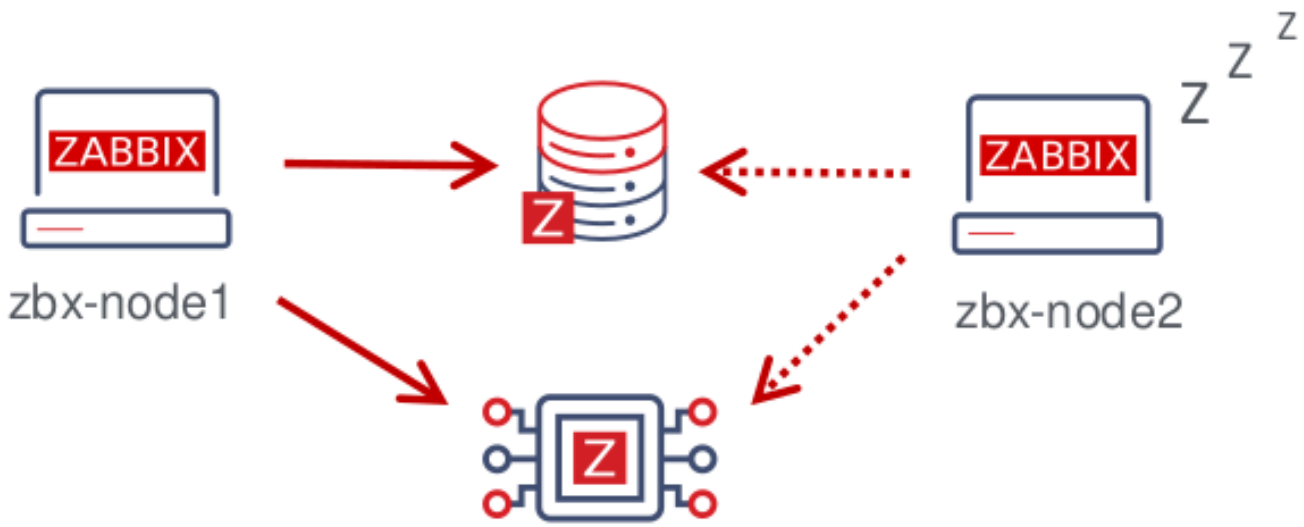
```
Server=zabbix-node-01,zabbix-node-02
```

For an active proxy, the node names must be listed in the Server parameter of the proxy, separated by a **semicolon**.

```
Server=zabbix-node-01;zabbix-node-02
```

Agent configuration

HA cluster nodes (servers) must be listed in the configuration of Zabbix agent or Zabbix agent 2.



To enable passive checks, the node names must be listed in the Server parameter, separated by a **comma**.

```
Server=zabbix-node-01,zabbix-node-02
```

To enable active checks, the node names must be listed in the ServerActive parameter. Note that for active checks the nodes must be separated by a comma from any other servers, while the nodes themselves must be separated by a **semicolon**, e.g.:

```
ServerActive=zabbix-node-01;zabbix-node-02
```

Failover to standby node

Zabbix will fail over to another node automatically if the active node stops. There must be at least one node in standby status for the failover to happen.

How fast will the failover be? All nodes update their last access time (and status, if it is changed) every 5 seconds. So:

- If the active node shuts down and manages to report its status as "stopped", another node will take over within **5 seconds**.
- If the active node shuts down/becomes unavailable without being able to update its status, standby nodes will wait for the **failover delay** + 5 seconds to take over.

The failover delay is configurable, with the supported range between 10 seconds and 15 minutes (one minute by default). To change the failover delay, you may run:

```
zabbix_server -R ha_set_failover_delay=5m
```

Managing HA cluster

The current status of the HA cluster can be managed using the dedicated runtime control options:

- `ha_status` - log HA cluster status in the Zabbix server log (and to stdout);
- `ha_remove_node=target` - remove an HA node identified by its `<target>` - name or ID of the node (name/ID can be obtained from the output of running `ha_status`), e.g.:

```
zabbix_server -R ha_remove_node=zabbix-node-02
```

Note that active/standby nodes cannot be removed.

- `ha_set_failover_delay=delay` - set HA failover delay (between 10 seconds and 15 minutes; time suffixes are supported (e.g., `30s`, `1m`)).

Node status can be monitored:

- In *Reports* > *System information*.
- In the *System information* dashboard widget.
- Using the `ha_status` runtime control option of the server (see above).

The `zabbix[cluster,discovery,nodes]` internal item can be used for node discovery, as it returns a JSON with the high-availability node information.

Disabling HA cluster

To disable a high availability cluster:

1. Make backup copies of configuration files.
2. Stop standby nodes.
3. Remove the HANodeName parameter from the active primary server.
4. Restart the primary server (it will start in standalone mode).

Upgrading HA cluster

To perform a major version upgrade for the HA nodes:

1. Stop all nodes.
2. Create a full database backup.
3. If the database uses replication make sure that all nodes are in sync and have no issues. Do not upgrade if replication is broken.
4. Select a single node that will perform database upgrade, change its configuration to standalone mode by commenting out `HANodeName` and upgrade it.
5. Make sure that database upgrade is fully completed (*System information* should display that Zabbix server is running).
6. Restart the node in HA mode.
7. Upgrade and start the rest of nodes (it is not required to change them to standalone mode as the database is already upgraded at this point).

In a minor version upgrade it is sufficient to upgrade the first node, make sure it has upgraded and running, and then start upgrade on the next node.

Implementation details

The high availability (HA) cluster is an opt-in solution and it is supported for Zabbix server. The native HA solution is designed to be simple in use, it will work across sites and does not have specific requirements for the databases that Zabbix recognizes. Users are free to use the native Zabbix HA solution, or a third-party HA solution, depending on what best suits the high availability requirements in their environment.

The solution consists of multiple `zabbix_server` instances or nodes. Every node:

- Is configured separately.
- Uses the same database.
- May have several modes: active, standby, unavailable, stopped.

Only one node can be active (working) at a time. A standby node runs only one process - the HA manager. A standby node does no data collection, processing or other regular server activities; they do not listen on ports; they have minimum database connections.

Both active and standby nodes update their last access time every 5 seconds. Each standby node monitors the last access time of the active node. If the last access time of the active node is over 'failover delay' seconds, the standby node switches itself to be the active node and assigns 'unavailable' status to the previously active node.

The active node monitors its own database connectivity - if it is lost for more than `failover delay`−5 seconds, it must stop all processing and switch to standby mode. The active node also monitors the status of the standby nodes - if the last access time of a standby node is over 'failover delay' seconds, the standby node is assigned the `unavailable` status.

The nodes are designed to be compatible across minor Zabbix versions.

Configuration parameters: Zabbix server

Overview

Configuration parameters allow customizing Zabbix server behavior.

Specify parameter values in the Zabbix server configuration file (`zabbix_server.conf`). All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
AlertScriptsPath	Specify the alert script location.
AllowRoot	Allow server to run as `root`.
AllowSoftwareUpdateCheck	Allow receiving information about software updates from `zabbix.com`.
AllowUnsupportedDBVersions	Allow server to work with unsupported database versions.
BridgeAdapterURL	Set the Zabbix bridge adapter HTTP(S) URL.
CacheSize	Set the configuration cache size (in bytes).
CacheUpdateFrequency	Set the configuration cache update frequency (in seconds).
DBHost	Specify the database host (or socket directory).
DBName	Specify the database name. **Mandatory**.
DBPassword	Specify the database password.
DBPort	Specify the port of the database host.
DBSchema	Specify the database schema name. Used for PostgreSQL.
DBSocket	Specify path to the MySQL socket file.
DBUser	Specify the database user.
DBTLSConnect	Specify a value to enforce a TLS connection to the database.
DBTLSCAFile	Specify full pathname of the file containing top-level CA(s) certificates for database-certificate verification.
DBTLSCertFile	Specify full pathname of the file containing the Zabbix server certificate for authenticating to database.
DBTLSKeyFile	Specify full pathname of the file containing the private key for authenticating to database.
DBTLSCipher	Specify the list of encryption ciphers that Zabbix server permits for TLS protocols up through TLS v1.2. Supported only for MySQL.
DBTLSCipher13	Specify the list of encryption ciphersuites that Zabbix server permits for the TLS v1.3 protocol. Supported only for MySQL, starting from version 8.0.16.
DebugLevel	Set the debug level.
EnableGlobalScripts	Enable global scripts on Zabbix server.
EnableMobileDevices	Enable mobile device support on Zabbix server.
ExportDir	Specify a directory to enable real-time export of events, history and trends.

Parameter	Description
ExportFileSize	Set the maximum export file size (in bytes).
ExportType	Specify entity types for real-time export (events, history, trends).
ExternalScripts	Specify the external script location.
Fping6Location	Specify the fping6 location.
FpingLocation	Specify the fping location.
FrontendAllowedIP	Specify IP addresses or CIDR ranges that are permitted to connect from the frontend.
HANodeName	Specify the node name in a high-availability cluster.
HistoryCacheSize	Set the history cache size.
HistoryIndexCacheSize	Set the history-index cache size.
HistoryProvider	Set the history storage provider (Elasticsearch, ClickHouse, or SQL).
HistoryStorageDateIndex	This parameter is **deprecated** (use HistoryProvider instead). Enable storing history values in different indices based on date.
HistoryStorageURL	This parameter is **deprecated** (use HistoryProvider instead). Specify the history storage URL.
HistoryStorageTypes	This parameter is **deprecated** (use HistoryProvider instead). Specify value types to be sent to history storage.
HousekeepingFrequency	Set the housekeeping procedure frequency (in hours).
Include	Specify a directory or individual files to be included in the configuration file.
JavaGateway	Specify the Zabbix Java gateway IP address (or hostname).
JavaGatewayPort	Specify the Zabbix Java gateway port.
ListenBacklog	Set the maximum number of pending connections in the TCP queue.
ListenIP	Specify the trapper listen IP address.
ListenPort	Specify the trapper listen port.
LoadModule	Specify the module to load at server startup.
LoadModulePath	Specify full path to the location of server modules.
LogFile	Specify the server log file.
LogFileSize	Set the maximum log file size.
LogSlowQueries	Set the maximum duration a database query may take before being logged (in milliseconds).
LogType	Specify the log output type.
MaxConcurrentChecksPerPoller	Set the maximum number of asynchronous checks that can be executed at once by each HTTP agent poller, agent poller, or SNMP poller.
MaxHousekeeperDelete	Set the maximum number of rows to be deleted per task in one housekeeping cycle.
NodeAddress	Specify IP or hostname with optional port to override how the frontend connects to the server.
PidFile	Specify the PID file.
ProblemHousekeepingFrequency	Set the frequency of deleting problems for deleted triggers.
ProxyConfigFrequency	Set the frequency of sending configuration data to Zabbix proxy by the server.
ProxyDataFrequency	Set the frequency of requesting history data from Zabbix proxy by the server.
ServiceManagerSyncFrequency	Set the frequency of synchronizing the configuration of a service manager.
SMSDevices	Specify modem files that are allowed to use Zabbix server.
SNMPTrapperFile	Specify the temporary file used for passing data from the SNMP trap daemon to the server.
SocketDir	Specify the directory for storing IPC sockets used by internal Zabbix services.
SourceIP	Specify the source IP address.
SSHKeyLocation	Specify the location of public and private keys for SSH checks and actions.
SSLCertLocation	Specify the location of SSL-client certificate files for client authentication.
SSLKeyLocation	Specify the location of SSL private key files for client authentication.
SSLCALocation	Specify the location of certificate authority (CA) files for SSL-server certificate verification.
StartAgentPollers	Set the number of pre-forked instances of asynchronous Zabbix agent pollers.
StartAlerters	Set the number of pre-forked instances of alerters.
StartBrowserPollers	Set the number of pre-forked instances of browser item pollers.
StartConnectors	Set the number of pre-forked instances of connector workers.
StartDBSyncers	Set the number of pre-forked instances of history syncers.
StartDiscoverers	Set the number of pre-forked instances of discovery workers.
StartEscalators	Set the number of pre-forked instances of escalators.
StartHistoryPollers	Set the number of pre-forked instances of history pollers.
StartHTTPAgentPollers	Set the number of pre-forked instances of asynchronous HTTP agent pollers.
StartHTTPPollers	Set the number of pre-forked instances of HTTP pollers.
StartIPMIPollers	Set the number of pre-forked instances of IPMI pollers.
StartJavaPollers	Set the number of pre-forked instances of Java pollers.
StartLLDProcessors	Set the number of pre-forked instances of low-level discovery (LLD) workers.
StartODBCPollers	Set the number of pre-forked instances of ODBC pollers.
StartPingers	Set the number of pre-forked instances of ICMP pingers.

Parameter	Description
StartPollersUnreachable	Set the number of pre-forked instances of pollers for unreachable hosts (including IPMI and Java).
StartPollers	Set the number of pre-forked instances of pollers.
StartPreprocessors	Set the number of pre-started instances of preprocessing workers.
StartProxyPollers	Set the number of pre-forked instances of pollers for passive proxies.
StartReportWriters	Set the number of pre-forked instances of report writers.
StartSNMPPollers	Set the number of pre-forked instances of asynchronous SNMP pollers.
StartSNMPTrapper	Set to "1" to start an SNMP trapper process.
StartTelemetryQueryPollers	Set the number of pre-forked instances of asynchronous Telemetry query pollers.
StartTimers	Set the number of pre-forked instances of timers.
StartTrappers	Set the number of pre-forked instances of trappers.
StartVMwareCollectors	Set the number of pre-forked VMware collector instances.
StatsAllowedIP	Specify IP addresses or DNS names of external Zabbix instances that are allowed to request stats.
TelemetryProvider	Set the telemetry data source.
Timeout	Specify how long to wait (in seconds) for establishing connection and exchanging data with Zabbix proxy, agent, web service, and for SNMP checks (except SNMP `walk[OID]` and `get[OID]` items).
TLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.
TLSCertFile	Specify full pathname of the file containing the server certificate or certificate chain, used for encrypted communications between Zabbix components.
TLSCipherAll	Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.
TLSCipherAll13	Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.
TLSCipherCert	Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for certificate-based encryption.
TLSCipherCert13	Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for certificate-based encryption.
TLSCipherPSK	Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for PSK-based encryption.
TLSCipherPSK13	Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for PSK-based encryption.
TLSCRLFile	Specify full pathname of the file containing revoked certificates. This parameter is used for encrypted communications between Zabbix components.
TLSFrontendAccept	Specify encryption level for incoming connections from the frontend.
TLSFrontendCertIssuer	Specify the allowed frontend certificate issuer.
TLSFrontendCertSubject	Specify the allowed frontend certificate subject.
TLSKeyFile	Specify full pathname of the file containing the server private key, used for encrypted communications between Zabbix components.
TLSListen	Require TLS-only connections on the trapper socket.
TmpDir	Specify the temporary directory.
TrapperTimeout	Specify trapper timeout in seconds for: - retrieval of historical data from Zabbix proxy; - sending configuration data to Zabbix proxy; - global script execution or remote command execution on Zabbix server.
TrendCacheSize	Set the trend cache size.
TrendFunctionCacheSize	Set the trend-function cache size.
UnavailableDelay	Set the frequency of checking host availability during the unavailability period.
UnreachableDelay	Set the frequency of checking host availability during the unreachability period.
UnreachablePeriod	Set after how many seconds of unreachability treat host as unavailable.
User	Drop privileges to a specific, existing user on the system.
ValueCacheSize	Set the size of the history value cache.
Vault	Specify the vault provider.
VaultAppRoleID	Specify the HashiCorp vault AppRole ID.
VaultAppSecretID	Specify the HashiCorp vault AppRole secret ID.
VaultDBPath	Specify the location for retrieving database credentials by keys.
VaultPrefix	Specify custom prefix for the vault path or query.
VaultTLSCertFile	Specify the name of the SSL certificate file used for client authentication.
VaultTLSKeyFile	Specify the name of the SSL private key file used for client authentication.
VaultToken	Specify the HashiCorp vault authentication token.
VaultURL	Specify the vault server URL.

Parameter	Description
VMwareCacheSize	Set the shared memory size for storing VMware data.
VMwareFrequency	Set the frequency of data gathering from a single VMware service in seconds.
VMwarePerfFrequency	Set the frequency of retrieving performance counter statistics from a single VMware service in seconds.
VMwareTimeout	Set the maximum number of seconds a `vmware collector` will wait for a response from VMware service.
WebDriverURL	Specify the WebDriver interface URL.
WebServiceURL	Specify the web service URL in the format `http[s]://host:port/report`.

Note that:

- The default values reflect daemon defaults, not the values in the shipped configuration files.
- Values support environment variables.
- Zabbix supports configuration files only in UTF-8 encoding without BOM.
- Comments starting with # are only supported in the beginning of the line.

Warning:
Too many processes can exhaust the per-process file descriptor limit for the preprocessing manager.

Such processes are data gathering by pollers, unreachable pollers, ODBC pollers, HTTP pollers, Java pollers, pingers, trappers, proxypollers together with IPMI manager, SNMP trapper, preprocessing workers, and discovery workers. Exhausting the file descriptor limit will cause Zabbix server to **stop**, typically shortly after startup but sometimes taking longer.

To avoid such issues, review the Zabbix server configuration file to optimize the number of concurrent checks and processes. Additionally, if necessary, ensure that the file descriptor limit is set sufficiently high by checking and adjusting system limits.

Parameter details

AlertScriptsPath

Specify the alert script location. Depends on the `datadir` compile-time installation variable.

Default: `/usr/local/share/zabbix/alertscripts`

AllowRoot

Allow server to run as `root`. If disabled and the server is started by `root`, the server will try to switch to the `zabbix` user instead. Has no effect if started under a regular user.

Default: 0
 Values: 0 - do not allow; 1 - allow

AllowSoftwareUpdateCheck

Allow Zabbix UI receiving information about software updates from `zabbix.com`.

Default: 1
 Values: 0 - do not allow; 1 - allow

AllowUnsupportedDBVersions

Allow server to work with unsupported database versions.

Default: 0
 Values: 0 - do not allow; 1 - allow

BridgeAdapterURL

Set the Zabbix bridge adapter HTTP(S) URL.

Example:

```
BridgeAdapterURL=https://bridge-adapter.example.com:10054/rpc
```

CacheSize

Set the configuration cache size (in bytes). The shared memory size for storing host, item and trigger data.

Default: 32M
 Range: 128K-64G

CacheUpdateFrequency

Set the configuration cache update frequency (in seconds). See also runtime control options.

Default: 10
 Range: 1-3600

DBHost

Specify the database host (or socket directory).

For MySQL:

- `localhost` or an empty string uses the default UNIX-domain socket.

For PostgreSQL:

- `localhost` is resolved via DNS (typically to 127.0.0.1).
- An empty string uses the default UNIX-domain socket.
- A path (for example, `/var/run/pgbouncer`) results in using the UNIX socket at that path.
- A comma-separated list can contain multiple `host:port` values. Each host is tried in order until a read-write connection is established.

Example: `DBHost=localhost:5431,127.0.0.1:20051,zabbix.domain,[::1]:30051,[12fc::1]`

Default: `localhost`

DBName

Specify the database name.

Mandatory: **Yes**

DBPassword

Specify the database password. Comment this line if no password is used.

DBPort

Specify the port of the Zabbix database host (see DBHost).
 `DBSocket` and `DBPort` are mutually exclusive in server configuration. Specify only one, or leave both undefined.

Default for MySQL: 3306

Default for PostgreSQL: 5432

Range: 1024-65535

DBSchema

Specify the database schema name. Used for PostgreSQL.

DBSocket

Specify the path to the MySQL socket file.
 `DBSocket` and `DBPort` are mutually exclusive in server configuration. Specify only one, or leave both undefined.

DBUser

Specify the database user.

DBTLSConnect

Specify a value to enforce a TLS connection to the database: - *required* - connect using TLS; - *verify_ca* - connect using TLS and verify certificate; - *verify_full* - connect using TLS, verify certificate and verify that database identity specified by `DBHost` matches its certificate.

With MySQL, starting from 5.7.11, and PostgreSQL the following values are supported: `required`, `verify_ca`, `verify_full`.
 With MariaDB, starting from version 10.2.6, the `required` and `verify_full` values are supported.
 By default, not set to any option and the behavior depends on database configuration.

DBTLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificates for database certificate verification.

Mandatory: no (yes, if DBTLSConnect set to `verify_ca` or `verify_full`)

DBTLSCertFile

Specify full pathname of the file containing the Zabbix server certificate for authenticating to database.

DBTLSKeyFile

Specify full pathname of the file containing the private key for authenticating to database.

DBTLSCipher

Specify the list of encryption ciphers that Zabbix server permits for TLS protocols up through TLS v1.2. Supported only for MySQL.

DBTLSCipher13

Specify the list of encryption ciphersuites that Zabbix server permits for the TLS v1.3 protocol. Supported only for MySQL, starting from version 8.0.16.

DebugLevel

Set the debug level: - *0* - basic information about starting and stopping of Zabbix processes; - *1* - critical information; - *2* - error information; - *3* - warnings; - *4* - for debugging (produces lots of information); - *5* - extended debugging (produces even more information).
 See also runtime control options.

Default: 3
 Range: 0-5

EnableGlobalScripts

Enable global scripts on Zabbix server.
 Note: Global script execution is enabled by default. However, for new installations starting with Zabbix 7.0, EnableGlobalScripts is explicitly set to 0 (disabled).

Default: 1
 Values: 0 - disable; 1 - enable

EnableMobileDevices

Enable mobile device support on Zabbix server.

When enabled, Zabbix server will process requests to add or remove mobile devices, and will send push notifications to those devices.

Default: 1
 Values: 0 - disable; 1 - enable

ExportDir

Specify a directory to enable real-time export of events, history and trends.

ExportFileSize

Set the maximum export file size (in bytes). Used for rotation if `ExportDir` is set.

Default: 1G
 Range: 1M-1G

ExportType

Specify entity types for real-time export (events, history, trends). A comma-delimited list. All types are exported by default. Valid only if `ExportDir` is set.
 Note that if `ExportType` is specified, but `ExportDir` is not, then this is a configuration error and the server will not start.

Example for history and trends export:

```
ExportType=history,trends
```

Example for event export only:

```
ExportType=events
```

ExternalScripts

Specify the external script location. Depends on the `datadir` compile-time installation variable.

Default: `/usr/local/share/zabbix/externalscripts`

Fping6Location

Specify the fping6 location. Make sure that the fping6 binary has root ownership and the SUID flag set. Make empty (`Fping6Location=`) if your fping utility is capable to process IPv6 addresses.

Default: `/usr/sbin/fping6`

FpingLocation

Specify the fping location. Make sure that the fping binary has root ownership and the SUID flag set.

Default: `/usr/sbin/fping`

FrontendAllowedIP

Specify IP addresses or CIDR ranges that are permitted to connect from the frontend. A comma-separated list. By default, all connections are accepted for frontend requests. If IPv6 support is enabled then `127.0.0.1`, `::127.0.0.1`, `::ffff:127.0.0.1` are treated equally and `::/0` will allow any IPv4 or IPv6 address. `0.0.0.0/0` can be used to allow any IPv4 address.

Example:

```
FrontendAllowedIP=127.0.0.1,192.168.1.0/24,::1,2001:db8::/32,zabbix.example.com
```

HANodeName

Specify the node name in a high-availability cluster. When empty the server is working in standalone mode and a node with empty name is created.

HistoryCacheSize

Set the history cache size (in bytes). The shared memory size for storing history data.

Default: 16M
 Range: 128K-16G

HistoryIndexCacheSize

Set the history-index cache size (in bytes). The shared memory size for indexing the history data stored in history cache. The index cache size needs roughly 100 bytes to cache one item.

Default: 4M
 Range: 128K-16G

HistoryProvider

Set the history storage provider using the format `<provider>[;<options>]`, where:

- `<provider>` - one of the supported history storage providers: `sql` (default), `elasticsearch`, or `clickhouse`.
- `<options>` - comma-delimited `key=value` pairs. Values can be quoted.

Options for all history storage providers:

- `value_types` - set the value types to store (`uint`, `dbl`, `str`, `log`, `text`, `json`). Each can be used once, and only for a single provider.
- `precache` - set 1 (default) to preload data from the history storage provider into the value cache in bulk. This can significantly improve value cache warmup after restart.

Options for Elasticsearch and ClickHouse:

- `url` - set the HTTP or HTTPS URL used to connect to the history storage provider.
- `log_slow_queries` - set the maximum duration (in milliseconds) a history storage provider query may take before being logged. This option becomes enabled starting with `DebugLevel=3`.
- `source_ip` - set the source IP address used for outgoing connections to the history storage provider.

Options for Elasticsearch:

- `date_index` - set 1 to enable storing history values in different indices based on date.

Options for ClickHouse:

- `db` - set the name of the database.
- `username` - set the username used to log in to the database.
- `password` - set the password used to log in to the database.
- `ssl_cert_file`, `ssl_key_file`, `ssl_key_password`, `ssl_verify_peer`, `ssl_verify_host`, `ssl_ca_location`, `ssl_cert_location`, `ssl_key_location` - set these options to configure a secure (TLS/SSL) connection to the database.

Example for ClickHouse:

```
HistoryProvider=clickhouse;value_types="uint,dbl,str,log,text,json",url=http://localhost:8123,db=zabbix,us
```

Example for Elasticsearch:

```
HistoryProvider=elasticsearch;value_types="uint,dbl,str,log,text,json",url=http://localhost:9200,date_inde
```

Example for ClickHouse and Elasticsearch:

```
HistoryProvider=clickhouse;value_types="uint,dbl,str",url=http://localhost:8123,db=zabbix,username=zabbix,
HistoryProvider=elasticsearch;value_types="log,text,json",url=http://localhost:9200,date_index=0
```

HistoryStorageDateIndex

This parameter is deprecated (use HistoryProvider instead).
 Enable storing history values in different indices based on date. This parameter is used for Elasticsearch setup.

Default: 0
 Values: 0 - disable; 1 - enable

HistoryStorageURL

This parameter is deprecated (use HistoryProvider instead).
 Specify the history storage URL. This parameter is used for Elasticsearch setup.

HistoryStorageTypes

This parameter is deprecated (use HistoryProvider instead).
 Specify value types to be sent to history storage. A comma-delimited list. This parameter is used for Elasticsearch setup.

Default: `uint,dbl,str,log,text,json`

HousekeepingFrequency

Set the housekeeping procedure frequency (in hours). Housekeeping is removing outdated information from the database.
 Note: To prevent housekeeper from being overloaded (for example, when history and trend periods are greatly reduced), no more than 4 times `HousekeepingFrequency` hours of outdated information are deleted in one housekeeping cycle, for each item. Thus, if `HousekeepingFrequency` is 1, no more than 4 hours of outdated information (starting from the oldest entry) will be deleted per cycle.
 Note: To lower load on server startup housekeeping is postponed for 30 minutes after server start. Thus, if `HousekeepingFrequency` is 1, the very first housekeeping procedure after server start will run after 30 minutes, and will repeat with one hour delay thereafter.
 It is possible to disable automatic housekeeping by setting `HousekeepingFrequency` to 0. In this case the housekeeping procedure can only be started by `housekeeper_execute` runtime control option and the period of outdated information deleted in one housekeeping cycle is 4 times the period since the last housekeeping cycle, but not less than 4 hours and not greater than 4 days.
 See also runtime control options and details on the housekeeping procedure.

Default: 1
 Range: 0-24

Include

Specify a directory or individual files to be included in the configuration file.
 During Zabbix installation, this parameter is set to `/usr/local/etc`, unless you change this during compile time.

General rules for the `Include` parameter:

- The path can be relative to the `zabbix_server.conf` file location.
- You can specify multiple `Include` parameters in one configuration file.
- Zabbix processes the `Include` parameter as soon as it reads it, so consider placing it at the end of the configuration file.

Rules for included files:

- Each file must be readable.
- Each file must have correct syntax. Otherwise, the component will not start.

Rules for including directories or file patterns:

- Do not assume any order of inclusion (e.g., files are not included in alphabetical order).
- Do not define the same parameter in more than one included file (e.g., to override a general setting with a specific one), since you cannot predict which value will be used.
- Watch out for backup copies that some text editors create automatically (e.g., editing `include/my_specific.conf` may create `include/my_specific.conf.BAK`), since both files will be included.

To include only relevant files from a directory, use the wildcard (`*`) for pattern matching.

Example:

```
Include=/absolute/path/to/config/files/*.conf
```

JavaGateway

Specify the IP address or hostname of Zabbix Java gateway. Only required if Java pollers are started.

JavaGatewayPort

Specify the Zabbix Java gateway port.

Default: 10052
 Range: 1024-32767

ListenBacklog

Set the maximum number of pending connections in the TCP queue.
 The default value is a hard-coded constant, which depends on the system.
 The maximum supported value depends on the system, too high values may be silently truncated to the 'implementation-specified maximum'.

Default: `SOMAXCONN`
 Range: 0 - INT_MAX

ListenIP

Specify the trapper listen IP address. A comma-delimited list.
 Trapper will listen on all network interfaces if this parameter is not specified.

Default: `0.0.0.0`

ListenPort

Specify the trapper listen port.

Default: 10051
 Range: 1024-32767

LoadModule

Specify the module to load at server startup. Modules are used to extend the functionality of the server. The module must be located in the directory specified by `LoadModulePath` or the path must precede the module name. If the preceding path is absolute (starts with `/`) then `LoadModulePath` is ignored.
 Formats:

- `LoadModule=<module.so>`
- `LoadModule=<path/module.so>`
- `LoadModule=</abs_path/module.so>`

Multiple `LoadModule` parameters can be specified.

LoadModulePath

Specify full path to the location of server modules. The default depends on compilation options.

LogFile

Specify the server log file.

Mandatory: Yes, if `LogType` is set to `file`; otherwise, no

LogFileSize

Set the maximum log file size (in MB).
0 - disable automatic log rotation.
 Note: If the log file size limit is reached and file rotation fails, for whatever reason, the existing log file is truncated and started anew.

Default: 1
 Range: 0-1024
 Mandatory: Yes, if `LogType` is set to `file`; otherwise, no

LogSlowQueries

Set the maximum duration a database query may take before being logged (in milliseconds).
 0 - don't log slow queries.
 This option becomes enabled starting with DebugLevel=3.

Default: 0
 Range: 0-3600000

LogType

Specify the log output type: - *file* - write log to file specified by `LogFile` parameter; - *system* - write log to syslog; - *console* - write log to standard output.

Default: `file`

MaxConcurrentChecksPerPoller

Set the maximum number of asynchronous checks that can be executed at once by each HTTP agent poller, agent poller or SNMP poller. See StartHTTPAgentPollers, StartAgentPollers, and StartSNMPPollers.

Default: 1000
 Range: 1-1000

MaxHousekeeperDelete

Set the maximum number of rows to be deleted per task in one housekeeping cycle from `history`, `history_str`, `history_log`, `history_uint`, `history_text`, `history_bin`, `history_json`, `trends`, `trends_uint` and `problem` tables.
 Housekeeping tasks exist in the `housekeeper` table in the format: `[housekeeperid]`, `[object]`, `[objectid]`.
 This parameter applies only to deleting data left behind by already deleted items.
 If set to 0 then no limit is used at all.
 In this case you must know what you are doing, so as not to overload the database.
 When a lot of items are deleted it increases the load to the database, because the housekeeper will need to remove all the history data that these items had.
 For example, if we only have to remove 1 item prototype from the template, but this template is linked to 50 hosts and for every host the prototype is expanded to 100 real items, 5000 items in total have to be removed (1*50*100).
 If 500 is set for `MaxHousekeeperDelete` (`MaxHousekeeperDelete=500`), the housekeeper process will have to remove up to 2500000 values (5000*500) for the deleted items from history and trends tables in one cycle.

See also details on the housekeeping procedure.

Default: 5000
 Range: 0-1000000

NodeAddress

Specify IP or hostname with optional port to override how the frontend connects to the server.
 Format: <address>[:<port>]

 If IP or hostname is not set, the value of ListenIP will be used. If ListenIP is not set, the value localhost will be used.
 If port is not set, the value of ListenPort will be used. If ListenPort is not set, the value 10051 will be used.

 This option can be overridden by the address specified in the frontend configuration.

 See also: HANodeName parameter; Enabling high availability.

Default: localhost:10051

PidFile

Specify the PID file.

Default: /tmp/zabbix_server.pid

ProblemHousekeepingFrequency

Set the frequency of deleting problems for deleted triggers (in seconds).
 See also details on the housekeeping procedure.

Default: 60
 Range: 1-3600

ProxyConfigFrequency

Set the frequency of sending configuration data to Zabbix proxy by the server. Used only for proxies in the passive mode.

Default: 10
 Range: 1-604800

ProxyDataFrequency

Set the frequency of requesting history data from Zabbix proxy by the server. Used only for proxies in the passive mode.

Default: 1
 Range: 1-3600

ServiceManagerSyncFrequency

Set the frequency of synchronizing the configuration of a service manager (in seconds).

Default: 60
 Range: 1-3600

SMSDevices

Specify modem files that are allowed to use Zabbix server. A comma-delimited list.
 SMS sending is not possible if this parameter is not set.

Example:

```
SMSDevices=/dev/ttyUSB0,/dev/ttyUSB1
```

SNMPTrapperFile

Specify the temporary file used for passing data from the SNMP trap daemon to the server.
 Must be the same as in zabbix_trap_receiver.pl or SNMPTT configuration file.

Default: /tmp/zabbix_traps.tmp

SocketDir

Specify the directory for storing IPC sockets used by internal Zabbix services.

Default: /tmp

SourceIP

Specify the source IP address for:

- outgoing connections to Zabbix proxy and Zabbix agent
- agentless connections (VMware, SSH, JMX, SNMP, Telnet and simple checks)
- HTTP agent connections
- script item JavaScript HTTP requests
- preprocessing JavaScript HTTP requests
- sending notification emails (connections to SMTP server)
- webhook notifications (JavaScript HTTP connections)
- connections to the Vault

SSHKeyLocation

Specify the location of public and private keys for SSH checks and actions.

SSLCertLocation

Specify the location of SSL-client certificate files for client authentication.
 This parameter is used in web monitoring, for communication with Vault, for history storage provider ClickHouse, and for telemetry data source ClickHouse.

SSLKeyLocation

Specify the location of SSL private key files for client authentication.
 This parameter is used in web monitoring, for communication with Vault, for history storage provider ClickHouse, and for telemetry data source ClickHouse.

SSLCALocation

Specify the location of certificate authority (CA) files for SSL-server certificate verification. If not set, system-wide directory will be used.
 Note that the value of this parameter will be set as libcurl option `CURLOPT_CAPATH`. For libcurl versions before 7.42.0, this only has effect if libcurl was compiled to use OpenSSL. For more information, see cURL web page.
 This parameter is used in web monitoring, SMTP authentication, HTTP agent items, for communication with Vault, for history storage provider ClickHouse, and for telemetry data source ClickHouse.

StartAgentPollers

Set the number of pre-forked instances of Zabbix agent pollers. See MaxConcurrentChecksPerPoller.

Default: 1
 Range: 0-1000

StartAlerters

Set the number of pre-forked instances of alerters.

Default: 3
 Range: 1-100

StartBrowserPollers

Set the number of pre-forked instances of browser item pollers.

Default: 1
 Range: 0-1000

StartConnectors

Set the number of pre-forked instances of connector workers. The connector manager process is started automatically when a connector worker is started.

Default: 0
 Range: 0-1000

StartDBSyncers

Set the number of pre-forked instances of history syncers.
 Note: Be careful when changing this value, increasing it may do more harm than good. Roughly, the default value should be enough to handle up to 4000 NVPS.

Default: 4
 Range: 1-100

StartDiscoverers

Set the number of pre-forked instances of discovery workers[1].

Default: 5
 Range: 0-1000

StartEscalators

Set the number of pre-forked instances of escalators.

Default: 1
 Range: 1-100

StartHistoryPollers

Set the number of pre-forked instances of history pollers.
 Only required for calculated checks.

Default: 5
 Range: 0-1000

StartHTTPAgentPollers

Set the number of pre-forked instances of HTTP agent pollers. See MaxConcurrentChecksPerPoller.

Default: 1
 Range: 0-1000

StartHTTPPollers

Set the number of pre-forked instances of HTTP pollers[1].

Default: 1
 Range: 0-1000

StartIPMIPollers

Set the number of pre-forked instances of IPMI pollers.

Default: 0
 Range: 0-1000

StartJavaPollers

Set the number of pre-forked instances of Java pollers[1].

Default: 0
 Range: 0-1000

StartLLDProcessors

Set the number of pre-forked instances of low-level discovery (LLD) workers[1].
 The LLD manager process is automatically started when an LLD worker is started.

Default: 2
 Range: 1-100

StartODBCPollers

Set the number of pre-forked instances of ODBC pollers[1].

Default: 1
 Range: 0-1000

StartPingers

Set the number of pre-forked instances of ICMP pingers[1].

Default: 1
 Range: 0-1000

StartPollersUnreachable

Set the number of pre-forked instances of pollers for unreachable hosts (including IPMI and Java)[1].
 At least one poller for unreachable hosts must be running if regular, IPMI or Java pollers are started.

Default: 1
 Range: 0-1000

StartPollers

Set the number of pre-forked instances of pollers[1].

Default: 5
 Range: 0-1000

StartPreprocessors

Set the number of pre-started instances of preprocessing worker[1] threads should be set to no less than the available CPU core count. More workers should be set if preprocessing is not CPU-bound and has lots of network requests.

Default: 16
 Range: 1-1000

StartProxyPollers

Set the number of pre-forked instances of pollers for passive proxies[1].

Default: 1
 Range: 0-250

StartReportWriters

Set the number of pre-forked instances of report writers.
 If set to 0, scheduled report generation is disabled.
 The report manager process is automatically started when a report writer is started.

Default: 0
 Range: 0-100

StartSNMPPollers

Set the number of pre-forked instances of SNMP pollers. See `MaxConcurrentChecksPerPoller`.

Default: 1
 Range: 0-1000

StartSNMPTrapper

Set to 1 to start an SNMP trapper process.

Default: 0
 Range: 0-1

StartTelemetryQueryPollers

Set the number of pre-forked instances of asynchronous Telemetry query pollers. See MaxConcurrentChecksPerPoller.

Default: 1
 Range: 0-1000

StartTimers

Set the number of pre-forked instances of timers.
 Timers process maintenance periods.

Default: 1
 Range: 1-1000

StartTrappers

Set the number of pre-forked instances of trappers[1].
 Trappers accept incoming connections from Zabbix sender, active agents and active proxies.

Default: 5
 Range: 0-1000

StartVMwareCollectors

Set the number of pre-forked VMware collector instances.

Default: 0
 Range: 0-250

StatsAllowedIP

Specify IP addresses (optionally in CIDR notation) or DNS names of external Zabbix instances that are allowed to request stats. A comma-delimited list.
 If this parameter is not set no stats requests will be accepted.
 If IPv6 support is enabled then `127.0.0.1`, `::127.0.0.1`, `::ffff:127.0.0.1` are treated equally and `::/0` will allow any IPv4 or IPv6 address. `0.0.0.0/0` can be used to allow any IPv4 address.

Example:

```
StatsAllowedIP=127.0.0.1,192.168.1.0/24,::1,2001:db8::/32,zabbix.example.com
```

TelemetryProvider

Set the telemetry data source using the format `<provider>[;<options>]` where:

- `<provider>` - one of the supported telemetry data sources: `clickhouse`
- `<options>` - comma-delimited `key=value` pairs. Values can be quoted.

Options for ClickHouse:

- `url` - set the HTTP or HTTPS URL used to connect to the telemetry data source.
- `db` - set the name of the database.
- `username` - set the username used to log in to the database.
- `password` - set the password used to log in to the database.
- `ssl_cert_file`, `ssl_key_file`, `ssl_key_password`, `ssl_verify_peer`, `ssl_verify_host`, `ssl_ca_location`, `ssl_cert_location`, `ssl_key_location` - set these options to configure a secure (TLS/SSL) connection to the database.
- `source_ip` - set the source IP address used for outgoing connections to the telemetry data source.
- `vault_path` - set the path to a Vault secret holding the `username` and `password` values. Do not use together with the `username` and `password` options.

Example for ClickHouse:

```
TelemetryProvider=clickhouse;url=http://localhost:8123,db=zabbix,username=zabbix,password=zabbix
```

Example for ClickHouse with Vault:

```
TelemetryProvider=clickhouse;url=http://localhost:8123,db=zabbix,vault_path=clickhouse
```

Timeout

Specify how long to wait (in seconds) for establishing connection and exchanging data with Zabbix proxy, agent, web service, and legacy SNMP checks (single OID number or string).

This parameter defines the duration for various communication operations:

- remote command execution on Zabbix agent
- SSH/Telnet command execution
- requests to Zabbix web service
- communication timeout for mediatype test request and the `history.push` method
- rescheduling of items when an IPMI interface becomes unavailable
- sending response to Zabbix proxy when failing to exchange data due to rights or encryption issues
- deadline for IPC asynchronous sockets and runtime control options
- JMX connections
- fetching statistics from a remote Zabbix proxy or server
- sending responses to Zabbix frontend
- asynchronous poller DNS requests
- response for active check heartbeat

- retrieval of Zabbix agent data (values) from active agents
- retrieval of data from Zabbix sender
- response when trapper fails to process a request
- sending active check list to Zabbix agent

This timeout will **not** be used for those checks that have flexible timeout settings configured in the frontend (on global, proxy, or per-item level). For example, SNMP `walk[OID]` and `get[OID]` items use the configured timeout from the frontend; legacy SNMP checks still use the server timeout value.

Default: 3
 Range: 1-30

TLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.

TLSCertFile

Specify full pathname of the file containing the server certificate or certificate chain, used for encrypted communications between Zabbix components.

TLSCipherAll

Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.

Example:

```
TLS_AES_256_GCM_SHA384:TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256
```

TLSCipherAll13

Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.

Example for GnuTLS:

```
NONE:+VERS-TLS1.2:+ECDHE-RSA:+RSA:+ECDHE-PSK:+PSK:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+SHA1:+CURVE-ALL
```

Example for OpenSSL:

```
EECDH+aRSA+AES128:RSA+aRSA+AES128:kECDHEPSK+AES128:kPSK+AES128
```

TLSCipherCert

Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for certificate-based encryption.

Example for GnuTLS:

```
NONE:+VERS-TLS1.2:+ECDHE-RSA:+RSA:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+SHA1:+CURVE-ALL:+COMP-NULL:+SIG
```

Example for OpenSSL:

```
EECDH+aRSA+AES128:RSA+aRSA+AES128
```

TLSCipherCert13

Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for certificate-based encryption.

TLSCipherPSK

Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for PSK-based encryption.

Example for GnuTLS:

```
NONE:+VERS-TLS1.2:+ECDHE-PSK:+PSK:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+SHA1:+CURVE-ALL:+COMP-NULL:+SIG
```

Example for OpenSSL:

```
kECDHEPSK+AES128:kPSK+AES128
```

TLSCipherPSK13

Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for PSK-based encryption.

Example:

```
TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256
```

TLSCRLFile

Specify full pathname of the file containing revoked certificates. This parameter is used for encrypted communications between Zabbix components.

TLSFrontendAccept

Specify encryption level for incoming connections from the frontend.

Multiple values can be specified, separated by comma:

- unencrypted - accept connections without encryption.
- cert - accept connections secured with TLS and a certificate.

Default: `unencrypted`

TLSFrontendCertIssuer

Specify the allowed frontend certificate issuer.

TLSFrontendCertSubject

Specify the allowed frontend certificate subject.

TLSKeyFile

Specify full pathname of the file containing the server private key, used for encrypted communications between Zabbix components.

TLSListen

Require TLS-only connections on the trapper socket.

Supported values:

- required - accept only TLS connections

TmpDir

Specify the temporary directory.

Default: `/tmp`

TrapperTimeout

Specify trapper timeout in seconds for:

- retrieval of historical data from Zabbix proxy
- sending configuration data to Zabbix proxy
- global script execution or remote command execution on Zabbix server

Default: 300
 Range: 1-300

TrendCacheSize

Set the trend cache size, in bytes.
 The shared memory size for storing trends.

Default: `4M`
 Range: 128K-16G

TrendFunctionCacheSize

Set the trend-function cache size, in bytes.
 The shared memory size for caching calculated trend function data.

Default: `4M`
 Range: 128K-2G

UnavailableDelay

Set the frequency of checking host availability during the unavailability period in seconds.

Default: 60
 Range: 1-3600

UnreachableDelay

Set the frequency of checking host availability during the unreachability period in seconds.

Default: 15
 Range: 1-3600

UnreachablePeriod

Set after how many seconds of unreachability treat host as unavailable.

Default: 45
 Range: 1-3600

User

Drop privileges to a specific, existing user on the system.
 Only has effect if run as 'root' and AllowRoot is disabled.

Default: zabbix

ValueCacheSize

Set the size of the history value cache (in bytes).
 The shared memory size for caching item history data requests.
 Setting to 0 disables the value cache (not recommended).
 When the value cache runs out of the shared memory a warning message is written to the server log every 5 minutes.

Default: 8M
 Range: 0, 128K-64G

Vault

Specify the vault provider: - *HashiCorp* - HashiCorp KV Secrets Engine version 2; - *CyberArk* - CyberArk Central Credential Provider
 Must match the vault provider set in the frontend.

Default: HashiCorp

VaultAppRoleID

Specify the HashiCorp vault AppRole ID.
 If set, enables AppRole authentication. Cannot be used with VaultToken.
 Supported only for HashiCorp vault.

VaultAppSecretID

Specify the HashiCorp vault AppRole secret ID.
 Supported only for HashiCorp vault.

VaultDBPath

Specify the location for retrieving database credentials by keys. It is a Vault path or query, depending on the Vault.

The keys used for HashiCorp are password and username.

Example path with VaultPrefix=/v1/secret/data/zabbix/:

```
database
```

Example path without VaultPrefix:

```
secret/zabbix/database
```

The keys used for CyberArk are Content and UserName.

Example:

```
AppID=zabbix_server&Query=Safe=passwordSafe;Object=zabbix_server_database
```

This option can only be used if DBUser and DBPassword are not specified.

VaultPrefix

Specify custom prefix for the vault path or query, depending on the Vault. The most suitable defaults will be used if not specified.
 Note that data is automatically appended after mountpoint for HashiCorp if VaultPrefix is not specified.

Example prefix for HashiCorp:

```
/v1/secret/data/zabbix/
```

Example prefix for Cyberark:

```
/AIMWebService/api/Accounts?
```

VaultTLSCertFile

Specify the name of the SSL certificate file used for client authentication. The certificate file must be in PEM1 format.
 If the certificate file contains also the private key, leave the SSL key file field empty.
 The directory containing this file is specified by the configuration parameter SSLCertLocation.
 This option can be omitted but is recommended for CyberArkCCP vault.

VaultTLSKeyFile

Specify the name of the SSL private key file used for client authentication. The private key file must be in PEM1 format.
 The directory containing this file is specified by the configuration parameter SSLKeyLocation.
 This option can be omitted but is recommended for CyberArkCCP vault.

VaultToken

Specify the HashiCorp vault authentication token. The HashiCorp Vault authentication token should have been generated exclusively for Zabbix server with read-only permission to the paths specified in Vault macros and read-only permission to the path specified in the optional VaultDBPath configuration parameter.
 It is an error if `VaultToken` and `VAULT_TOKEN` environment variable are defined at the same time.
 Cannot be used with VaultAppRoleID.
 Supported only for HashiCorp vault.

VaultURL

Specify the vault server URL. The system-wide CA certificates directory will be used if SSLCALocation is not specified.

Default: `https://127.0.0.1:8200`

VMwareCacheSize

Set the shared memory size for storing VMware data. A VMware internal check `zabbix[vmware,buffer,...]` can be used to monitor the VMware cache usage (see Internal checks).
 Note that shared memory is not allocated if there are no `vmware collector` instances configured to start.

Default: 8M
 Range: 256K-2G

VMwareFrequency

Set the frequency of data gathering from a single VMware service (in seconds). This frequency should be set to the least update interval of any VMware monitoring item.

Default: 60
 Range: 10-86400

VMwarePerfFrequency

Set the frequency of retrieving performance counter statistics from a single VMware service (in seconds).
 This frequency should be set to the least update interval of any VMware monitoring item that uses VMware performance counters.

Default: 60
 Range: 10-86400

VMwareTimeout

Set the maximum number of seconds a `vmware collector` will wait for a response from VMware service (vCenter or ESX hypervisor).

Default: 10
 Range: 1-300

WebDriverURL

Specify the WebDriver interface URL.

Example (used with Selenium WebDriver standalone server):

```
WebDriverURL=http://localhost:4444
```

WebServiceURL

Specify the web service URL in the format `http[s]://host:port/report`.

Example:

```
WebServiceURL=http://localhost:10053/report
```

The `http://` scheme may be omitted for non-TLS (HTTP) connections only; with TLS configured, `https://` must be used.

Footnotes

[1] Note that too many data gathering processes (pollers, unreachable pollers, ODBC pollers, HTTP pollers, Java pollers, pingers, trappers, proxypollers) together with IPMI manager, SNMP trapper, preprocessing workers, and discovery workers can exhaust the per-process file descriptor limit for the preprocessing manager.

Warning:
Exhausting the file descriptor limit will cause Zabbix server to stop, typically shortly after startup but sometimes taking longer. To avoid such issues, review the Zabbix server configuration file to optimize the number of concurrent checks and processes. Additionally, if necessary, ensure that the file descriptor limit is set sufficiently high by checking and adjusting system limits.

Proxy

Overview Zabbix proxy is a process that may collect monitoring data from one or more monitored devices and send the information to the Zabbix server, essentially working on behalf of the server. All collected data is buffered locally and then transferred to the Zabbix server the proxy belongs to.

Deploying a proxy is optional, but may be very beneficial to distribute the load of a single Zabbix server. If only proxies collect data, processing on the server becomes less CPU and disk I/O hungry.

A Zabbix proxy is the ideal solution for centralized monitoring of remote locations, branches and networks with no local administrators.

Zabbix proxy requires a separate database.

> **Attention:**
> Note that databases supported with Zabbix proxy are SQLite, MySQL and PostgreSQL.

See also:

- Distributed monitoring using proxies
- Proxy load balancing and high availability
- Synchronization of monitoring configuration
- Server-proxy data exchange protocol
- Configuration parameters: Zabbix proxy
- Set up a Zabbix Cloud proxy

Running proxy If installed as package

Zabbix proxy runs as a daemon process. The proxy can be started by executing:

```
systemctl start zabbix-proxy
```

This will work on most of GNU/Linux systems. On other systems you may need to run:

```
/etc/init.d/zabbix-proxy start
```

Similarly, for stopping/restarting/viewing status of Zabbix proxy, use the following commands:

```
systemctl stop zabbix-proxy
systemctl restart zabbix-proxy
systemctl status zabbix-proxy
```

Start up manually

If the above does not work you have to start it manually. Find the path to the `zabbix_proxy` binary and execute:

```
zabbix_proxy
```

You can use the following command-line parameters with Zabbix proxy:

Parameter	Description
`-c --config <config-file>`	Path to the configuration file.
`-f --foreground`	Run Zabbix proxy in foreground.
`-R --runtime-control <option>`	Perform administrative functions.
`-T --test-config`	Validate configuration file and exit.
`-h --help`	Print help information and exit.
`-V --version`	Display version number.

Examples of running Zabbix proxy with command-line parameters:

```
zabbix_proxy -c /usr/local/etc/zabbix_proxy.conf
zabbix_proxy --help
zabbix_proxy -V
```

Runtime control

Runtime control options:

Option	Description	Target
`config_cache_reload`	Reload configuration cache. Ignored if cache is being currently loaded. Active Zabbix proxy will connect to the Zabbix server and request configuration data. Passive Zabbix proxy will request configuration data from Zabbix server the next time when the server connects to the proxy.	
`history_cache_clear=<target>`	Clear history cache for the item specified by its ID. Affects all values of the item, except the first and last value.	**target** - ID of the item.
`diaginfo[=<section>]`	Gather diagnostic information in the proxy log file.	`historycache` - history cache statistics; `preprocessing` - preprocessing manager statistics; `locks` - list of mutexes (is empty on *BSD* systems).
`snmp_cache_reload`	Reload SNMP cache - clear SNMP engine properties (engine time, engine boots, engine id, credentials) for all hosts. Use to force a global cache clear when troubleshooting SNMP issues.	
`housekeeper_execute`	Start the housekeeping procedure. Ignored if the housekeeping procedure is currently in progress.	
`log_level_increase[=<target>]`	Increase log level, affects all processes if target is not specified. Not supported on *BSD* systems.	**process type** - all processes of specified type (e.g., `poller`). See all proxy process types. **process type,N** - process type and number (e.g., `poller,3`) **pid** - process identifier (1 to 65535). For larger values specify target as 'process type,N'.
`log_level_decrease[=<target>]`	Decrease log level, affects all processes if target is not specified. Not supported on *BSD* systems.	
`prof_enable[=<target>]`	Enable profiling. Affects all processes if target is not specified. Enabled profiling provides details of all rwlocks/mutexes by function name.	**process type** - all processes of specified type (e.g., `history syncer`). See all proxy process types. **process type,N** - process type and number (e.g., `history syncer,1`). **pid** - process identifier (1 to 65535). For larger values specify target as 'process type,N'. **scope** - `rwlock`, `mutex`, `processing` can be used with the process type and number (e.g., `history syncer,1,processing`) or all processes of type (e.g., `history syncer,rwlock`).
`prof_disable[=<target>]`	Disable profiling. Affects all processes if target is not specified.	**process type** - all processes of specified type (e.g., `history syncer`). See all proxy process types. **process type,N** - process type and number (e.g., `history syncer,1`). **pid** - process identifier (1 to 65535). For larger values specify target as 'process type,N'.

When a process type is specified as a target for runtime control, the command applies only to running processes of that type. If no process of the specified type is currently running, the command returns an error:

```
zabbix_proxy -R log_level_increase='data sender'
Cannot redirect signal: "data sender" process does not exist
```

Example of using runtime control to reload the proxy configuration cache:

```
zabbix_proxy -c /usr/local/etc/zabbix_proxy.conf -R config_cache_reload
```

Example of using runtime control to clear the history cache for an item:

```
zabbix_proxy -c /usr/local/etc/zabbix_proxy.conf -R history_cache_clear=42243
```

Examples of using runtime control to gather diagnostic information:

```
### Gather all available diagnostic information in the proxy log file:
zabbix_proxy -R diaginfo

### Gather history cache statistics in the proxy log file:
zabbix_proxy -R diaginfo=historycache
```

Example of using runtime control to reload the SNMP cache:

```
zabbix_proxy -R snmp_cache_reload
```

Attention:
When an SNMPv3 interface is updated via the Zabbix UI, Zabbix will automatically reload the new SNMPv3 credentials for that interface in most cases; use `-R snmp_cache_reload` only if polling still fails after credential changes (for example, due to engineBoots/engineID inconsistencies or non-RFC devices), or when you need to force a global SNMP cache clear for troubleshooting.

Example of using runtime control to trigger execution of housekeeper:

```
zabbix_proxy -c /usr/local/etc/zabbix_proxy.conf -R housekeeper_execute
```

Examples of using runtime control to change log level:

```
### Increase log level of all processes:
zabbix_proxy -c /usr/local/etc/zabbix_proxy.conf -R log_level_increase

### Increase log level of second poller process:
zabbix_proxy -c /usr/local/etc/zabbix_proxy.conf -R log_level_increase=poller,2

### Increase log level of process with PID 1234:
zabbix_proxy -c /usr/local/etc/zabbix_proxy.conf -R log_level_increase=1234

### Decrease log level of all http poller processes:
zabbix_proxy -c /usr/local/etc/zabbix_proxy.conf -R log_level_decrease="http poller"
```

Process user

Zabbix proxy is designed to run as a non-root user. It will run as whatever non-root user it is started as. So you can run proxy as any non-root user without any issues.

If you will try to run it as 'root', it will switch to a hardcoded 'zabbix' user, which must be present on your system. You can only run proxy as 'root' if you modify the `AllowRoot` parameter in the proxy configuration file accordingly.

Configuration file

The default path to the configuration file depends on the operating system and Zabbix installation method:

- For installation from sources: `/usr/local/etc/zabbix_proxy.conf`, or as set by compile-time variables `--sysconfdir` or `--prefix`.
- For installation from packages: `/etc/zabbix/zabbix_proxy.conf`.

See the configuration file options for details on configuring `zabbix_proxy`.

Proxy process types and threads

- `agent poller` - asynchronous poller process for passive checks with a worker thread;
- `apm manager` - manager process for application performance monitoring (APM);
- `apm worker` - process for handling application performance monitoring (APM) tasks from the `apm manager`;
- `availability manager` - process for host availability updates;
- `browser poller` - poller for browser item checks;
- `configuration syncer` - process for managing in-memory cache of configuration data;
- `data sender` - proxy data sender;

- discovery manager - manager process for discovery of devices;
- discovery worker - process for handling discovery tasks from the discovery manager;
- history syncer - history DB writer;
- housekeeper - process for removing outdated item history;
- http agent poller - asynchronous poller process for HTTP checks with a worker thread;
- http poller - web monitoring poller;
- icmp pinger - poller for icmpping checks;
- internal poller - poller for internal checks;
- ipmi manager - IPMI poller manager;
- ipmi poller - poller for IPMI checks;
- java poller - poller for Java checks;
- odbc poller - poller for ODBC checks;
- poller - normal poller for passive checks;
- preprocessing manager - manager of preprocessing tasks with preprocessing worker threads;
- preprocessing worker - thread for data preprocessing;
- self-monitoring - process for collecting internal server statistics;
- snmp poller - asynchronous poller process for SNMP checks with a worker thread (walk[OID] and get[OID] items only);
- snmp trapper - trapper for SNMP traps;
- supervisor - coordinates the startup and shutdown of server processes and threads, and monitors their readiness;
- task manager - process for remote execution of tasks requested by other components (e.g., close problem, acknowledge problem, check item value now, remote command functionality);
- telemetry query poller - poller for telemetry query item checks;
- trapper - trapper for active checks, traps, and proxy communication;
- unreachable poller - poller for unreachable devices;
- vmware collector - VMware data collector responsible for data gathering from VMware services.

The proxy log file can be used to observe these process types.

The proxy log file is created with read-write permissions for the file owner only. Additionally the file is readable by the owner group. All other permissions are denied.

Various types of Zabbix proxy processes can be monitored using the zabbix[process,<type>,<mode>,<state>] internal item.

History syncer transaction statistics

The history syncer process title displays detailed statistics about history syncer transactions.

```
205276 ?        S      0:00  zabbix_proxy: history syncer #1 [processed 1 values in 0.001179 (0.001167,0.0
205277 ?        S      0:00  zabbix_proxy: history syncer #2 [processed 0 values in 0.000022 (0.000000,0.0
```

The timings, in processed...in N (<timings>) sec, are:

- Time spent writing item values into database.
- Time spent updating item data (state, errors).

Housekeeping procedure

Zabbix proxy has the housekeeper process that removes outdated item history and trends. It runs in cycles, with frequency determined by HousekeepingFrequency, and the per-cycle deletion limit determined by ProxyLocalBuffer and ProxyOfflineBuffer. Unlike in the Zabbix server housekeeping procedure, the proxy housekeeper process does not use the housekeeper table—it deletes all outdated data once per housekeeping cycle.

Supported platforms Zabbix proxy runs on the same list of supported platforms as Zabbix server.

Memory buffer The memory buffer allows you to store new data (item values, network discovery, host autoregistration) in the buffer and upload to Zabbix server without accessing the database. The memory buffer was introduced for the proxy in Zabbix 7.0.

In installations before Zabbix 7.0 the collected data was stored in the database before uploading to Zabbix server. For these installations this remains the default behavior after upgrading to Zabbix 7.0.

For optimized performance, it is recommended to configure the use of memory buffer on the proxy. This is possible by modifying the value of ProxyBufferMode from disk (hardcoded default for existing installations) to hybrid (recommended) or memory. It is also required to set the memory buffer size (ProxyMemoryBufferSize parameter).

In hybrid mode the buffer is protected from data loss by flushing unsent data to the database if the proxy is stopped, the buffer is full or data too old. When all values have been flushed into database, the proxy goes back to using memory buffer.

In memory mode, the memory buffer will be used, however, there is no protection against data loss. If the proxy is stopped, or the memory gets overfilled, the unsent data will be dropped.

The hybrid mode (`ProxyBufferMode=hybrid`) is applied to all new installations since Zabbix 7.0.

Additional parameters such as `ProxyMemoryBufferSize` and `ProxyMemoryBufferAge` define the memory buffer size and the maximum age of data in the buffer, respectively.

Note that with conflicting configuration the proxy will print an error and fail to start, for example, if:

- `ProxyBufferMode` is set to `hybrid` or `memory` and `ProxyMemoryBufferSize` is 0.
- `ProxyBufferMode` is set to `hybrid` or `memory` and `ProxyLocalBuffer` is not 0.

Locale Note that the proxy requires a UTF-8 locale so that some textual items can be interpreted correctly. Most modern Unix-like systems have a UTF-8 locale as default, however, there are some systems where that may need to be set specifically.

Calculation of queues during maintenance

> **Attention:**
> The Zabbix proxy is not aware of maintenance periods; see Calculation of queues durin

Using Zabbix frontend behind proxy If Zabbix frontend runs behind proxy server, the cookie path in the proxy configuration file needs to be rewritten in order to match the reverse-proxied path. See examples below. If the cookie path is not rewritten, users may experience authorization issues, when trying to login to Zabbix frontend.

Example configuration for NGINX

```
### ..
location / {
### ..
proxy_cookie_path /zabbix /;
proxy_pass http://192.168.0.94/zabbix/;
### ..
```

Example configuration for Apache

```
### ..
ProxyPass "/" http://host/zabbix/
ProxyPassReverse "/" http://host/zabbix/
ProxyPassReverseCookiePath /zabbix /
ProxyPassReverseCookieDomain host zabbix.example.com
### ..
```

Distributed monitoring using proxies

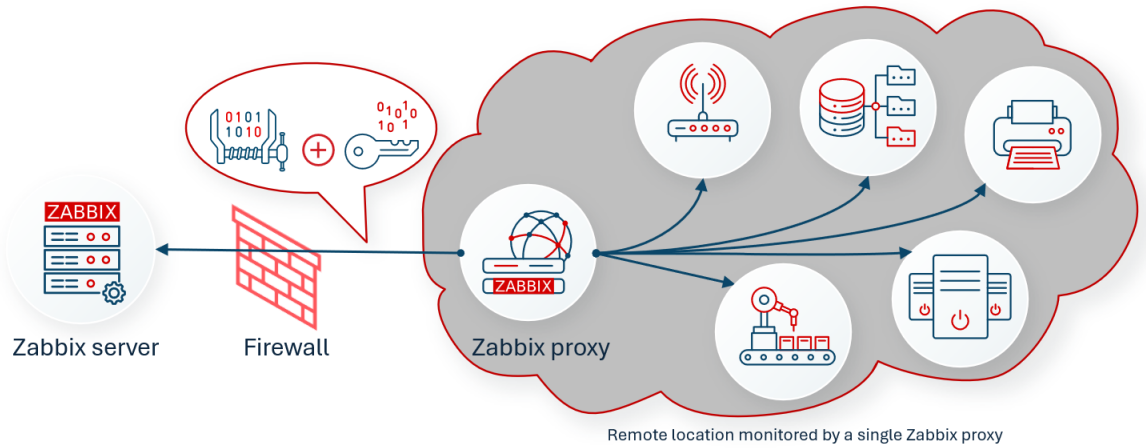
Overview

Zabbix proxies can be used to create an effective and reliable distributed monitoring setup.

Proxies are deployed locally to collect performance and availability data. Then the data is reported to a centralized server.

A Zabbix proxy can be used to:

- Monitor remote locations
- Monitor locations having unreliable communications
- Offload Zabbix server when monitoring thousands of devices
- Collect telemetry data for application performance monitoring (APM)



The proxy requires only one TCP connection to the Zabbix server. This way it is easier to get around a firewall as you only need to configure one firewall rule.

Attention:
Zabbix proxy must use a separate database. Pointing it to the Zabbix server database will break the configuration.

All data collected by the proxy is stored locally before transmitting it over to the server. This way no data is lost due to any temporary communication problems with the server. The `ProxyLocalBuffer` and `ProxyOfflineBuffer` parameters in the proxy configuration file control for how long the data is kept locally.

Attention:
It may happen that a proxy, which receives the latest configuration changes directly from Zabbix server database, has a more up-to-date configuration than Zabbix server whose configuration may not be updated as fast due to the value of CacheUpdateFrequency. As a result, proxy may start gathering data and send them to Zabbix server that ignores these data.

Zabbix proxy is a data collector. It does not calculate triggers, process events or send alerts. For an overview of what proxy functionality is, review the following table:

Function		Supported by proxy
Items		
	Zabbix agent checks	**Yes**
	Zabbix agent checks (active)	**Yes** [1]
	Simple checks	**Yes**
	Trapper items	**Yes**
	SNMP checks	**Yes**
	SNMP traps	**Yes**
	IPMI checks	**Yes**
	JMX checks	**Yes**
	Log file monitoring	**Yes**
	Internal checks	**Yes**
	SSH checks	**Yes**
	Telnet checks	**Yes**
	External checks	**Yes**
	Dependent items	**Yes**
	Script items	**Yes**
	Browser items	**Yes**
	Telemetry query items	**Yes**
Built-in web monitoring		**Yes**
Item value preprocessing		**Yes**
Network discovery		**Yes**
Active agent autoregistration		**Yes**
Low-level discovery		**Yes** [2]
Remote commands		**Yes**
Telemetry data collection (APM)		**Yes** [3]
Calculating triggers		*No*

Function	Supported by proxy
Processing events	*No*
Event correlation	*No*
Sending alerts	*No*

Note:
[1] To make sure that an agent asks the proxy (and not the server) for active checks, the proxy must be listed in the **ServerActive** parameter in the agent configuration file.
 [2] For LLD, Zabbix proxy only collects and preprocesses the data and then sends it to Zabbix server for further processing.
 [3] Telemetry data is collected only if APM is enabled on the proxy (see the *APM* tab in proxy configuration); Zabbix server cannot receive telemetry data.

Proxy features

When making a choice of using/not using a proxy, several considerations must be taken into account.

	Proxy
Lightweight	**Yes**
GUI	No
Works independently	**Yes**
Easy maintenance	**Yes**
Automatic DB creation	**Yes**[1]
Local administration	No
Ready for embedded hardware	**Yes**
One way TCP connections	**Yes**
Centralized configuration	**Yes**
Generates notifications	No

[1] Automatic DB creation feature works only with SQLite. Other supported databases require manual setup.

Attention:
The Zabbix proxy is not aware of maintenance periods; see Calculation of queues during maintenance for details.

Protection from overloading

If Zabbix server was down for some time, and proxies have collected a lot of data, and then the server starts, it may get overloaded (history cache usage stays at 95-100% for some time). This overload could result in a performance hit, where checks are processed slower than they should. Protection from this scenario was implemented to avoid problems that arise due to overloading history cache.

When Zabbix server history cache is full the history cache write access is being throttled, stalling server data gathering processes. The most common history cache overload case is after server downtime when proxies are uploading gathered data. To avoid this proxy throttling was added (currently it cannot be disabled).

When history cache usage reaches 80%, the Zabbix server enters throttling mode. In throttling mode the server accepts proxy data only when history cache usage is below 60%, rotating the accepted proxies. Once history cache usage drops below 20%, the server switches back to normal mode.

Additionally, in normal mode the Zabbix server throttles individual proxies sending very large packets (10,000+ records) if history cache usage exceeds 60%. This decision is applied at the instant the server evaluates a proxy upload and therefore may not always be reflected immediately in history cache usage graphs (the `zabbix[wcache,history,pused]` internal item and its update interval may miss short spikes).

This throttling mode will continue until either cache usage hits 80% again or falls down to 20% or the throttling list is empty. In the first case the server will stop accepting proxy data again. In the other two cases the server will start working normally, accepting data from all proxies.

The above information can be illustrated in the following table:

History write cache usage	Zabbix server mode	Zabbix server action
Reaches 80%	Wait	Stops accepting proxy data, but maintains a *throttling list* (prioritized list of proxies to be contacted later).

History write cache usage	Zabbix server mode	Zabbix server action
Reaches 60%	Normal, but prepared to be throttled	May refuse very large proxy uploads (more than 10k records) when deciding whether to accept data; continues accepting other proxy data.
Drops to 20%	Normal	Drops the throttling list and starts accepting proxy data normally.

You may use the zabbix[wcache,history,pused] internal item to correlate this behavior of Zabbix server with a metric.

Configuration

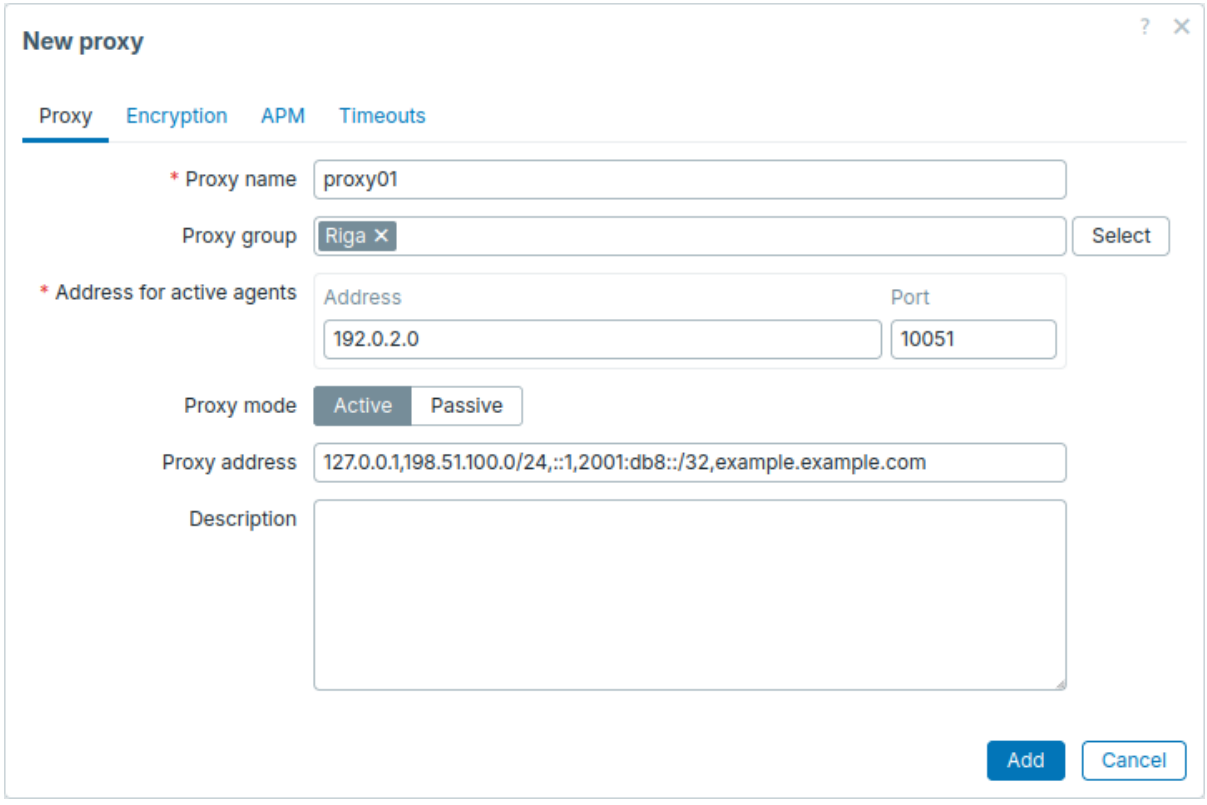
Once you have installed and configured a proxy, it is time to configure it in the Zabbix frontend.

Adding proxies

Note:
Instead of installing a proxy on your own infrastructure, you can deploy one in Zabbix Cloud. To do so, click *Create proxy in Cloud* in *Administration* > *Proxies*; Zabbix Cloud will open in a new browser tab. After creating a proxy node there, configure the proxy in Zabbix frontend as described below. For more information, see Set up a Zabbix Cloud proxy.

To configure a proxy in Zabbix frontend:

1. Go to *Administration* > *Proxies*.
2. Click *Create proxy*.
3. Enter the proxy parameters in the form.



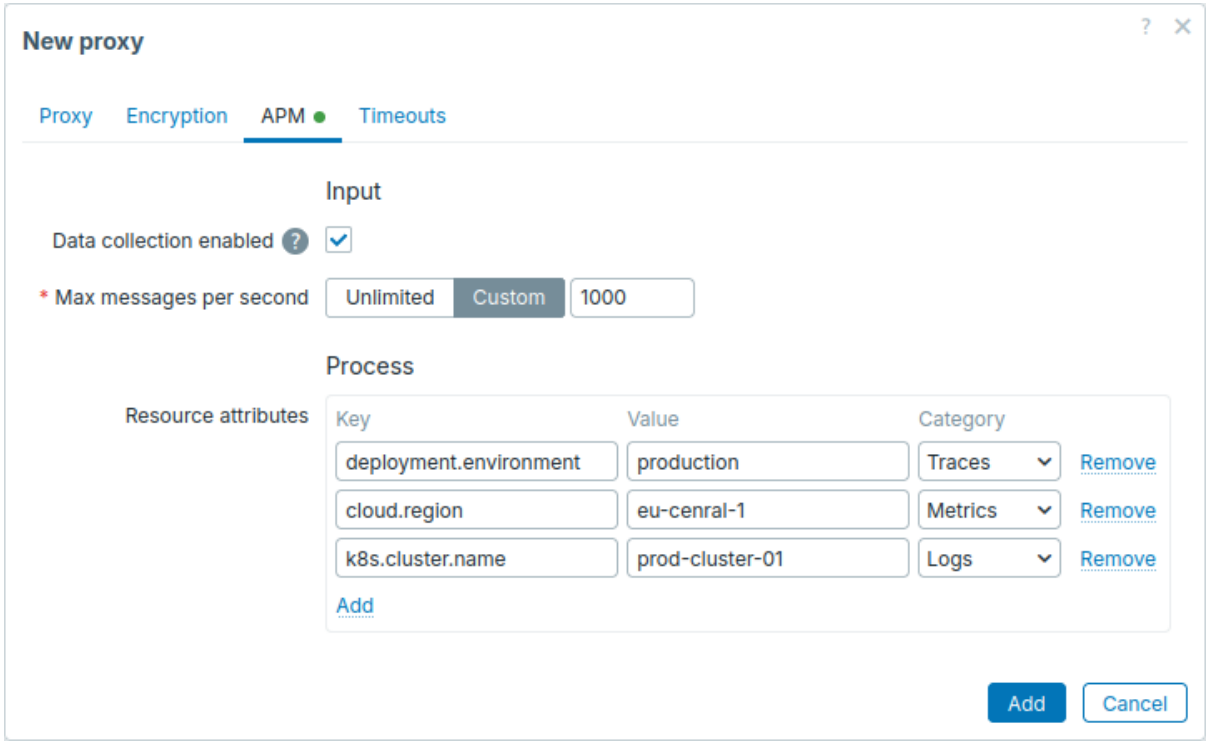
Parameter	Description
Proxy name	Enter the proxy name. It must be the same name as in the *Hostname* parameter in the proxy configuration file.
Proxy group	Select one proxy group for proxy load balancing/high availability.
Address for active agents	Enter address the monitored active agents or senders must connect to. Supported **only** for Zabbix 7.0 agents or later. This address is used to connect to both active and passive proxies. This field is only available if a proxy group is selected in the *Proxy group* field.
Address	IP address/DNS name to connect to.
Port	TCP port number (10051 by default) to connect to. User macros are supported.

Parameter	Description
Proxy mode	Select the proxy mode. **Active** - the proxy will connect to the Zabbix server and request configuration data **Passive** - Zabbix server connects to the proxy *Note* that without encrypted communications (sensitive) proxy configuration data may become available to parties having access to the Zabbix server trapper port when using an active proxy. This is possible because anyone may pretend to be an active proxy and request configuration data if authentication does not take place or proxy addresses are not limited in the *Proxy address* field.
Proxy address	If specified then active proxy requests are only accepted from this list of comma-delimited IP addresses, optionally in CIDR notation, or DNS names of active Zabbix proxy. This field is only available if an active proxy is selected in the *Proxy mode* field. Macros are not supported.
Interface	Enter interface details for a passive proxy. This field is only available if a passive proxy is selected in the Proxy mode field.
Address	IP address/DNS name of the passive proxy.
Port	TCP port number of the passive proxy (10051 by default). User macros are supported.
Description	Enter the proxy description.

The **Encryption** tab allows you to require encrypted connections with the proxy.

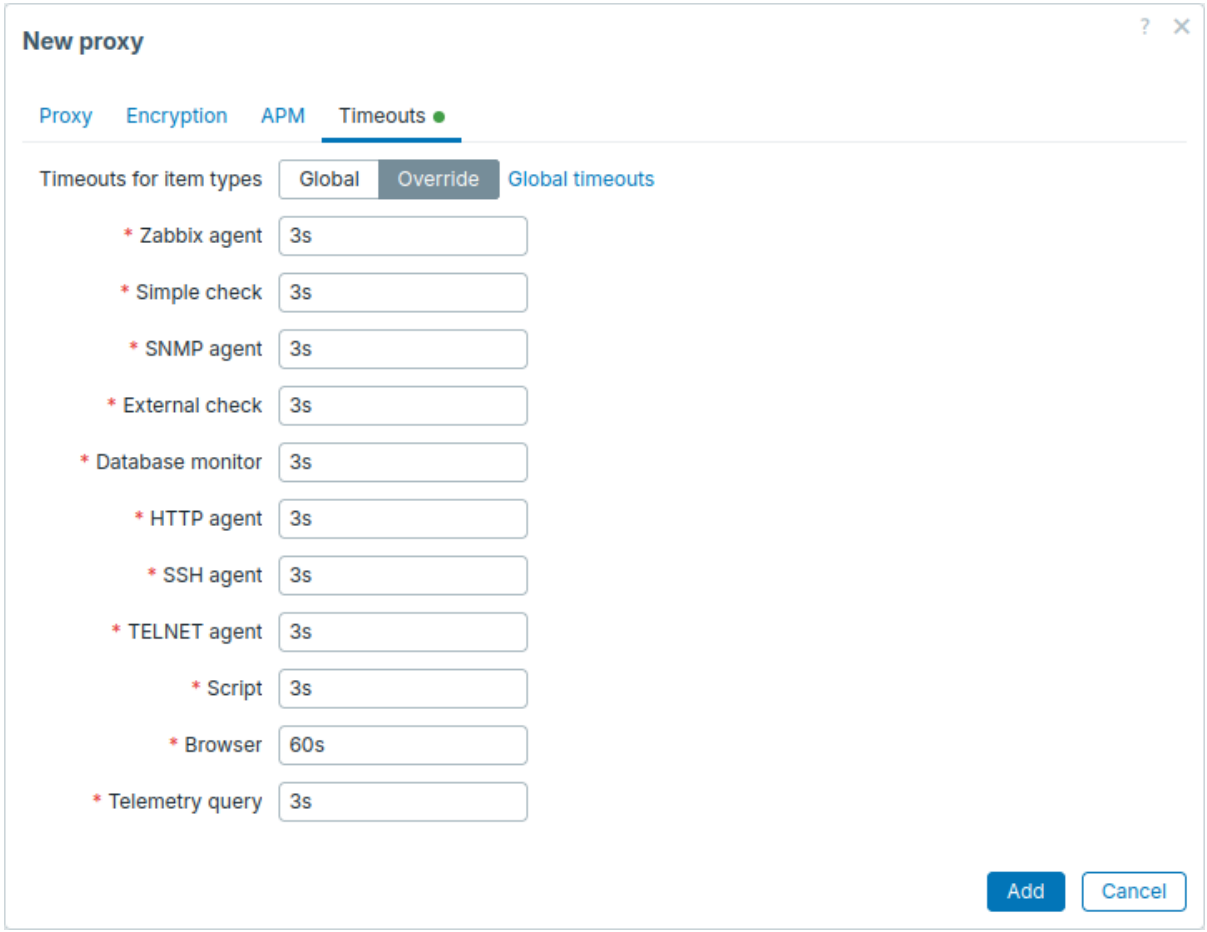
Parameter	Description
Connections to proxy	How the server connects to the passive proxy: no encryption (default), using PSK (pre-shared key) or certificate.
Connections from proxy	Select what type of connections are allowed from the active proxy. Several connection types can be selected at the same time (useful for testing and switching to other connection type). Default is "No encryption".
Issuer	Allowed issuer of certificate. Certificate is first validated with CA (certificate authority). If it is valid, signed by the CA, then the *Issuer* field can be used to further restrict allowed CA. This field is optional, intended to use if your Zabbix installation uses certificates from multiple CAs.
Subject	Allowed subject of certificate. Certificate is first validated with CA. If it is valid, signed by the CA, then the *Subject* field can be used to allow only one value of *Subject* string. If this field is empty then any valid certificate signed by the configured CA is accepted.
PSK identity	Pre-shared key identity string. Do not put sensitive information in the PSK identity, it is transmitted unencrypted over the network to inform a receiver which PSK to use.
PSK	Pre-shared key (hex-string). Maximum length: 512 hex-digits (256-byte PSK) if Zabbix uses GnuTLS or OpenSSL library, 64 hex-digits (32-byte PSK) if Zabbix uses mbed TLS (PolarSSL) library. Example: 1f87b595725ac58dd977beef14b97461a7c1045b9a1c963065002c5473194952

The **APM** tab allows you to configure application performance monitoring (APM) on the proxy. The proxy can collect telemetry data (traces, metrics, and logs) sent by applications or OpenTelemetry collectors over the OpenTelemetry protocol (OTLP/gRPC), and write the data to the telemetry data source.



Parameter	Description
Data collection enabled	Mark the checkbox to enable telemetry data collection on the proxy. Before enabling telemetry data collection, set the `StartAPMCollector` parameter to 1 in the proxy configuration file (requires Zabbix proxy with APM support).
Max messages per second	Set the maximum number of messages per second the proxy accepts: **Unlimited** - accept all incoming messages. **Custom** - accept up to the specified number of messages per second. A message is an OTLP request sent by an application, an OpenTelemetry collector, or another telemetry source. Its size depends on the sender; for example, an OpenTelemetry collector may batch thousands of spans into one message, while a simple application may send smaller messages (e.g., one per trace). If the limit is exceeded, the proxy responds with the gRPC status code `UNAVAILABLE` (temporary error), so the sender may retry sending the message later, depending on its configuration.
Resource attributes	Specify the resource attributes to add to incoming telemetry data of the selected category (e.g., to mark all data received by this proxy with its environment or region). You can then use these attributes to filter telemetry data, for example, in telemetry query item conditions. If incoming data already contains a resource attribute with the same key, the incoming attribute is kept instead of the one specified here.
Key	Enter the attribute key (e.g., `deployment.environment`, `cloud.region`). Must be unique for the selected *Category*.
Value	Enter the attribute value (e.g., `production`, `eu-central-1`). Can be empty.
Category	Select the category of incoming data the attribute applies to: *Traces*, *Metrics*, or *Logs*.

The **Timeouts** tab allows you to override global timeouts for item types that support it.



Parameter	Description
Timeouts for item types	Set item timeout (based on its type): **Global** - use global timeout (displayed in the grayed out *Timeout* field per item type); **Override** - set a custom timeout (in the *Timeout* field for each item type). Allowed range: 1 - 600s (default: inherited from global timeouts). Time suffixes, e.g., 30s, 1m, and user macros are supported. Clicking the *Global timeouts* link allows you to configure global timeouts. Note that the *Global timeouts* link is visible only to users of *Super admin* type with permissions to the *Administration > General* frontend section. Note that while proxy-level timeouts override the global ones, they will be overridden by individual item timeouts, if configured.

Note:

If the proxy major version does not match the server major version, icon will be displayed next to *Timeouts for item types*, with the hover message "Timeouts disabled because the proxy and server versions do not match". In such cases, the proxy will use the `Timeout` parameter from the proxy configuration file.

The editing form of an existing proxy has the following additional buttons:

- *Refresh configuration* - refresh configuration of the proxy;
- *Clone* - create a new proxy based on the properties of the existing proxy;
- *Delete* - delete the proxy.

Host configuration

You can specify that an individual host should be monitored by a proxy or proxy group in the host configuration form, using the *Monitored by* field.



Host mass update is another way of specifying that hosts should be monitored by a proxy or proxy group.

Proxy load balancing and high availability

Overview

Zabbix proxies can be organized into proxy groups to enable proxy load balancing and high availability.

Proxy load balancing and high availability is the automatic redistribution of hosts between proxies within a proxy group:

- If a proxy goes offline, its hosts will be moved to other proxies, thus maintaining high proxy availability.
- If a proxy has a much higher/lower number of hosts than other proxies, its hosts will be moved to other proxies to balance proxy load.

Host redistribution works only between proxies in a group that meet the following conditions:

- Proxies are running Zabbix 7.0 or later.
- Proxy version matches Zabbix server version. If using Zabbix agent (passive), proxy version must match agent version. Active agents only require Zabbix 7.0 or later.
- The proxy group has an online state.
- Hosts are configured to be monitored by a proxy group rather than by individual proxies.

Note:
Proxy group health can be monitored with internal checks by any host assigned to a proxy group. However, to monitor the health of a single proxy in a group, assign the host to that proxy; otherwise, results may be inconsistent.

Host redistribution

Proxy load balancing and high availability is managed by Zabbix server via the proxy group manager, which continuously monitors the state of all proxies in each proxy group and their host distribution.

Proxy high availability within a group is ensured via proxy failover: when a proxy goes offline, its hosts are immediately redistributed to other proxies. Proxy load balancing also occurs, as hosts are reassigned to proxies with the fewest assigned hosts.

Additionally, proxy load balancing is triggered when proxy host count differs from the group average by at least 10 hosts and a factor of 2 (host excess or host deficit). If the imbalance persists after a grace period (10 x failover delay), the proxy group is queued for host redistribution.

The proxy group manager redistributes hosts using the following logic:

1. Calculate the average number of hosts per proxy.
2. For proxies with host excess—move the excess hosts to the unassigned pool of proxies.
3. For proxies with host deficit—calculate how many hosts are needed to reach balance.
4. Remove the required number of hosts from proxies with the most hosts.
5. Move unassigned hosts to proxies with the fewest hosts.

Examples of host redistribution:

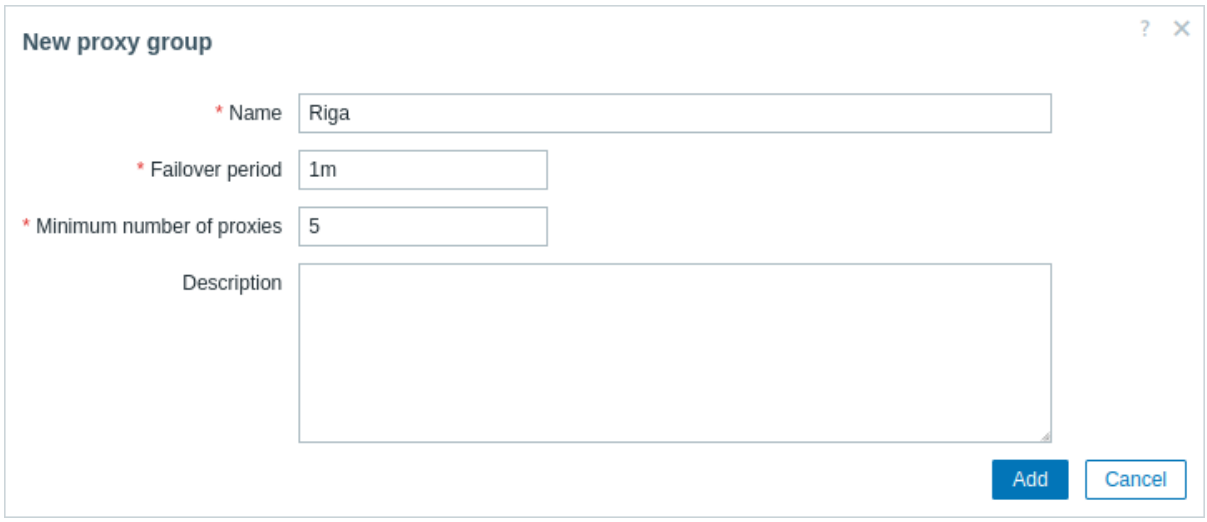
Hosts on proxy	Group average	Host reassignment
100	50	Yes
60	50	No
40	50	No
25	50	Yes
15	5	Yes
10	5	No

Attention:
Having fewer than 10 hosts monitored by a proxy group may lead to uneven host distribution among proxies in the group.

Configuring a proxy group

To configure a proxy group in Zabbix frontend:

1. Go to *Administration* > *Proxy groups*.
2. Click *Create proxy group*.
3. Enter the proxy group parameters in the form.



Parameter	Description
Name	Name of the proxy group.
Failover period	Period in seconds during which a proxy in the proxy group must communicate with Zabbix server to be considered online (default: 1m; range: 10s-15m). If the proxy does not communicate within this period, the proxy state is changed to *Offline*, and its hosts are immediately redistributed to other proxies. Proxy load balancing begins after 10 x this period. Supports time suffixes (e.g., 30s, 1m) and user macros.
Minimum number of proxies	Minimum number of online proxies required to keep the proxy group online (default: 1; range: 1-1000). Supports user macros. This value should be less than the total number of proxies in the group. For example, in a group of 10 proxies, setting the minimum to 10 will cause the group to go offline if any proxy fails. Note that online proxies in an offline group continue to function normally, but load balancing/high availability will not take place.
Description	Description of the proxy group.
Proxies	Displays a list of up to five proxies (as links or in plain text, depending on user permissions to proxies) when editing a group with proxies.

Configuring proxy load balancing

To use proxy load balancing, you need to configure a proxy group in Zabbix frontend (see above) and make sure that hosts are monitored by a proxy group, not individual proxies (you may use host mass update to move hosts from proxies to the proxy group).

If using Zabbix agent, also configure it as follows:

- For passive checks, list all proxy group proxies in the Server parameter.
- For active checks, it's recommended to list all proxy group proxies or Zabbix server in the ServerActive parameter. Note that only Zabbix agent 7.0 (or later) will work with proxy groups in its active mode.

If the `ServerActive` parameter includes only one proxy from a proxy group (or Zabbix server), the agent will still be able to connect to the correct proxy. When the agent service starts and connects to the specified proxy, the agent will receive and cache the full list of proxy IPs and their current load within the group. Then, active checks will be redirected to the correct online proxy for the host, based on the current proxy-host assignment within the proxy group.

Warning:
Having only a single proxy specified in Zabbix agent `ServerActive` parameter may lead to lost monitoring data if the agent is started/rebooted while the specified proxy is offline.

Note:
When using Zabbix sender, data requests are also redirected to the correct online proxy for the host, based on the current proxy-host assignment within the proxy group. However, if you're sending values of multiple hosts from an input file, use the −g option to prevent sending data to the wrong proxy.

Zabbix agent must also be able to connect to all proxies in the proxy group through the firewall. Otherwise, active checks may hang or fail during redirection or failover. For example:

- During active checks, a proxy may redirect the agent to another proxy. If that proxy is blocked by a firewall, the communication will hang while waiting for a response.
- In stable high-availability setups without recent rebalancing, agents may never contact backup proxies. If firewall rules changed and weren't tested, failover may fail.

Testing proxy load balancing

To test proxy load balancing:

1. Configure a proxy group.
2. Make sure that the proxy group has an online state.
3. Make sure that hosts are monitored by a proxy group, not individual proxies (you may use host mass update to move hosts from proxies to the proxy group).
4. Wait a few seconds for the configuration update and host distribution among proxies in the proxy group. Observe the change by refreshing the host list in *Administration* > *Proxies*.

Important notes

- SNMP traps are not supported by proxies in proxy group.
- Checks that depend on external configuration (e.g., scripts for external checks or ODBC configuration for database checks) must have the same configuration on all proxies in proxy group.
- Database checks require extended permissions on the database object/server.
- VMware hosts monitored by a proxy group will be randomly spread between proxies in the group. This causes each proxy to cache all VMware data, resulting in additional load on vCenter.
- Hosts created based on autoregistration data from a proxy in a proxy group are set to be monitored by that proxy group. However, hosts created based on network discovery data from a proxy in a proxy group are set to be monitored by that proxy.

Synchronization of monitoring configuration

Overview

This page provides details on the monitoring configuration update for the proxy, i.e. how changes made to the monitoring configuration on the server are synchronized to the proxy.

Incremental update

The proxy configuration update is incremental. During a configuration sync only the modified entities are updated (thus, if no entities have been modified, nothing will be sent). This approach allows you to save resources and set a smaller interval (almost instant) for the proxy configuration update.

Proxy configuration changes are tracked using revision numbers. Only entities with revisions larger than the proxy configuration revision are included in configuration data sent to the proxy.

The entities for a configuration sync are as follows:

Entity	Details
autoregistration tls data	All autoregistration TLS data.
expressions	All expressions (regular expressions, expression tables).
global configuration	Global configuration defined in the 'config' table
host	All properties, interfaces, inventory, items, item preprocessing, item parameters, web scenarios of a host.
host macros	All macros defined on a host and all template IDs linked to it.
proxy discovery rule	Discovery rules and checks assigned to a proxy.

That means:

- If an item is changed on a **host**, all configuration of that host will be synced.

- If a **regular expression** is changed, all regular expressions will be synced.

An exception are the host macros which are sent also if anything on the host has been changed.

The `-R config_cache_reload` command on the proxy will also initiate an incremental update.

Note that a full configuration sync will take place on a proxy start/restart, HA failover, if the session token has changed, or if the configuration update failed on the proxy, for example, if the connection was broken while receiving configuration data.

Configuration parameters

The **ProxyConfigFrequency** parameter determines how often the proxy configuration is synced with the server (10 seconds by default).

Note that ProxyConfigFrequency is:

- server parameter for passive proxies
- proxy parameter for active proxies

On active proxies ProxyConfigFrequency is a new parameter since Zabbix 6.4 and must be used instead of the now-deprecated ConfigFrequency.

Attention:
If both ProxyConfigFrequency and ConfigFrequency are used, the proxy will log an error and terminate.

Calculation of queues during maintenance

Attention:
The Zabbix proxy is not aware of maintenance periods; see Calculation of queues during maintenance for details.

Server-proxy data exchange protocol

Server - proxy data exchange is based on JSON format.

Request and response messages must begin with header and data length.

Passive proxy

Configuration request

The server will first send an empty `proxy config` request. This request is sent every `ProxyConfigFrequency` (server configuration parameter) seconds.

The proxy responds with the current proxy version, session token and configuration revision. The server responds with the configuration data that need to be updated.

name	value type	description
server→proxy:		
request	*string*	'proxy config'
proxy→server:		
version	*string*	Proxy version (<major>.<minor>.<build>).
session	*string*	Proxy configuration session token.
config_revision	*number*	Proxy configuration revision.
server→proxy:		
full_sync	*number*	1 - if full configuration data is sent; absent otherwise (optional).
data	*array*	Object of table data. Absent if configuration has not been changed (optional).
<table>	*object*	One or more objects with <table> data (optional, depending on changes).
fields	*array*	Array of field names.
-	*string*	Field name.
data	*array*	Array of rows.
-	*array*	Array of columns.
-	*string*,*number*	Column value with type depending on column type in database schema.

name	value type	description
macro.secrets	*object*	Secret macro information, absent if there are no changes in vault macros (optional).
config_revision	*number*	Configuration cache revision - sent with configuration data (optional).
del_hostids	*array*	Array of removed hostids (optional).
-	*number*	Host identifier.
del_macro_hostids	*array*	Array of hostids with all macros removed (optional).
-	*number*	Host identifier.
proxy→server:		
response	*string*	Request success information ('success' or 'failed').
version	*string*	Proxy version (<major>.<minor>.<build>).

Example:

server→proxy:

```
{
  "request":"proxy config"
}
```

proxy→server:

```
{
  "version": "8.2.0",
  "session": "0033124949800811e5686dbfd9bcea98",
  "config_revision": 0
}
```

server→proxy:

```
{
    "full_sync": 1,
    "data": {
        "hosts": {
            "fields": ["hostid", "host", "status", "ipmi_authtype", "ipmi_privilege", "ipmi_username", "ip
            "data": [
                [10084, "Zabbix server", 0, -1, 2, "", "", "Zabbix server", 1, 1, "", "", "", ""]
            ]
        },
        "interface": {
            "fields": ["interfaceid", "hostid", "main", "type", "useip", "ip", "dns", "port", "available"]
            "data": [
                [1, 10084, 1, 1, 1, "127.0.0.1", "", "10053", 1]
            ]
        },
        "interface_snmp": {
            "fields": ["interfaceid", "version", "bulk", "community", "securityname", "securitylevel", "au
            "data": []
        },
        "host_inventory": {
            "fields": ["hostid", "type", "type_full", "name", "alias", "os", "os_full", "os_short", "seria
            "data": [
                [10084, "", "", ..., "56.95387", "24.22067", ..., ""]
            ]
        },
        "items": {
            "fields": ["itemid", "type", "snmp_oid", "hostid", "key_", "delay", "history", "status", "valu
            "data": [
                [44161, 7, "", 10084, "agent.hostmetadata", "10s", "90d", 0, 1, "", "", "", "", 0, "", "",
                [44162, 0, "", 10084, "agent.ping", "10s", "90d", 0, 3, "", "", "", "", 0, "", "", "", "",
            ]
```

```
        },
        "item_rtdata": {
            "fields": ["itemid", "lastlogsize", "mtime"],
            "data": [
                [44161, 0, 0],
                [44162, 0, 0]
            ]
        },
        "item_preproc": {
            "fields": ["item_preprocid", "itemid", "step", "type", "params", "error_handler", "error_handl
            "data": []
        },
        "item_parameter": {
            "fields": ["item_parameterid", "itemid", "name", "value"],
            "data": []
        },
        "globalmacro": {
            "fields": ["globalmacroid", "macro", "value", "type"],
            "data": [
                [2, "{$SNMP_COMMUNITY}", "public", 0]
            ]
        },
        "hosts_templates": {
            "fields": ["hosttemplateid", "hostid", "templateid", "link_type"],
            "data": []
        },
        "hostmacro": {
            "fields": ["hostmacroid", "hostid", "macro", "value", "type", "automatic"],
            "data": [
                [5676, 10084, "{$M}", "AppID=zabbix_server&Query=Safe=passwordSafe;Object=zabbix:Content",
            ]
        },
        "drules": {
            "fields": ["druleid", "name", "iprange", "delay"],
            "data": [
                [2, "Local network", "127.0.0.1", "10s"]
            ]
        },
        "dchecks": {
            "fields": ["dcheckid", "druleid", "type", "key_", "snmp_community", "ports", "snmpv3_securityn
            "data": [
                [2, 2, 9, "system.uname", "", "10052", "", 0, "", "", 0, 0, 0, "", 1, 0]
            ]
        },
        "regexps": {
            "fields": ["regexpid", "name"],
            "data": [
                [1, "File systems for discovery"],
                [2, "Network interfaces for discovery"],
                [3, "Storage devices for SNMP discovery"],
                [4, "Windows service names for discovery"],
                [5, "Windows service startup states for discovery"]
            ]
        },
        "expressions": {
            "fields": ["expressionid", "regexpid", "expression", "expression_type", "exp_delimiter", "case
            "data": [
                [1, 1, "^(btrfs|ext2|ext3|ext4|reiser|xfs|ffs|ufs|jfs|jfs2|vxfs|hfs|apfs|refs|ntfs|fat32|z
                [3, 3, "^(Physical memory|Virtual memory|Memory buffers|Cached memory|Swap space)$", 4, ",
                [5, 4, "^(MMCSS|gupdate|SysmonLog|clr_optimization_v2.0.50727_32|clr_optimization_v4.0.303
                [6, 5, "^(automatic|automatic delayed)$", 3, ",", 1],
                [7, 2, "^Software Loopback Interface", 4, ",", 1],
```

```
                [8, 2, "^(In)?[Ll]oop[Bb]ack[0-9._]*$", 4, ",", 1],
                [9, 2, "^NULL[0-9.]*$", 4, ",", 1],
                [10, 2, "^[Ll]o[0-9.]*$", 4, ",", 1],
                [11, 2, "^[Ss]ystem$", 4, ",", 1],
                [12, 2, "^Nu[0-9.]*$", 4, ",", 1]
            ]
        },
        "settings": {
            "fields": ["name", "type", "value_str", "value_int"],
            "data": [
                ["autoreg_tls_accept", 2, "", 1],
                ["hk_history_global", 2, "", 0],
                ["snmptrap_logging", 2, "", 1],
                ["proxy_secrets_provider", 2, "", 0],
                ["hk_history", 1, "31d", 0],
                ["timeout_db_monitor", 1, "3s", 0],
                ["timeout_external_check", 1, "3s", 0],
                ["timeout_http_agent", 1, "3s", 0],
                ["timeout_simple_check", 1, "3s", 0],
                ["timeout_snmp_agent", 1, "3s", 0],
                ["timeout_ssh_agent", 1, "3s", 0],
                ["timeout_telnet_agent", 1, "3s", 0],
                ["timeout_zabbix_agent", 1, "3s", 0],
                ["timeout_browser", 1, "30s", 0]
            ]
        },
        "httptest": {
            "fields": ["httptestid", "name", "delay", "agent", "authentication", "http_user", "http_passwo
            "data": []
        },
        "httptestitem": {
            "fields": ["httptestitemid", "httptestid", "itemid", "type"],
            "data": []
        },
        "httptest_field": {
            "fields": ["httptest_fieldid", "httptestid", "type", "name", "value"],
            "data": []
        },
        "httpstep": {
            "fields": ["httpstepid", "httptestid", "name", "no", "url", "timeout", "posts", "required", "s
            "data": []
        },
        "httpstepitem": {
            "fields": ["httpstepitemid", "httpstepid", "itemid", "type"],
            "data": []
        },
        "httpstep_field": {
            "fields": ["httpstep_fieldid", "httpstepid", "type", "name", "value"],
            "data": []
        },
        "config_autoreg_tls": {
            "fields": ["autoreg_tlsid", "tls_psk_identity", "tls_psk"],
            "data": [
                [1, "", ""]
            ]
        }
    },
    "macro.secrets": {
        "AppID=zabbix_server&Query=Safe=passwordSafe;Object=zabbix": {
            "Content": "738"
        }
    },
```

```
    "config_revision": 2
}
```

proxy→server:

```
{
  "response": "success",
  "version": "8.2.0"
}
```

Data request

The `proxy data` request is used to obtain host interface availability, historical, discovery and autoregistration data from proxy. This request is sent every `ProxyDataFrequency` (server configuration parameter) seconds.

name	value type	description
server→proxy:		
request	*string*	'proxy data'
proxy→server:		
session	*string*	Data session token.
interface availability	*array*	*(optional)* Array of interface availability data objects.
interfaceid	*number*	Interface identifier.
available	*number*	Interface availability: **0**, *INTERFACE_AVAILABLE_UNKNOWN* - unknown **1**, *INTERFACE_AVAILABLE_TRUE* - available **2**, *INTERFACE_AVAILABLE_FALSE* - unavailable
error	*string*	Interface error message or empty string.
history data	*array*	*(optional)* Array of history data objects.
itemid	*number*	Item identifier.
clock	*number*	Item value timestamp (seconds).
ns	*number*	Item value timestamp (nanoseconds).
value	*string*	*(optional)* Item value.
id	*number*	Value identifier (ascending counter, unique within one data session).
timestamp	*number*	*(optional)* Timestamp of log type items.
source	*string*	*(optional)* Eventlog item source value.
severity	*number*	*(optional)* Eventlog item severity value.
eventid	*number*	*(optional)* Eventlog item eventid value.
state	*string*	*(optional)* Item state: **0**, *ITEM_STATE_NORMAL* **1**, *ITEM_STATE_NOTSUPPORTED*
lastlogsize	*number*	*(optional)* Last log size of log type items.
mtime	*number*	*(optional)* Modification time of log type items.
discovery data	*array*	*(optional)* Array of discovery data objects.
clock	*number*	Discovery data timestamp.
druleid	*number*	Discovery rule identifier.
dcheckid	*number*	Discovery check identifier or null for discovery rule data.

name	value type	description
type	*number*	Discovery check type: **-1** discovery rule data **0**, *SVC_SSH* - SSH service check **1**, *SVC_LDAP* - LDAP service check **2**, *SVC_SMTP* - SMTP service check **3**, *SVC_FTP* - FTP service check **4**, *SVC_HTTP* - HTTP service check **5**, *SVC_POP* - POP service check **6**, *SVC_NNTP* - NNTP service check **7**, *SVC_IMAP* - IMAP service check **8**, *SVC_TCP* - TCP port availability check **9**, *SVC_AGENT* - Zabbix agent **10**, *SVC_SNMPv1* - SNMPv1 agent **11**, *SVC_SNMPv2* - SNMPv2 agent **12**, *SVC_ICMPPING* - ICMP ping **13**, *SVC_SNMPv3* - SNMPv3 agent **14**, *SVC_HTTPS* - HTTPS service check **15**, *SVC_TELNET* - Telnet availability check
ip	*string*	Host IP address.
dns	*string*	Host DNS name.
port	*number*	*(optional)* Service port number.
key_	*string*	*(optional)* Item key for discovery check of type **9** *SVC_AGENT*
value	*string*	*(optional)* Value received from the service, can be empty for most of services.
status	*number*	*(optional)* Service status: **0**, *DOBJECT_STATUS_UP* - Service UP **1**, *DOBJECT_STATUS_DOWN* - Service DOWN
auto reg- is- tra- tion	*array*	*(optional)* Array of autoregistration data objects.
clock	*number*	Autoregistration data timestamp.
host	*string*	Host name.
ip	*string*	*(optional)* Host IP address.
dns	*string*	*(optional)* Resolved DNS name from IP address.
port	*string*	*(optional)* Host port.
host_metadata	*string*	*(optional)* Host metadata sent by agent (based on HostMetadata or HostMetadataItem agent configuration parameter).
tasks	*array*	*(optional)* Array of tasks.
type	*number*	Task type: **0**, *ZBX_TM_TASK_PROCESS_REMOTE_COMMAND_RESULT* - remote command result
status	*number*	Remote-command execution status: **0**, *ZBX_TM_REMOTE_COMMAND_COMPLETED* - remote command completed successfully **1**, *ZBX_TM_REMOTE_COMMAND_FAILED* - remote command failed
error	*string*	*(optional)* Error message.
parent_taskid	*number*	Parent task ID.
more	*number*	*(optional)* 1 - there are more history data to send.
clock	*number*	*(optional)* Data transfer timestamp (seconds).
ns	*number*	*(optional)* Data transfer timestamp (nanoseconds).
version	*string*	Proxy version (<major>.<minor>.<build>).
server→proxy:		
response	*string*	Request success information (’success’ or ’failed’).
tasks	*array*	*(optional)* Array of tasks.

name	value type	description
type	*number*	Task type: **1**, *ZBX_TM_TASK_PROCESS_REMOTE_COMMAND* - remote command
clock	*number*	Task creation time.
ttl	*number*	Time in seconds after which the task expires.
commandtype	*number*	Remote-command type: **0**, *ZBX_SCRIPT_TYPE_CUSTOM_SCRIPT* - use custom script **1**, *ZBX_SCRIPT_TYPE_IPMI* - use IPMI **2**, *ZBX_SCRIPT_TYPE_SSH* - use SSH **3**, *ZBX_SCRIPT_TYPE_TELNET* - use Telnet **4**, *ZBX_SCRIPT_TYPE_GLOBAL_SCRIPT* - use global script (currently functionally equivalent to custom script)
command	*string*	Remote command to execute.
execute_on	*number*	Execution target for custom scripts: **0**, *ZBX_SCRIPT_EXECUTE_ON_AGENT* - execute script on agent **1**, *ZBX_SCRIPT_EXECUTE_ON_SERVER* - execute script on server **2**, *ZBX_SCRIPT_EXECUTE_ON_PROXY* - execute script on proxy
port	*number*	*(optional)* Port for Telnet and SSH commands.
authtype	*number*	*(optional)* Authentication type for SSH commands.
username	*string*	*(optional)* User name for Telnet and SSH commands.
password	*string*	*(optional)* Password for Telnet and SSH commands.
publickey	*string*	*(optional)* Public key for SSH commands.
privatekey	*string*	*(optional)* Private key for SSH commands.
parent_taskid	*number*	Parent task ID.
hostid	*number*	Target host ID.

Example:

server→proxy:

```
{
  "request": "proxy data"
}
```

proxy→server:

```
{
    "session": "12345678901234567890123456789012"
    "interface availability": [
        {
            "interfaceid": 1,
            "available": 1,
            "error": ""
    },
        {
            "interfaceid": 2,
            "available": 2,
            "error": "Get value from agent failed: cannot connect to [[127.0.0.1]:10049]: [111] Connection
    },
        {
            "interfaceid": 3,
            "available": 1,
            "error": ""
    },
        {
            "interfaceid": 4,
            "available": 1,
            "error": ""
    }
    ],
```

```
    "history data":[
        {
            "itemid":"12345",
            "clock":1478609647,
            "ns":332510044,
            "value":"52956612",
            "id": 1
        },
        {
            "itemid":"12346",
            "clock":1478609647,
            "ns":330690279,
            "state":1,
            "value":"Cannot find information for this network interface in /proc/net/dev.",
            "id": 2
        }
    ],
    "discovery data":[
        {
            "clock":1478608764,
            "drule":2,
            "dcheck":3,
            "type":12,
            "ip":"10.3.0.10",
            "dns":"vdebian",
            "status":1
        },
        {
            "clock":1478608764,
            "drule":2,
            "dcheck":null,
            "type":-1,
            "ip":"10.3.0.10",
            "dns":"vdebian",
            "status":1
        }
    ],
    "auto registration":[
        {
            "clock":1478608371,
            "host":"Logger1",
            "ip":"10.3.0.1",
            "dns":"localhost",
            "port":"10050"
        },
        {
            "clock":1478608381,
            "host":"Logger2",
            "ip":"10.3.0.2",
            "dns":"localhost",
            "port":"10050"
        }
    ],
    "tasks":[
        {
            "type": 0,
            "status": 0,
            "parent_taskid": 10
        },
        {
            "type": 0,
            "status": 1,
```

```
            "error": "No permissions to execute task.",
            "parent_taskid": 20
        }
    ],
    "version":"8.2.0"
}
```

server→proxy:

```
{
  "response": "success",
  "tasks":[
      {
         "type": 1,
         "clock": 1478608371,
         "ttl": 600,
         "commandtype": 2,
         "command": "restart_service1.sh",
         "execute_on": 2,
         "port": 80,
         "authtype": 0,
         "username": "userA",
         "password": "password1",
         "publickey": "MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCqGKukO1De7zhZj6+H0qtjTkVxwTCpvKe",
         "privatekey": "lsuusFncCzWBQ7RKNUSesmQRMSGkVb1/3j+skZ6UtW+5u09lHNsj6tQ5QCqGKukO1De7zhd",
         "parent_taskid": 10,
         "hostid": 10070
      },
      {
         "type": 1,
         "clock": 1478608381,
         "ttl": 600,
         "commandtype": 1,
         "command": "restart_service2.sh",
         "execute_on": 0,
         "authtype": 0,
         "username": "",
         "password": "",
         "publickey": "",
         "privatekey": "",
         "parent_taskid": 20,
         "hostid": 10084
      }
  ]
}
```

Tasks request

The `proxy tasks` exchange handles delivery and acknowledgement of remote command tasks between the server and a passive proxy. When the server polls the proxy for task results, it sends an empty `proxy tasks` request at the interval defined by `ProxyDataFrequency`. The proxy responds with its current version and any pending task results (including timestamps). The server then acknowledges receipt by returning a `response` and may include new tasks for the proxy to execute.

name	value type	description
server→proxy:		
request	*string*	'proxy tasks'
proxy→server:		
version	*string*	Proxy version (<major>.<minor>.<build>).
clock	*number*	*(optional)* Data transfer timestamp (seconds).
ns	*number*	*(optional)* Data transfer timestamp (nanoseconds).
server→proxy:		
response	*string*	Request success information ('success' or 'failed').

Example:

server→proxy:

```
{
  "request":"proxy tasks"
}
```

proxy→server:

```
{
  "version":"7.0.0",
  "clock":1721059872,
  "ns":848141522
}
```

server→proxy:

```
{
  "response":"success"
}
```

Active proxy

Configuration request

The `proxy config` request is sent by active proxy to obtain proxy configuration data. This request is sent every `ProxyConfigFrequency` (proxy configuration parameter) seconds.

name	value type	description
proxy→server:		
request	*string*	'proxy config'
host	*string *	Proxy name.
version	*string*	Proxy version (<major>.<minor>.<build>).
session	*string*	Proxy configuration session token.
config_revision	*number*	Proxy configuration revision.
hostmap_revision	*number*	Current host-to-proxy mapping revision.
server→proxy:		
full_sync	*number*	1 - if full configuration data is sent, absent otherwise (optional).
data	*array*	Object of table data. Absent if configuration has not been changed (optional).
<table>	*object*	One or more objects with <table> data (optional, depending on changes).
fields	*array*	Array of field names.
-	*string*	Field name.
data	*array*	Array of rows.
-	*array*	Array of columns.
-	*string,number*	Column value with type depending on column type in database schema.
macro.secrets	*object*	Secret macro information, absent if there are no changes in vault macros (optional).
proxy_group	*string*	The name of the proxy group to which the proxy belongs.
config_revision	*number*	Configuration cache revision - sent with configuration data (optional).
del_hostids	*array*	Array of removed hostids (optional).
-	*number*	Host identifier.
del_macro_hostids	*array*	Array of hostids with all macros removed (optional).
-	*number*	Host identifier.

Example:

proxy→server:

```
{
  "request": "proxy config",
```

```
  "host": "Zabbix proxy",
  "version":"8.2.0",
  "session": "fd59a09ff4e9d1fb447de1f04599bcf6",
  "config_revision": 0
}
```


server→proxy:


```
{
"full_sync": 1,
"data": {
"hosts": {
"fields": ["hostid", "host", "status", "ipmi_authtype", "ipmi_privilege", "ipmi_username", "ipmi_password"
"data": [
[10084, "Zabbix server", 0, -1, 2, "", "", "Zabbix server", 1, 1, "", "", "", ""]
]
},
"interface": {
"fields": ["interfaceid", "hostid", "main", "type", "useip", "ip", "dns", "port", "available"],
"data": [
[1, 10084, 1, 1, 1, "127.0.0.1", "", "10053", 1]
]
},
"interface_snmp": {
"fields": ["interfaceid", "version", "bulk", "community", "securityname", "securitylevel", "authpassphrase
"data": []
},
"host_inventory": {
"fields": ["hostid", "type", "type_full", "name", "alias", "os", "os_full", "os_short", "serialno_a", "ser
"data": [
[10084, "", "", ..., "56.95387", "24.22067", ..., ""]
]
},
"items": {
"fields": ["itemid", "type", "snmp_oid", "hostid", "key_", "delay", "history", "status", "value_type", "tr
"data": [
[44161, 7, "", 10084, "agent.hostmetadata", "10s", "90d", 0, 1, "", "", "", "", 0, "", "", "", "", 0, null
[44162, 0, "", 10084, "agent.ping", "10s", "90d", 0, 3, "", "", "", "", 0, "", "", "", "", 0, 1, 0, "", nu
]
},
"item_rtdata": {
"fields": ["itemid", "lastlogsize", "mtime"],
"data": [
[44161, 0, 0],
[44162, 0, 0]
]
},
"item_preproc": {
"fields": ["item_preprocid", "itemid", "step", "type", "params", "error_handler", "error_handler_params"],
"data": []
},
"item_parameter": {
"fields": ["item_parameterid", "itemid", "name", "value"],
"data": []
},
"globalmacro": {
"fields": ["globalmacroid", "macro", "value", "type"],
"data": [
[2, "{$SNMP_COMMUNITY}", "public", 0]
]
},
"hosts_templates": {
"fields": ["hosttemplateid", "hostid", "templateid", "link_type"],
```

```
"data": []
},
"hostmacro": {
"fields": ["hostmacroid", "hostid", "macro", "value", "type", "automatic"],
"data": [
[5676, 10084, "{$M}", "AppID=zabbix_server&Query=Safe=passwordSafe;Object=zabbix:Content", 2, 0]
]
},
"drules": {
"fields": ["druleid", "name", "iprange", "delay"],
"data": [
[2, "Local network", "127.0.0.1", "10s"]
]
},
"dchecks": {
"fields": ["dcheckid", "druleid", "type", "key_", "snmp_community", "ports", "snmpv3_securityname", "snmpv
"data": [
[2, 2, 9, "system.uname", "", "10052", "", 0, "", "", 0, 0, 0, "", 1, 0]
]
},
"regexps": {
"fields": ["regexpid", "name"],
"data": [
[1, "File systems for discovery"],
[2, "Network interfaces for discovery"],
[3, "Storage devices for SNMP discovery"],
[4, "Windows service names for discovery"],
[5, "Windows service startup states for discovery"]
]
},
"expressions": {
"fields": ["expressionid", "regexpid", "expression", "expression_type", "exp_delimiter", "case_sensitive"]
"data": [
[1, 1, "^(btrfs|ext2|ext3|ext4|reiser|xfs|ffs|ufs|jfs|jfs2|vxfs|hfs|apfs|refs|ntfs|fat32|zfs)$", 3, ",", 0
[3, 3, "^(Physical memory|Virtual memory|Memory buffers|Cached memory|Swap space)$", 4, ",", 1],
[5, 4, "^(MMCSS|gupdate|SysmonLog|clr_optimization_v2.0.50727_32|clr_optimization_v4.0.30319_32)$", 4, ","
[6, 5, "^(automatic|automatic delayed)$", 3, ",", 1],
[7, 2, "^Software Loopback Interface", 4, ",", 1],
[8, 2, "^(In)?[Ll]oop[Bb]ack[0-9._]*$", 4, ",", 1],
[9, 2, "^NULL[0-9.]*$", 4, ",", 1],
[10, 2, "^[Ll]o[0-9.]*$", 4, ",", 1],
[11, 2, "^[Ss]ystem$", 4, ",", 1],
[12, 2, "^Nu[0-9.]*$", 4, ",", 1]
]
},
"settings": {
"fields": ["name", "type", "value_str", "value_int"],
"data": [
["autoreg_tls_accept", 2, "", 1],
["hk_history_global", 2, "", 0],
["snmptrap_logging", 2, "", 1],
["proxy_secrets_provider", 2, "", 0],
["hk_history", 1, "31d", 0],
["timeout_db_monitor", 1, "3s", 0],
["timeout_external_check", 1, "3s", 0],
["timeout_http_agent", 1, "3s", 0],
["timeout_simple_check", 1, "3s", 0],
["timeout_snmp_agent", 1, "3s", 0],
["timeout_ssh_agent", 1, "3s", 0],
["timeout_telnet_agent", 1, "3s", 0],
["timeout_zabbix_agent", 1, "3s", 0],
["timeout_browser", 1, "30s", 0]
```

```
]
},
"httptest": {
"fields": ["httptestid", "name", "delay", "agent", "authentication", "http_user", "http_password", "hostid
"data": []
},
"httptestitem": {
"fields": ["httptestitemid", "httptestid", "itemid", "type"],
"data": []
},
"httptest_field": {
"fields": ["httptest_fieldid", "httptestid", "type", "name", "value"],
"data": []
},
"httpstep": {
"fields": ["httpstepid", "httptestid", "name", "no", "url", "timeout", "posts", "required", "status_codes"
"data": []
},
"httpstepitem": {
"fields": ["httpstepitemid", "httpstepid", "itemid", "type"],
"data": []
},
"httpstep_field": {
"fields": ["httpstep_fieldid", "httpstepid", "type", "name", "value"],
"data": []
},
"config_autoreg_tls": {
"fields": ["autoreg_tlsid", "tls_psk_identity", "tls_psk"],
"data": [
[1, "", ""]
]
}
},
"macro.secrets": {
"AppID=zabbix_server&Query=Safe=passwordSafe;Object=zabbix": {
"Content": "738"
}
},
"config_revision": 2
}
```


Data request

The `proxy data` request is sent by proxy to provide host interface availability, history, discovery and autoregistration data. This request is sent every `DataSenderFrequency` (proxy configuration parameter) seconds. Note that active proxy will still poll Zabbix server every second for remote command tasks (with an empty `proxy data` request).

name	value type	description
proxy→server:		
request	*string*	'proxy data'
host	*string*	Proxy name.
session	*string*	Data session token.
interface availability	*array*	*(optional)* Array of interface availability data objects.
interfaceid	*number*	Interface identifier.
available	*number*	Interface availability: **0**, *INTERFACE_AVAILABLE_UNKNOWN* - unknown **1**, *INTERFACE_AVAILABLE_TRUE* - available **2**, *INTERFACE_AVAILABLE_FALSE* - unavailable
error	*string*	Interface error message or empty string.

name	value type	description
history	*array*	*(optional)* Array of history data objects.
data		
itemid	*number*	Item identifier.
clock	*number*	Item value timestamp (seconds).
ns	*number*	Item value timestamp (nanoseconds).
value	*string*	*(optional)* Item value.
id	*number*	Value identifier (ascending counter, unique within one data session).
timestamp	*number*	*(optional)* Timestamp of log type items.
source	*string*	*(optional)* Eventlog item source value.
severity	*number*	*(optional)* Eventlog item severity value.
eventid	*number*	*(optional)* Eventlog item eventid value.
state	*string*	*(optional)* Item state: **0**, *ITEM_STATE_NORMAL* **1**, *ITEM_STATE_NOTSUPPORTED*
lastlogsize	*number*	*(optional)* Last log size of log type items.
mtime	*number*	*(optional)* Modification time of log type items.
discovery	*array*	*(optional)* Array of discovery data objects.
data		
clock	*number*	Discovery data timestamp.
druleid	*number*	Discovery rule identifier.
dcheckid	*number*	Discovery check identifier or null for discovery rule data.
type	*number*	Discovery check type: **-1** discovery rule data **0**, *SVC_SSH* - SSH service check **1**, *SVC_LDAP* - LDAP service check **2**, *SVC_SMTP* - SMTP service check **3**, *SVC_FTP* - FTP service check **4**, *SVC_HTTP* - HTTP service check **5**, *SVC_POP* - POP service check **6**, *SVC_NNTP* - NNTP service check **7**, *SVC_IMAP* - IMAP service check **8**, *SVC_TCP* - TCP port availability check **9**, *SVC_AGENT* - Zabbix agent **10**, *SVC_SNMPv1* - SNMPv1 agent **11**, *SVC_SNMPv2* - SNMPv2 agent **12**, *SVC_ICMPPING* - ICMP ping **13**, *SVC_SNMPv3* - SNMPv3 agent **14**, *SVC_HTTPS* - HTTPS service check **15**, *SVC_TELNET* - Telnet availability check
ip	*string*	Host IP address.
dns	*string*	Host DNS name.
port	*number*	*(optional)* Service port number.
key_	*string*	*(optional)* Item key for discovery check of type **9** *SVC_AGENT*
value	*string*	*(optional)* Value received from the service, can be empty for most services.
status	*number*	*(optional)* Service status: **0**, *DOBJECT_STATUS_UP* - Service UP **1**, *DOBJECT_STATUS_DOWN* - Service DOWN
autoregistration	*array*	*(optional)* Array of autoregistration data objects.
clock	*number*	Autoregistration data timestamp.
host	*string*	Host name.
ip	*string*	*(optional)* Host IP address.
dns	*string*	*(optional)* Resolved DNS name from IP address.
port	*string*	*(optional)* Host port.
host_metadata	*string*	*(optional)* Host metadata sent by agent (based on HostMetadata or HostMetadataItem agent configuration parameter).
tasks	*array*	*(optional)* Array of tasks.
type	*number*	Task type: **0**, *ZBX_TM_TASK_PROCESS_REMOTE_COMMAND_RESULT* - remote command result

name	value type	description
status	*number*	Remote-command execution status: **0**, *ZBX_TM_REMOTE_COMMAND_COMPLETED* - remote command completed successfully **1**, *ZBX_TM_REMOTE_COMMAND_FAILED* - remote command failed
error	*string*	*(optional)* Error message.
parent_taskid	*number*	Parent task ID.
more	*number*	*(optional)* 1 - there are more history data to send
clock	*number*	*(optional)* Data transfer timestamp (seconds).
ns	*number*	*(optional)* Data transfer timestamp (nanoseconds).
version	*string*	Proxy version (<major>.<minor>.<build>).
server→proxy:		
response	*string*	Request success information ('success' or 'failed').
upload	*string*	Upload control for historical data (history, autoregistration, host availability, network discovery). Possible values: **enabled** - normal operation **disabled** - server is not accepting data (possibly due to internal cache over limit)
tasks	*array*	*(optional)* Array of tasks.
type	*number*	Task type: **1**, *ZBX_TM_TASK_PROCESS_REMOTE_COMMAND* - remote command
clock	*number*	Task creation time.
ttl	*number*	Time in seconds after which the task expires.
commandtype	*number*	Remote-command type: **0**, *ZBX_SCRIPT_TYPE_CUSTOM_SCRIPT* - use custom script **1**, *ZBX_SCRIPT_TYPE_IPMI* - use IPMI **2**, *ZBX_SCRIPT_TYPE_SSH* - use SSH **3**, *ZBX_SCRIPT_TYPE_TELNET* - use Telnet **4**, *ZBX_SCRIPT_TYPE_GLOBAL_SCRIPT* - use global script (currently functionally equivalent to custom script)
command	*string*	Remote command to execute.
execute_on	*number*	Execution target for custom scripts: **0**, *ZBX_SCRIPT_EXECUTE_ON_AGENT* - execute script on agent **1**, *ZBX_SCRIPT_EXECUTE_ON_SERVER* - execute script on server **2**, *ZBX_SCRIPT_EXECUTE_ON_PROXY* - execute script on proxy
port	*number*	*(optional)* Port for Telnet and SSH commands.
authtype	*number*	*(optional)* Authentication type for SSH commands.
username	*string*	*(optional)* User name for Telnet and SSH commands.
password	*string*	*(optional)* Password for Telnet and SSH commands.
publickey	*string*	*(optional)* Public key for SSH commands.
privatekey	*string*	*(optional)* Private key for SSH commands.
parent_taskid	*number*	Parent task ID.
hostid	*number*	Target host ID.

Example:

proxy→server:

```
{
"request": "proxy data",
"host": "Zabbix proxy",
"session": "818cdd1b537bdc5e50c09ed4969235b6",
"interface availability": [{
"interfaceid": 1,
"available": 1,
"error": ""
}],
```

```
"history data": [{
"id": 1114,
"itemid": 44162,
"clock": 1665730632,
"ns": 798953105,
"value": "1"
}, {
"id": 1115,
"itemid": 44161,
"clock": 1665730633,
"ns": 811684663,
"value": "58"
}],
"auto registration": [{
"clock": 1665730633,
"host": "Zabbix server",
"ip": "127.0.0.1",
"dns": "localhost",
"port": "10053",
"host_metadata": "58",
"tls_accepted": 1
}],
"discovery data": [{
"clock": 1665732232,
"drule": 2,
"dcheck": 2,
"ip": "127.0.0.1",
"dns": "localhost",
"port": 10052,
"status": 1
}, {
"clock": 1665732232,
"drule": 2,
"dcheck": null,
"ip": "127.0.0.1",
"dns": "localhost",
"status": 1
}],
"host data": [{
"hostid": 10084,
"active_status": 1
}],
"tasks": [{
"type": 3,
"clock": 1665730985,
"ttl": 0,
"status": -1,
"info": "Remote commands are not enabled",
"parent_taskid": 3
}],
"version": "8.2.0",
"clock": 1665730643,
"ns": 65389964
}
```


server→proxy:


```
{
"upload": "enabled",
"response": "success",
"tasks": [{
"type": 2,
"clock": 1665730986,
```

```
"ttl": 600,
"commandtype": 0,
"command": "ping -c 3 127.0.0.1; case $? in [01]) true;; *) false;; esac",
"execute_on": 2,
"port": 0,
"authtype": 0,
"username": "",
"password": "",
"publickey": "",
"privatekey": "",
"alertid": 0,
"parent_taskid": 4,
"hostid": 10084
}]
}
```

Configuration parameters: Zabbix proxy

Overview

Configuration parameters allow customizing Zabbix proxy behavior.

Specify parameter values in the Zabbix proxy configuration file (`zabbix_proxy.conf`). All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
AllowRoot	Allow proxy to run as `root`.
AllowUnsupportedDBVersions	Allow proxy to work with unsupported database versions.
APMListenIP	Specify the IP address that the APM collector should listen on for incoming telemetry data.
APMListenPort	Specify the port that the APM collector should listen on for incoming telemetry data.
APMTLSCAFile	Specify the full pathname of the file containing the top-level CA(s) certificates for verifying the certificates of applications or OpenTelemetry collectors that send data to the APM collector.
APMTLSCertFile	Specify the full pathname of the file containing the APM collector certificate or certificate chain.
APMTLSKeyFile	Specify the full pathname of the file containing the APM collector private key.
CacheSize	Set the configuration cache size.
ConfigFrequency	This parameter is **deprecated** (use ProxyConfigFrequency instead). Set the frequency of retrieving configuration data from Zabbix server.
DataSenderFrequency	Set the frequency of sending collected data to Zabbix server.
DBHost	Specify the database host (or socket directory).
DBName	Specify the database name or path to the database file for SQLite3.
DBPassword	Specify the database password.
DBPort	Specify the port of the Zabbix database host.
DBSchema	Specify the database schema name. Used for PostgreSQL.
DBSocket	Specify path to the MySQL socket file.
DBUser	Specify the database user.
DBTLSConnect	Specify a value to enforce a TLS connection to the database.
DBTLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificates for database certificate verification.
DBTLSCertFile	Specify full pathname of the file containing the Zabbix proxy certificate for authenticating to database.
DBTLSKeyFile	Specify full pathname of the file containing the private key for authenticating to database.
DBTLSCipher	Specify the list of encryption ciphers that Zabbix proxy permits for TLS protocols up through TLS v1.2. Supported only for MySQL.
DBTLSCipher13	Specify the list of encryption ciphersuites that Zabbix proxy permits for the TLS v1.3 protocol. Supported only for MySQL, starting from version 8.0.16.
DebugLevel	Set the debug level.
EnableRemoteCommands	Allow remote commands from Zabbix server.
ExternalScripts	Specify the external script location.

Parameter	Description
Fping6Location	Specify the fping6 location.
FpingLocation	Specify the fping location.
HistoryCacheSize	Set the history cache size.
HistoryIndexCacheSize	Set the history index cache size.
Hostname	Specify the proxy name.
HostnameItem	Specify the item used for setting Hostname (if undefined).
HousekeepingFrequency	Set the housekeeping procedure frequency (in hours).
Include	Specify a directory or individual files to be included in the configuration file.
JavaGateway	Specify the Zabbix Java gateway IP address (or hostname).
JavaGatewayPort	Specify the Zabbix Java gateway port.
ListenBacklog	Set the maximum number of pending connections in the TCP queue.
ListenIP	Specify the trapper listen IP address.
ListenPort	Specify the trapper listen port.
LoadModule	Specify the module to load at proxy startup.
LoadModulePath	Specify full path to the location of proxy modules.
LogFile	Specify the proxy log file.
LogFileSize	Set the maximum log file size.
LogRemoteCommands	Enable logging of executed shell commands as warnings.
LogSlowQueries	Set the maximum duration a database query may take before being logged (in milliseconds).
LogType	Specify the log output type.
MaxConcurrentChecksPerPoller	Set the maximum number of asynchronous checks that can be executed at once by each HTTP agent poller, agent poller, or SNMP poller.
PidFile	Specify the PID file.
ProxyBufferMode	Specify the storage mechanism for history, discovery, and autoregistration data (disk/memory/hybrid).
ProxyConfigFrequency	Set the frequency of retrieving configuration data from Zabbix server in seconds.
ProxyLocalBuffer	Set the duration of keeping data locally (in hours), even if the data have already been synced with the server.
ProxyMemoryBufferAge	Set the maximum age of data in the proxy memory buffer in seconds.
ProxyMemoryBufferSize	Set the shared memory cache size for collected history, discovery, and auto registration data.
ProxyMode	Set the proxy operating mode (active/passive).
ProxyOfflineBuffer	Set the duration of keeping data (in hours) in case of no connectivity with Zabbix server.
Server	If ProxyMode is set to `active`: specify Zabbix server IP address or DNS name (address:port) or cluster (address:port;address2:port) to get configuration data from and send data to. If ProxyMode is set to `passive`: specify IP addresses or DNS names of Zabbix server.
SNMPTrapperFile	Specify the temporary file used for passing data from the SNMP trap daemon to the proxy.
SocketDir	Specify the directory for storing IPC sockets used by internal Zabbix services.
SourceIP	Specify the source IP address.
SSHKeyLocation	Specify the location of public and private keys for SSH checks and actions.
SSLCertLocation	Specify the location of SSL-client certificate files for client authentication.
SSLKeyLocation	Specify the location of SSL private key files for client authentication.
SSLCALocation	Specify the location of certificate authority (CA) files for SSL server certificate verification.
StartAgentPollers	Set the number of pre-forked instances of asynchronous Zabbix agent pollers.
StartAPMCollector	Set to 1 to start the APM collector.
StartBrowserPollers	Set the number of pre-forked instances of browser item pollers.
StartDBSyncers	Set the number of pre-forked instances of history syncers.
StartDiscoverers	Set the number of pre-forked instances of discovery workers.
StartHTTPAgentPollers	Set the number of pre-forked instances of asynchronous HTTP agent pollers.
StartHTTPPollers	Set the number of pre-forked instances of HTTP pollers.
StartIPMIPollers	Set the number of pre-forked instances of IPMI pollers.
StartJavaPollers	Set the number of pre-forked instances of Java pollers.
StartODBCPollers	Set the number of pre-forked instances of ODBC pollers.
StartPingers	Set the number of pre-forked instances of ICMP pingers.
StartPollersUnreachable	Set the number of pre-forked instances of pollers for unreachable hosts (including IPMI and Java).
StartPollers	Set the number of pre-forked instances of pollers.
StartPreprocessors	Set the number of pre-started instances of preprocessing workers.
StartSNMPPollers	Set the number of pre-forked instances of asynchronous SNMP pollers.
StartSNMPTrapper	Set to ”1” to start an SNMP trapper process.
StartTelemetryQueryPollers	Set the number of pre-forked instances of asynchronous Telemetry query pollers.
StartTrappers	Set the number of pre-forked instances of trappers.
StartVMwareCollectors	Set the number of pre-forked VMware collector instances.

Parameter	Description
StatsAllowedIP	Specify IP addresses or DNS names of external Zabbix instances that are allowed to request stats.
TelemetryProvider	Set the telemetry data source.
Timeout	Specify how long to wait (in seconds) for establishing connection and exchanging data with Zabbix proxy, agent, web service, and for SNMP checks (except SNMP `walk[OID]` and `get[OID]` items).
TLSAccept	Specify the level of encryption required for incoming connections from Zabbix server.
TLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.
TLSCertFile	Specify full pathname of the file containing the server certificate or certificate chain, used for encrypted communications between Zabbix components.
TLSCipherAll	Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.
TLSCipherAll13	Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.
TLSCipherCert	Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for certificate-based encryption.
TLSCipherCert13	Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for certificate-based encryption.
TLSCipherPSK	Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for PSK-based encryption.
TLSCipherPSK13	Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for PSK-based encryption.
TLSConnect	Specify the level of encryption required for connections to Zabbix server.
TLSCRLFile	Specify full pathname of the file containing revoked certificates. This parameter is used for encrypted communications between Zabbix components.
TLSKeyFile	Specify full pathname of the file containing the proxy private key, used for encrypted communications between Zabbix components.
TLSListen	Require TLS-only connections on the trapper socket.
TLSPSKFile	Specify full pathname of the file containing the proxy pre-shared key, used for encrypted communications with Zabbix server.
TLSPSKIdentity	Specify the pre-shared key identity string, used for encrypted communications with Zabbix server.
TLSServerCertIssuer	Specify the allowed server certificate issuer.
TLSServerCertSubject	Specify the allowed server certificate subject.
TmpDir	Specify the temporary directory.
TrapperTimeout	Specify timeout in seconds for: - retrieval of configuration data from the Zabbix server; - global script execution or remote command execution.
UnavailableDelay	Set the frequency of checking host availability during the unavailability period.
UnreachableDelay	Set the frequency of checking host availability during the unreachability period.
UnreachablePeriod	Set after how many seconds of unreachability treat host as unavailable.
User	Drop privileges to a specific, existing user on the system.
Vault	Specify the vault provider.
VaultAppRoleID	Specify the HashiCorp vault AppRole ID.
VaultAppSecretID	Specify the HashiCorp vault AppRole secret ID.
VaultDBPath	Specify the location for retrieving database credentials by keys.
VaultPrefix	Specify custom prefix for the vault path or query.
VaultTLSCertFile	Specify the name of the SSL certificate file used for client authentication.
VaultTLSKeyFile	Specify the name of the SSL private key file used for client authentication.
VaultToken	Specify the HashiCorp vault authentication token.
VaultURL	Specify the vault server URL.
VMwareCacheSize	Set the shared memory size for storing VMware data.
VMwareFrequency	Set the frequency of data gathering from a single VMware service in seconds.
VMwarePerfFrequency	Set the frequency of retrieving performance counter statistics from a single VMware service in seconds.
VMwareTimeout	Set the maximum number of seconds a `vmware collector` will wait for a response from VMware service.
WebDriverURL	Specify the WebDriver interface URL.

Note that:

- The default values reflect daemon defaults, not the values in the shipped configuration files.
- Values support environment variables.
- Zabbix supports configuration files only in UTF-8 encoding without BOM.
- Comments starting with # are only supported in the beginning of the line.

Parameter details

AllowRoot

Allow proxy to run as `root`. If disabled and the proxy is started by `root`, the proxy will try to switch to the `zabbix` user instead. Has no effect if started under a regular user.

Default: 0
 Values: 0 - do not allow; 1 - allow

AllowUnsupportedDBVersions

Allow proxy to work with unsupported database versions.

Default: 0
 Values: 0 - do not allow; 1 - allow

APMListenIP

Specify the IP address that the APM collector should listen on for incoming telemetry data.
 If not specified, the APM collector will listen on all network interfaces.

Default: `0.0.0.0`

APMListenPort

Specify the port that the APM collector should listen on for incoming telemetry data.
 The default port is the standard OpenTelemetry (OTLP/gRPC) port, used by default by OpenTelemetry SDKs and collectors.

Default: 4317
 Range: 1024-32767

APMTLSCAFile

Specify the full pathname of the file containing the top-level CA(s) certificates for verifying the certificates of applications or OpenTelemetry collectors that send data to the APM collector.
 If this parameter is set, the APM collector accepts only encrypted connections; otherwise, it accepts only unencrypted connections.

APMTLSCertFile

Specify the full pathname of the file containing the APM collector certificate or certificate chain.
 This parameter must be set together with `APMTLSKeyFile`, and only if `APMTLSCAFile` is set.

APMTLSKeyFile

Specify the full pathname of the file containing the APM collector private key.
 This parameter must be set together with `APMTLSCertFile`, and only if `APMTLSCAFile` is set.

CacheSize

Set the configuration cache size, in bytes. The shared memory size for storing host and item data.

Default: `32M`
 Range: 128K-64G

ConfigFrequency

This parameter is **deprecated** (use `ProxyConfigFrequency` instead).
Set the frequency of retrieving configuration data from Zabbix server (in seconds).
Active proxy parameter. Ignored for passive proxies (see `ProxyMode` parameter).

Default: 3600
 Range: 1-604800

DataSenderFrequency

Set the frequency of sending collected data from proxy to Zabbix server. Note that an active proxy will still poll Zabbix server every second for remote command tasks.
Active proxy parameter. Ignored for passive proxies (see `ProxyMode` parameter).

Default: 1
 Range: 1-3600

DBHost

Specify the database host (or socket directory).

For MySQL:

- `localhost` or an empty string uses the default UNIX-domain socket.

For PostgreSQL:

- `localhost` is resolved via DNS (typically to `127.0.0.1`).
- An empty string uses the default UNIX-domain socket.
- A path (for example, `/var/run/pgbouncer`) results in using the UNIX socket at that path.
- A comma-separated list can contain multiple `host:port` values. Each host is tried in order until a read-write connection is established. Example: `DBHost=localhost:5431,127.0.0.1:20051,zabbix.domain,[::1]:30051,[12fc::1]`

Default: `localhost`

DBName

Specify the database name or path to the database file for SQLite3 (the multi-process architecture of Zabbix does not allow the use of an in-memory database, e.g., `:memory:`, `file::memory:?cache=shared` or `file:memdb1?mode=memory&cache=shared`).
Warn
Do not attempt to use the same database the Zabbix server is using.

Mandatory: Yes

DBPassword

Specify the database password. Comment this line if no password is used. Ignored for SQLite.

DBPort

Specify the port of the Zabbix database host (see DBHost). Ignored for SQLite.
 DBSocket and DBPort are mutually exclusive in proxy configuration. Specify only one, or leave both undefined.

Default for MySQL: 3306

Default for PostgreSQL: 5432

Range: 1024-65535

DBSchema

Specify the database schema name. Used for PostgreSQL.

DBSocket

Specify path to the MySQL socket file.
DBSocket and DBPort are mutually exclusive in proxy configuration. Specify only one, or leave both undefined.

DBUser

Specify the database user. Ignored for SQLite.

DBTLSConnect

Specify a value to enforce a TLS connection to the database: - *required* - connect using TLS; - *verify_ca* - connect using TLS and verify certificate; - *verify_full* - connect using TLS, verify certificate and verify that database identity specified by DBHost matches its certificate
On MySQL starting from 5.7.11 and PostgreSQL the following values are supported: ”required”, ”verify”, ”verify_full”.
On MariaDB starting from version 10.2.6 ”required” and ”verify_full” values are supported.
By default, not set to any option and the behavior depends on database configuration.

DBTLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificates for database certificate verification.

Mandatory: no (yes, if DBTLSConnect set to *verify_ca* or *verify_full*)

DBTLSCertFile

Specify full pathname of the file containing the Zabbix proxy certificate for authenticating to database.

DBTLSKeyFile

Specify full pathname of the file containing the private key for authenticating to the database.

DBTLSCipher

Specify the list of encryption ciphers that Zabbix proxy permits for TLS protocols up through TLS v1.2. Supported only for MySQL.

DBTLSCipher13

Specify the list of encryption ciphersuites that Zabbix proxy permits for the TLS v1.3 protocol. Supported only for MySQL, starting from version 8.0.16.

DebugLevel

Set the debug level: - *0* - basic information about starting and stopping of Zabbix processes; - *1* - critical information; - *2* - error information; - *3* - warnings; - *4* - for debugging (produces lots of information); - *5* - extended debugging (produces even more information).

Default: 3
 Range: 0-5

EnableRemoteCommands

Allow remote commands from Zabbix server.

Default: 0
 Values: 0 - do not allow; 1 - allow

ExternalScripts

Specify the external script location (depends on the `datadir` compile-time installation variable).

Default: `/usr/local/share/zabbix/externalscripts`

Fping6Location

Specify the fping6 location. Make sure that the fping6 binary has root ownership and the SUID flag set. Make empty (`Fping6Location=`) if your `fping` utility is capable to process IPv6 addresses.

Default: `/usr/sbin/fping6`

FpingLocation

Specify the fping location. Make sure that the fping binary has root ownership and the SUID flag set.

Default: `/usr/sbin/fping`

HistoryCacheSize

Set the history cache size, in bytes. The shared memory size for storing history data.

Default: `16M`
 Range: 128K-16G

HistoryIndexCacheSize

Set the history index cache size, in bytes. The shared memory size for indexing the history data stored in history cache. The index cache size needs roughly 100 bytes to cache one item.

Default: `4M`
 Range: 128K-16G

Hostname

Specify a unique, case-sensitive proxy name. Make sure the proxy name is known to the server.
Allowed characters: alphanumeric, '.', ' ', '_' and '-'. Maximum length: 128

Default: Set by HostnameItem

HostnameItem

Specify an item for setting `Hostname` (if undefined).
 This item will be run on the proxy similarly as on an agent.
 Ignored if `Hostname` is set.
Does not support user parameters, performance counters, or aliases; does support `system.run[]`.

Default: `system.hostname`

HousekeepingFrequency

Set the housekeeping procedure frequency (in hours). Housekeeping is removing outdated information from the database.
Note: To lower load on proxy startup housekeeping is postponed for 30 minutes after proxy start. Thus, if `HousekeepingFrequency` is 1, the very first housekeeping procedure after proxy start will run after 30 minutes, and will repeat every hour thereafter.
It is possible to disable automatic housekeeping by setting HousekeepingFrequency to 0. In this case the housekeeping procedure can only be started by `housekeeper_execute` runtime control option.
See also runtime control options and details on the housekeeping procedure.

Default: 1
 Range: 0-24

Include

Specify a directory or individual files to be included in the configuration file.
 During Zabbix installation, this parameter is set to `/usr/local/etc`, unless you change this during compile time.

General rules for the `Include` parameter:

- The path can be relative to the `zabbix_proxy.conf` file location.
- You can specify multiple `Include` parameters in one configuration file.
- Zabbix processes the `Include` parameter as soon as it reads it, so consider placing it at the end of the configuration file.

Rules for included files:

- Each file must be readable.
- Each file must have correct syntax. Otherwise, the component will not start.

Rules for including directories or file patterns:

- Do not assume any order of inclusion (e.g., files are not included in alphabetical order).
- Do not define the same parameter in more than one included file (e.g., to override a general setting with a specific one), since you cannot predict which value will be used.
- Watch out for backup copies that some text editors create automatically (e.g., editing `include/my_specific.conf` may create `include/my_specific.conf.BAK`), since both files will be included.

To include only relevant files from a directory, use the wildcard (*) for pattern matching.

Example:

```
Include=/absolute/path/to/config/files/*.conf
```

JavaGateway

Specify the IP address or hostname of Zabbix Java gateway. Only required if Java pollers are started.

JavaGatewayPort

Specify the Zabbix Java gateway port.

Default: 10052
 Range: 1024-32767

ListenBacklog

Set the maximum number of pending connections in the TCP queue.
The default value is a hard-coded constant, which depends on the system.
The maximum supported value depends on the system, too high values may be silently truncated to the 'implementation-specified maximum'.

Default: SOMAXCONN
 Range: 0 - INT_MAX

ListenIP

Specify the trapper listen IP address. A comma-delimited list.
Trapper will listen on all network interfaces if this parameter is not specified.

Default: 0.0.0.0

ListenPort

Specify the trapper listen port.

Default: 10051
 Range: 1024-32767

LoadModule

Specify the module to load at proxy startup. Modules are used to extend the functionality of the proxy. The module must be located in the directory specified by LoadModulePath or the path must precede the module name. If the preceding path is absolute (starts with /) then LoadModulePath is ignored.
 Formats: - LoadModule=<module.so> - LoadModule=<path/module.so> - LoadModule=</abs_path/module.so>

Multiple LoadModule parameters can be specified.

LoadModulePath

The full path to the location of proxy modules. The default depends on compilation options.

LogFile

Specify the proxy log file.

Mandatory: Yes, if LogType is set to file; otherwise, no

LogFileSize

Set the maximum log file size (in MB).
0 - disable automatic log rotation.
Note: If the log file size limit is reached and file rotation fails, for whatever reason, the existing log file is truncated and started anew.

Default: 1
 Range: 0-1024

LogRemoteCommands

Enable the logging of executed shell commands as warnings.

Default: 0
 Values: 0 - disabled, 1 - enabled

LogSlowQueries

Set the maximum duration a database query may take before being logged (in milliseconds).
0 - don't log slow queries.
This option becomes enabled starting with `DebugLevel=3`.

Default: 0
 Range: 0-3600000

LogType

Specify the log output type: - *file* - write log to the file specified by `LogFile` parameter; - *system* - write log to syslog; - *console* - write log to standard output.

Default: `file`

MaxConcurrentChecksPerPoller

Set the maximum number of asynchronous checks that can be executed at once by each HTTP agent poller, agent poller, or SNMP poller. See `StartHTTPAgentPollers`, `StartAgentPollers`, and `StartSNMPPollers`.

Default: 1000
 Range: 1-1000

PidFile

Specify the PID file.

Default: `/tmp/zabbix_proxy.pid`

ProxyBufferMode

Specify the storage mechanism for history, discovery, and autoregistration data: - *disk* - data is stored in database and uploaded from database; - *memory* - data is stored in memory and uploaded from memory. If buffer runs out of memory the old data will be discarded. On shutdown the buffer is discarded. - *hybrid* - the proxy buffer normally works like in the memory mode until it runs out of memory or the oldest record exceeds the configured age. If that happens the buffer is flushed to database and it works like in disk mode until all data have been uploaded and it starts working with memory again. On shutdown the memory buffer is flushed to database.

See also: Proxy memory buffer.

Default: `disk`
 Values: disk; memory; hybrid

ProxyConfigFrequency

Set the frequency of retrieving configuration data from Zabbix server in seconds.
Active proxy parameter. Ignored for passive proxies (see `ProxyMode` parameter).

Default: 10
 Range: 1-604800

ProxyLocalBuffer

Set the duration of keeping data locally (in hours), even if the data have already been synced with the server..
This parameter may be used if local data will be used by third-party applications.

Default: 0
 Range: 0-720

ProxyMemoryBufferAge

Set the maximum age of data in the proxy memory buffer in seconds.
 When enabled (not zero) and records in proxy memory buffer are older, then it forces proxy buffer to switch to database mode until all records are uploaded to server. This parameter must be less or equal to `ProxyOfflineBuffer` parameter.

Default: 0
 Range: 0;600-864000

ProxyMemoryBufferSize

Set the shared memory cache size for collected history, discovery, and auto registration data (in bytes). If enabled (not zero) proxy will keep history discovery and autoregistration data in memory unless cache is full or stored records are older than defined `ProxyMemoryBufferAge`. This parameter cannot be used together with `ProxyLocalBuffer` parameter.

Default: 0
 Range: 0;128K-2G

ProxyMode

Set the proxy operating mode.
0 - proxy in the active mode
1 - proxy in the passive mode
Note that (sensitive) proxy configuration data may become available to parties having access to the Zabbix server trapper port when using an active proxy. This is possible because anyone may pretend to be an active proxy and request configuration data; authentication does not take place.

Default: 0
 Range: 0-1

ProxyOfflineBuffer

Set the duration of keeping data (in hours) in case of no connectivity with Zabbix server.
Older data will be lost.

Default: 1
 Range: 1-720

Server

If ProxyMode is set to `active`: specify Zabbix server IP address or DNS name (address:port) or cluster (address:port;address2:port) to get configuration data from and send data to.
If port is not specified, the default port is used.
Cluster nodes must be separated by a semicolon.

If ProxyMode is set to `passive`: specify IP addresses, optionally in CIDR notation, or DNS names of Zabbix server. A comma-delimited list.
 Incoming connections will be accepted only from the addresses listed here. If IPv6 support is enabled then `127.0.0.1`, `::127.0.0.1`, `::ffff:127.0.0.1` are treated equally.
`::/0` will allow any IPv4 or IPv6 address. `0.0.0.0/0` can be used to allow any IPv4 address.

Example:

```
Server=127.0.0.1,192.168.1.0/24,::1,2001:db8::/32,zabbix.example.com
```

Mandatory: yes

SNMPTrapperFile

Specify the temporary file used for passing data from the SNMP trap daemon to the proxy.
Must be the same as in `zabbix_trap_receiver.pl` or SNMPTT configuration file.

Default: `/tmp/zabbix_traps.tmp`

SocketDir

Specify the directory for storing IPC sockets used by internal Zabbix services.

Default: `/tmp`

SourceIP

Specify the source IP address for:

- outgoing connections to Zabbix server
- agentless connections (VMware, SSH, JMX, SNMP, Telnet and simple checks)
- HTTP agent connections
- script item JavaScript HTTP requests
- preprocessing JavaScript HTTP requests
- connections to the Vault

SSHKeyLocation

Specify the location of public and private keys for SSH checks and actions.

SSLCertLocation

Specify the location of SSL-client certificate files for client authentication.
This parameter is used in web monitoring, for communication with Vault, and for telemetry data source ClickHouse.

SSLKeyLocation

Specify the location of SSL private key files for client authentication.
This parameter is used in web monitoring, for communication with Vault, and for telemetry data source ClickHouse.

SSLCALocation

Specify the location of certificate authority (CA) files for SSL server certificate verification.
Note that the value of this parameter will be set as the `CURLOPT_CAPATH` libcurl option. For libcurl versions before 7.42.0, this only has effect if libcurl was compiled to use OpenSSL. For more information, see the cURL web page.
This parameter is used in web monitoring, SMTP authentication, HTTP agent items, for communication with Vault, and for telemetry data source ClickHouse.

StartAgentPollers

Set the number of pre-forked instances of Zabbix agent `pollers`. See `MaxConcurrentChecksPerPoller`.

Default: 1
 Range: 0-1000

StartAPMCollector

Set to 1 to start a pre-forked instance of an APM collector (the `apm manager` and `apm worker` processes).

Default: 0
 Range: 0-1

StartBrowserPollers

Set the number of pre-forked instances of browser item `pollers`.

Default: 1
 Range: 0-1000

StartDBSyncers

Set the number of pre-forked instances of `history syncers`.
Note: Be careful when changing this value, increasing it may do more harm than good.

Default: 4
 Range: 1-100

StartDiscoverers

Set the number of pre-forked instances of `discovery workers`.

Default: 5
 Range: 0-1000

StartHTTPAgentPollers

Set the number of pre-forked instances of `HTTP agent pollers`. See `MaxConcurrentChecksPerPoller`.

Default: 1
 Range: 0-1000

StartHTTPPollers

Set the number of pre-forked instances of `HTTP pollers`.

Default: 1
 Range: 0-1000

StartIPMIPollers

Set the number of pre-forked instances of `IPMI pollers`.

Default: 0
 Range: 0-1000

StartJavaPollers

Set the number of pre-forked instances of `Java pollers`.

Default: 0
 Range: 0-1000

StartODBCPollers

Set the number of pre-forked instances of `ODBC pollers`.

Default: 1
 Range: 0-1000

StartPingers

Set the number of pre-forked instances of `ICMP pingers`.

Default: 1
 Range: 0-1000

StartPollersUnreachable

Set the number of pre-forked instances of pollers for unreachable hosts (including IPMI and Java). At least one poller for unreachable hosts must be running if regular, IPMI or Java pollers are started.

Default: 1
 Range: 0-1000

StartPollers

Set the number of pre-forked instances of `pollers`.

Default: 5
 Range: 0-1000

StartPreprocessors

Set the number of pre-started instances of preprocessing `worker` threads should be set to no less than the available CPU core count. More workers should be set if preprocessing is not CPU-bound and has lots of network requests.

Default: 16
 Range: 1-1000

StartSNMPPollers

Set the number of pre-forked instances of `SNMP pollers`. See `MaxConcurrentChecksPerPoller`.

Default: 1
 Range: 0-1000

StartTelemetryQueryPollers

Set the number of pre-forked instances of asynchronous Telemetry query pollers. See MaxConcurrentChecksPerPoller.

Default: 1
 Range: 0-1000

StartSNMPTrapper

Set to "1" to start an `SNMP trapper` process.

Default: 0
 Range: 0-1

StartTrappers

Set the number of pre-forked instances of `trappers`.
Trappers accept incoming connections from Zabbix sender and active agents.

Default: 5
 Range: 0-1000

StartVMwareCollectors

Set the number of pre-forked `vmware collector` instances.

Default: 0
 Range: 0-250

StatsAllowedIP

Specify IP addresses (optionally in CIDR notation) or DNS names of external Zabbix instances that are allowed to request stats. A comma-delimited list.
 If this parameter is not set no stats requests will be accepted.
If IPv6 support is enabled then `127.0.0.1`, `::127.0.0.1`, `::ffff:127.0.0.1` are treated equally and `::/0` will allow any IPv4 or IPv6 address. `0.0.0.0/0` can be used to allow any IPv4 address.

Example:

```
StatsAllowedIP=127.0.0.1,192.168.1.0/24,::1,2001:db8::/32,zabbix.example.com
```

TelemetryProvider

Set the telemetry data source using the format `<provider>[;<options>]` where:

- `<provider>` - one of the supported telemetry data sources: `clickhouse`
- `<options>` - comma-delimited `key=value` pairs. Values can be quoted.

Options for ClickHouse:

- `url` - set the HTTP or HTTPS URL used to connect to the telemetry data source.
- `db` - set the name of the database.
- `username` - set the username used to log in to the database.
- `password` - set the password used to log in to the database.
- `ssl_cert_file`, `ssl_key_file`, `ssl_key_password`, `ssl_verify_peer`, `ssl_verify_host`, `ssl_ca_location`, `ssl_cert_location`, `ssl_key_location` - set these options to configure a secure (TLS/SSL) connection to the database.
- `source_ip` - set the source IP address used for outgoing connections to the telemetry data source.
- `vault_path` - set the path to a Vault secret holding the `username` and `password` values. Do not use together with the `username` and `password` options.

Example for ClickHouse:

```
TelemetryProvider=clickhouse;url=http://localhost:8123,db=zabbix,username=zabbix,password=zabbix
```

Example for ClickHouse with Vault:

```
TelemetryProvider=clickhouse;url=http://localhost:8123,db=zabbix,vault_path=clickhouse
```

Timeout

Specify how long to wait (in seconds) for establishing connection and exchanging data with Zabbix server, agent, web service, and legacy SNMP checks (single OID number or string).

This parameter defines the duration for various communication operations:

- Remote command execution on Zabbix agent
- SSH / Telnet command execution
- Rescheduling of items when IPMI interface becomes unavailable
- Sending response to Zabbix server when failing to exchange data due to rights or encryption issues
- Deadline for IPC asynchronous sockets and runtime control options
- Asynchronous poller DNS requests

- Response for active check heartbeat
- Retrieval of Zabbix agent data (values) from active agents
- Retrieval of data from Zabbix sender
- Sending active check list to Zabbix agent

This timeout will **not** be used for those checks that have flexible timeout settings configured in the frontend (on global, proxy, or per-item level). For example, SNMP `walk[OID]` and `get[OID]` items use the configured timeout from the frontend; legacy SNMP checks still use the server timeout value.

Default: 3
 Range: 1-30

TLSAccept

Specify the level of encryption required for incoming connections from Zabbix server. Used for a passive proxy, ignored on an active proxy. Multiple values can be specified, separated by comma: - *unencrypted* - accept connections without encryption (default); - *psk* - accept connections with TLS and a pre-shared key (PSK); - *cert* - accept connections with TLS and a certificate.

Mandatory: yes for passive proxy, if TLS certificate or PSK parameters are defined (even for *unencrypted* connection); otherwise, no

TLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.

TLSCertFile

Specify full pathname of the file containing the proxy certificate or certificate chain, used for encrypted communications between Zabbix components.

TLSCipherAll

Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.

Example:

```
TLS_AES_256_GCM_SHA384:TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256
```

TLSCipherAll13

Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.

Example for GnuTLS:

```
NONE:+VERS-TLS1.2:+ECDHE-RSA:+RSA:+ECDHE-PSK:+PSK:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+SHA1:+CURVE-ALL
```

Example for OpenSSL:

```
EECDH+aRSA+AES128:RSA+aRSA+AES128:kECDHEPSK+AES128:kPSK+AES128
```

TLSCipherCert

Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for certificate-based encryption.

Example for GnuTLS:

```
NONE:+VERS-TLS1.2:+ECDHE-RSA:+RSA:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+SHA1:+CURVE-ALL:+COMP-NULL:+SIG
```

Example for OpenSSL:

```
EECDH+aRSA+AES128:RSA+aRSA+AES128
```

TLSCipherCert13

Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for certificate-based encryption.

TLSCipherPSK

Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for PSK-based encryption.

Example for GnuTLS:

```
NONE:+VERS-TLS1.2:+ECDHE-PSK:+PSK:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+SHA1:+CURVE-ALL:+COMP-NULL:+SIG
```

Example for OpenSSL:

```
kECDHEPSK+AES128:kPSK+AES128
```

TLSCipherPSK13

Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for PSK-based encryption.

Example:

```
TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256
```

TLSConnect

Specify the level of encryption required for proxy connections to Zabbix server. Used for an active proxy, ignored on a passive proxy. Only one value can be specified: - *unencrypted* - connect without encryption (default); - *psk* - connect using TLS and a pre-shared key (PSK); - *cert* - connect using TLS and a certificate.

Mandatory: yes for active proxy, if TLS certificate or PSK parameters are defined (even for *unencrypted* connection); otherwise, no

TLSCRLFile

Specify full pathname of the file containing revoked certificates. This parameter is used for encrypted communications between Zabbix components.

TLSKeyFile

Specify full pathname of the file containing the proxy private key, used for encrypted communications between Zabbix components.

TLSListen

Require TLS-only connections on the trapper socket.

Supported values:

- required - accept only TLS connections

TLSPSKFile

Specify full pathname of the file containing the proxy pre-shared key, used for encrypted communications with Zabbix server.

TLSPSKIdentity

Specify the pre-shared key identity string, used for encrypted communications with Zabbix server.

TLSServerCertIssuer

Specify the allowed server certificate issuer.

TLSServerCertSubject

Specify the allowed server certificate subject.

TmpDir

Specify the temporary directory.

Default: `/tmp`

TrapperTimeout

Specify timeout in seconds for:
 - retrieval of configuration data from the Zabbix server;
 - global script execution or remote command execution.

Default: 300
 Range: 1-300

UnavailableDelay

Set the frequency of checking host availability during the unavailability period in seconds.

Default: 60
 Range: 1-3600

UnreachableDelay

Set the frequency of checking host availability during the unreachability period in seconds.

Default: 15
 Range: 1-3600

UnreachablePeriod

Set after how many seconds of unreachability treat host as unavailable.

Default: 45
 Range: 1-3600

User

Drop privileges to a specific, existing user on the system.
Only has effect if run as `root` and `AllowRoot` is disabled.

Default: `zabbix`

Vault

Specify the vault provider: - *HashiCorp* - HashiCorp KV Secrets Engine version 2; - *CyberArk* - CyberArk Central Credential Provider
Must match the vault provider set in the frontend.

Default: `HashiCorp`

VaultAppRoleID

Specify the HashiCorp vault AppRole ID.
 If set, enables AppRole authentication. Cannot be used with VaultToken.
 Supported only for HashiCorp vault.

VaultAppSecretID

Specify the HashiCorp vault AppRole secret ID.
 Supported only for HashiCorp vault.

VaultDBPath

Specify the location for retrieving database credentials by keys. It is a Vault path or query, depending on the Vault.

The keys used for HashiCorp are `password` and `username`.

Example path with `VaultPrefix=/v1/secret/data/zabbix/`:

```
database
```

Example path without VaultPrefix:

```
secret/zabbix/database
```

The keys used for CyberArk are `Content` and `UserName`.

Example:

```
AppID=zabbix_server&Query=Safe=passwordSafe;Object=zabbix_proxy_database
```

This option can only be used if `DBUser` and `DBPassword` are not specified.

VaultPrefix

Specify custom prefix for the vault path or query, depending on the Vault. The most suitable defaults will be used if not specified.
Note that 'data' is automatically appended after mountpoint for HashiCorp if `VaultPrefix` is not specified.
Note that 'data' is automatically appended after mountpoint for HashiCorp if `VaultPrefix` is not specified.

Example prefix for HashiCorp:

```
/v1/secret/data/zabbix/
```

Example prefix for Cyberark:

```
/AIMWebService/api/Accounts?
```

VaultTLSCertFile

Specify the name of the SSL certificate file used for client authentication. The certificate file must be in PEM1 format.
If the certificate file contains also the private key, leave the SSL key file field empty.
The directory containing this file is specified by the SSLCertLocation configuration parameter.
This option can be omitted, but is recommended for CyberArkCCP vault.

VaultTLSKeyFile

Specify the name of the SSL private key file used for client authentication. The private key file must be in PEM1 format.
The directory containing this file is specified by the SSLKeyLocation configuration parameter.
This option can be omitted, but is recommended for CyberArkCCP vault.

VaultToken

Specify the HashiCorp vault authentication token that should have been generated exclusively for Zabbix proxy with read-only permission to the path specified in the optional VaultDBPath configuration parameter.
It is an error if `VaultToken` and the

VAULT_TOKEN environment variable are defined at the same time.
 Cannot be used with VaultAppRoleID.
 Supported only for HashiCorp vault.

VaultURL

Specify the vault server URL. The system-wide CA certificates directory will be used if SSLCALocation is not specified.

Default: https://127.0.0.1:8200

VMwareCacheSize

Set the shared memory size for storing VMware data.
A VMware internal check zabbix[vmware,buffer,...] can be used to monitor the VMware cache usage (see Internal checks).
Note that shared memory is not allocated if there are no vmware collector instances configured to start.

Default: 8M
 Range: 256K-2G

VMwareFrequency

Set the frequency of data gathering from a single VMware service (in seconds).
This frequency should be set to the least update interval of any VMware monitoring item.

Default: 60
 Range: 10-86400

VMwarePerfFrequency

Set the frequency of retrieving performance counter statistics from a single VMware service (in seconds).
This frequency should be set to the least update interval of any VMware monitoring item that uses VMware performance counters.

Default: 60
 Range: 10-86400

VMwareTimeout

Set the maximum number of seconds a vmware collector will wait for a response from VMware service (vCenter or ESX hypervisor).

Default: 10
 Range: 1-300

WebDriverURL

Specify the WebDriver interface URL.

Example (used with Selenium WebDriver standalone server):

```
WebDriverURL=http://localhost:4444
```

Agent

Overview Zabbix agent is deployed on a monitoring target to actively monitor local resources and applications (hard drives, memory, processor statistics, etc.).

The agent gathers operational information locally and reports data to Zabbix server for further processing. In case of failures (such as a hard disk running full or a crashed service process), Zabbix server can actively alert the administrators of the particular machine that reported the failure.

Zabbix agents are highly efficient because of the use of native system calls for gathering statistical information.

It is possible to run multiple instances of the agent on a host.

See also:

- Passive and active check protocols
- Running agent as root
- Configuration parameters: Zabbix agent
- Configuration parameters: Zabbix agent (Windows)
- Zabbix agent on Microsoft Windows
- Zabbix agent on IBM AIX

Passive and active checks

Zabbix agents can perform passive and active checks:

- Passive checks - Zabbix agent responds to a request from Zabbix server (or proxy). For example, the server requests data (e.g., CPU load), and the agent returns the result.
- Active checks - Zabbix agent collects and sends data without waiting for a request from Zabbix server (or proxy). First, it retrieves a list of monitoring items from server (CPU load, available memory, etc.), then collects the required data and periodically sends new values back to it.

The agent check type is configured by selecting the respective monitoring item type. Zabbix agent processes items of type *Zabbix agent* or *Zabbix agent (active)*.

Supported platforms For supported platforms, refer to the Requirements page.

Agent on Unix-like systems Zabbix agent on Unix-like systems is run on the host being monitored.

Installation

Zabbix agent can be installed on Linux-based systems using one of the following methods:

- Zabbix packages - select the Agent component (after selecting your Zabbix version, OS distribution, and OS version) and follow the provided instructions;
- Zabbix sources - download the source files and compile Zabbix agent by configuring the sources with the `--enable-agent` option.

Attention:
In general, 32-bit Zabbix agents will work on 64-bit systems, but may fail in some cases.

Note:
Pre-compiled Zabbix agent binaries are available for download for macOS, IBM AIX, FreeBSD, OpenBSD, and Solaris. Legacy binaries, compatible with current Zabbix server/proxy version, are available for NetBSD and HP-UX. For installing and running the agent on IBM AIX, see Zabbix agent on IBM AIX.

If installed as package

Zabbix agent runs as a daemon process. The agent can be started by executing:

```
systemctl start zabbix-agent
```

This will work on most of GNU/Linux systems. On other systems you may need to run:

```
/etc/init.d/zabbix-agent start
```

To stop, restart, or check the status of Zabbix agent, use the following commands:

```
systemctl stop zabbix-agent
systemctl restart zabbix-agent
systemctl status zabbix-agent
```

If installed from binary archive

After extracting the archive, locate the `zabbix_agentd` binary and the `zabbix_agentd.conf` configuration file. Their location inside the archive depends on the platform (for example, `sbin/` and `conf/`, or `usr/sbin/` and `etc/zabbix/`). Run the binary directly from the extracted directory and specify the configuration file with the `-c` option. For example:

```
/opt/zabbix/sbin/zabbix_agentd -c /opt/zabbix/etc/zabbix/zabbix_agentd.conf
```

On UNIX-like systems, Zabbix agent is designed to run as a non-root user. If started as root, it switches to the `zabbix` user (or the user specified by the [User](<PARAMS>#user) parameter) unless [AllowRoot=1](<PARAMS>#allowroot) is set in the configuration file. This user must exist on the system; otherwise, the agent does not start.

On IBM AIX, the agent can also be run with the System Resource Controller (SRC) and started at system boot; see Zabbix agent on IBM AIX.

Start up manually

You can start Zabbix agent by running the `zabbix_agentd` binary directly, for example:

```
zabbix_agentd
```

This works when the directory containing `zabbix_agentd` is included in `PATH`.

Agent on Windows systems Zabbix agent on Windows runs as a Windows service.

Installation

Zabbix agent can be installed on Windows using one of the following methods:

- Pre-compiled Zabbix agent binaries - download the Zabbix agent MSI installer package and follow the instructions on the Windows agent installation from MSI page;
- Zabbix sources - download the source files and follow the instructions on the Building Zabbix agent on Windows page.

See the Zabbix agent on Microsoft Windows page for additional details on installing Zabbix agent (from a ZIP archive) as a Windows service.

Options The following command-line parameters can be used with Zabbix agent:

Parameter	Description
UNIX and Windows agent	
`-c --config <config-file>`	Path to the configuration file. You may use this option to specify a configuration file that is not the default one.
`-f --foreground`	Run Zabbix agent in foreground (default: `true`).
`-p --print`	Print known items and exit. Note that to return user parameter results as well, you must specify the configuration file (if it is not in the default location).
`-t --test <item key>`	Test specified item and exit. Note that to return user parameter results as well, you must specify the configuration file (if it is not in the default location).
`-T --test-config`	Validate configuration file and exit.
`-h --help`	Display help information.
`-V --version`	Display version number.
UNIX agent only	
`-R --runtime-control <option>`	Perform administrative functions. See runtime control.
Windows agent only	
`-m --multiple-agents`	Use multiple agent instances (with `-i`, `-d`, `-s`, `-x` options). To distinguish service names of instances, each service name will include the Hostname value from the specified configuration file.
`-S --startup-type <value>`	Set the startup type of the Zabbix Windows agent service. Allowed values: `automatic` - *(default)* start the service automatically at Windows startup; `delayed` - delay starting the service after the automatically started services have completed startup (available on Windows Server 2008/Vista and later versions); `manual` - start the service manually (by a user or application); `disabled` - disable the service, so that it cannot be started by a user or application. You may use this option together with the `-i` option, or separately to modify the startup type of an already installed service.
`-i --install`	Install Zabbix Windows agent as service.
`-d --uninstall`	Uninstall Zabbix Windows agent service.
`-s --start`	Start Zabbix Windows agent service.
`-x --stop`	Stop Zabbix Windows agent service.

Specific examples of using command-line parameters:

- Printing all built-in agent items with values.
- Testing a user parameter with `mysql.ping` key defined in the specified configuration file.
- Installing a Zabbix Agent service for Windows using the default path to the configuration file.
- Installing a Zabbix Agent [Hostname] service for Windows using the configuration file located in the same folder as agent executable and making the service name unique by extending it by `Hostname` value from the configuration file.
- Modifying the startup type of an installed Zabbix Agent service for Windows using the configuration file located in the same folder as agent executable.

```
zabbix_agentd --print
zabbix_agentd -t "mysql.ping" -c /etc/zabbix/zabbix_agentd.conf
zabbix_agentd.exe -i
```

```
zabbix_agentd.exe -i -m -c zabbix_agentd.conf
zabbix_agentd.exe -c zabbix_agentd.conf -S delayed
```

Runtime control

With runtime control options you may change the log level of agent processes.

Option	Description	Target
`log_level_increase[=<target>]`	Increase log level. If target is not specified, all processes are affected.	Target can be specified as: **process type** - all processes of specified type (e.g., `listener`) See all agent process types. **process type,N** - process type and number (e.g., `listener,3`) **pid** - process identifier (1 to 65535). For larger values specify target as 'process type,N'.
`log_level_decrease[=<target>]`	Decrease log level. If target is not specified, all processes are affected.	
`userparameter_reload`	Reload values of the `UserParameter` and `Include` options from the current configuration file.	

Examples:

- Increasing log level of all processes
- Increasing log level of the third listener process
- Increasing log level of process with PID 1234
- Decreasing log level of all active check processes

```
zabbix_agentd -R log_level_increase
zabbix_agentd -R log_level_increase=listener,3
zabbix_agentd -R log_level_increase=1234
zabbix_agentd -R log_level_decrease="active checks"
```

Note:
Runtime control is not supported on OpenBSD, NetBSD and Windows.

Agent process types Agent process types are as follows:

- `active checks` - process for performing active checks;
- `collector` - process for data collection;
- `listener` - process for listening to passive checks.

The agent log file can be used to observe these process types.

The agent log file is created with read-write permissions for the file owner only. Additionally the file is readable by the owner group. All other permissions are denied.

Process user Zabbix agent on UNIX is designed to run as a non-root user. It will run as whatever non-root user it is started as. So you can run agent as any non-root user without any issues.

If you will try to run it as 'root', it will switch to a hardcoded 'zabbix' user, which must be present on your system. You can only run agent as 'root' if you modify the `AllowRoot` parameter in the agent configuration file accordingly.

Configuration file The default path to the configuration file depends on the operating system and Zabbix installation method:

- On UNIX, the default is:
  - For installation from sources: `/usr/local/etc/zabbix_agentd.conf`, or as set by compile-time variables `--sysconfdir` or `--prefix`.
  - For installation from packages and appliance: `/etc/zabbix/zabbix_agentd.conf`.
- On Windows, the default is `C:\Program Files\Zabbix Agent\zabbix_agentd.conf`.

If a single instance of the agent is running on a host, it can use the default configuration file or a configuration file specified on the command line. In case of multiple instances, each agent instance must have its own configuration file (one of the instances can use the default configuration file).

For details on configuring Zabbix agent see the configuration file options for Unix Zabbix agent or Windows Zabbix agent.

Locale Note that the agent requires a UTF-8 locale so that some textual agent items can return the expected content. Most modern Unix-like systems have a UTF-8 locale as default, however, there are some systems where that may need to be set specifically.

Exit code Zabbix agent returns 0 in case of successful exit and 1 in case of failure.

Passive and active check protocols

This section provides protocol details on passive and active checks performed by Zabbix agent and Zabbix agent 2.

Zabbix uses a JSON-based communication protocol for communicating with the agents.

Zabbix agent and Zabbix agent 2 protocols have been unified since Zabbix 7.0. The difference between Zabbix agent and Zabbix agent 2 requests/responses is expressed by the ”variant” tag value.

Passive checks

A passive check is a simple data request. Zabbix server or proxy asks for some data (for example, CPU load) and Zabbix agent sends back the result to the server.

Passive checks are executed asynchronously - it is not required to receive the response to one request before other checks are started. DNS resolving is asynchronous as well.

The agent poller will attempt to connect to all addresses returned by the DNS lookup. This ensures that if one IP address is unreachable, the poller will try the next available address, increasing the likelihood of a successful connection. This enhancement applies to both Zabbix server and proxy.

The maximum concurrency of asynchronous checks is 1000 (defined by MaxConcurrentChecksPerPoller).

The number of asynchronous agent pollers is defined by the StartAgentPollers parameter.

Server request

For definition of header and data length please refer to protocol details.

```
{
  "request": "passive checks",
  "data": [
    {
      "key": "agent.version",
      "timeout": 3
    }
  ]
}
```

Field		Type	Mandatory	Value
request		*string*	yes	"passive checks"
data		*array of object*	yes	Passive check item.
	key	*string*	yes	Item key with expanded macros.
	timeout	*number*	yes	Communication timeout.

Agent response

```
{
  "version": "8.2.0",
  "variant": 2,
  "data": [
    {
      "value": "8.2.0"
    }
```

```
  ]
}
```

Field		Type	Mandatory	Value
version		*string*	yes	The agent version number.
variant		*number*	yes	The agent variant (*1* - Zabbix agent, *2* - Zabbix agent 2).
data		*array of object*	yes	Contains the result of the check.
	value	*string*	no	The item value if the check was successful.
	error	*string*	no	The error message if the check was not successful.

For example, for supported items:

1. Server opens a TCP connection
2. Server sends **<HEADER><DATALEN>{"request":"passive checks","data":[{"key":"agent.ping","timeout":3}]}**
3. Agent reads the request and responds with **<HEADER><DATALEN>{"version":"8.2.0","variant":2,"data":[{"value":1}]}**
4. Server processes data to get the value, '1' in our case
5. TCP connection is closed

For not supported items:

1. Server opens a TCP connection
2. Server sends **<HEADER><DATALEN>{"request":"passive checks","data":[{"key":"vfs.fs.size[/nono]","timeout":3}]}**
3. Agent reads the request and responds with **<HEADER><DATALEN>{"version":"8.2.0","variant":2,"data":[{"error":"Unsupp**
   item key."}]}
4. Server processes data, changes item state to not supported with the specified error message
5. TCP connection is closed

Failover to old protocol

To make sure that Zabbix server or proxy can work with agents from pre-7.2 versions, which have plaintext protocol, a failover to the old protocol is implemented.

Passive checks are performed using the JSON protocol (7.0 and later) after restart or when the interface configuration is changed. If no valid JSON is received in response (agent sent "ZBX_NOTSUPPORTED"), Zabbix will cache the interface as old protocol and **retry** the check by sending only the item key.

Note that every hour Zabbix server/proxy will again try working with the new protocol with all interfaces, falling back to the old protocol if required.

Active checks

Active checks require more complex processing. The agent must first retrieve from the server/proxy a list of items and/or remote commands for independent processing.

The servers/proxies to get the active checks from are listed in the `ServerActive` parameter of the agent configuration file. The frequency of asking for these checks is set by the `RefreshActiveChecks` parameter in the same configuration file. However, if refreshing active checks fails, it is retried after hardcoded 60 seconds.

Note:
Since Zabbix 6.4 the agent (in active mode) no longer receives from the server/proxy a full copy of the configuration once every two minutes (default). Instead, in order to decrease network traffic and resources usage, an incremental configuration sync is performed every 5 seconds (default) upon which the server/proxy provides a full copy of the configuration **only** if the agent has not yet received it, or something has changed in host configuration, global macros or global regular expressions.

The agent then periodically sends the new values to the server(s). If the agent received any remote commands to execute, the execution result will also be sent. Note that remote command execution on an active agent is supported since Zabbix agent 7.0.

Note:
If an agent is behind the firewall you might consider using only Active checks because in this case you wouldn't need to modify the firewall to allow initial incoming connections.

Getting the list of items

Agent request

The active checks request is used to obtain the active checks to be processed by agent. This request is sent by the agent upon start and then with RefreshActiveChecks intervals.

```
{
  "request": "active checks",
  "host": "Zabbix server",
  "host_metadata": "mysql,nginx",
  "interface": "zabbix.server.lan",
  "ip": "159.168.1.1",
  "port": 12050,
  "version": "8.2.0",
  "variant": 2,
  "config_revision": 1,
  "session": "e3dcbd9ace2c9694e1d7bbd030eeef6e"
}
```

Field	Type	Mandatory	Value
request	*string*	yes	`active checks`
host	*string*	yes	Host name.
host_metadata	*string*	no	The configuration parameter HostMetadata or HostMetadataItem metric value.
interface	*string*	no	The configuration parameter HostInterface or HostInterfaceItem metric value.
ip	*string*	no	The configuration parameter ListenIP first IP if set.
port	*number*	no	The configuration parameter ListenPort value if set and not default agent listening port.
version	*string*	yes	The agent version number.
variant	*number*	yes	The agent variant (*1* - Zabbix agent, *2* - Zabbix agent 2).
config_revision	*number*	no	Configuration identifier for incremental configuration sync.
session	*string*	no	Session identifier for incremental configuration sync.

Server response

The active checks response is sent by the server back to agent after processing the active checks request.

```
{
  "response": "success",
  "config_revision": 2,
  "data": [
    {
      "key": "system.uptime",
      "itemid": 1234,
      "delay": "10s",
      "lastlogsize": 0,
      "mtime": 0
    },
    {
      "key": "agent.version",
      "itemid": 5678,
      "delay": "10m",
      "lastlogsize": 0,
      "mtime": 0,
      "timeout": "30s"
    }
  ],
  "commands": [
    {
      "command": "df -h --output=source,size / | awk 'NR>1 {print $2}'",
      "id": 1324,
      "wait": 1
    }
  ]
}
```

Field	Type	Mandatory	Value
response	*string*	yes	`success` | `failed`
info	*string*	no	Error information in case of failure.
data	*array of objects*	no	Active check items. Omitted if host configuration is unchanged.
key	*string*	no	Item key with expanded macros.
itemid	*number*	no	Item identifier.
delay	*string*	no	Item update interval. Flexible/scheduling intervals are supported by both Zabbix agent and Zabbix agent 2 since Zabbix 7.0.
lastlogsize	*number*	no	Item lastlogsize.
mtime	*number*	no	Item mtime.
timeout	*string*	no	Item timeout.
refresh_unsupported	*number*	no	Unsupported item refresh interval.
regexp	*array of objects*	no	Global regular expressions.
name	*string*	no	Global regular expression name.
expression	*string*	no	Global regular expression.
expression_type	*number*	no	Global regular expression type.
exp_delimiter	*string*	no	Global regular expression delimiter.
case_sensitive	*number*	no	Global regular expression case sensitivity setting.
commands	*array of objects*	no	Remote commands to execute. Included if remote command execution has been triggered by an action operation or manual script execution. Note that remote command execution on an active agent is supported since Zabbix agent 7.0. Older active agents will ignore any remote commands included in the active checks server response.
command	*string*	no	Remote command.
id	*number*	no	Remote command identifier.
wait	*number*	no	Remote command mode of execution ("0" (nowait) for commands from action operations; "1" (wait) for commands from manual script execution).
timeout	*number*	no	Remote command execution timeout in server/proxy configuration.
config_revision	*number*	no	Configuration identifier for incremental configuration sync. Omitted if host configuration is unchanged. Incremented if host configuration is changed.

The server must respond with success.

For example:

1. Agent opens a TCP connection
2. Agent asks for the list of checks
3. Server responds with a list of items and remote commands to execute
4. Agent parses the response
5. TCP connection is closed
6. Agent starts periodical collection of data and executes remote commands (supported since Zabbix agent 7.0)

Attention:
Note that (sensitive) configuration data may become available to parties having access to the Zabbix server trapper port when using an active check. This is possible because anyone may pretend to be an active agent and request item configuration data; authentication does not take place unless you use encryption options.

Sending in collected data

Agent sends

The agent data request contains the gathered item values and the values for executed remote commands (if any).

```
{
  "request": "agent data",
  "data": [
    {
      "id": 1,
      "itemid": 5678,
```

```
      "value": "7.0.0",
      "clock": 1712830783,
      "ns": 76808644
    },
    {
      "id": 2,
      "itemid": 1234,
      "value": "69672",
      "clock": 1712830783,
      "ns": 77053975
    }
  ],
  "commands": [
    {
      "id": 1324,
      "value": "16G"
    }
  ],
  "session": "8495cd52070e6ca52b371f29c8574165",
  "host": "Zabbix server",
  "version": "8.2.0",
  "variant": 2
}
```


Field	Type	Mandatory	Value
request	*string*	yes	`agent data`
data	*array of objects*	yes	Item values.
id	*number*	yes	The value identifier (incremental counter used for checking duplicated values in the case of network problems).
itemid	*number*	yes	The item identifier.
value	*string*	no	The item value.
lastlogsize	*number*	no	The item lastlogsize.
mtime	*number*	no	The item mtime.
state	*number*	no	The item state.
source	*string*	no	The value event log source.
eventid	*number*	no	The value event log eventid.
severity	*number*	no	The value event log severity.
timestamp	*number*	no	The value event log timestamp.
clock	*number*	yes	The value timestamp (seconds since Epoch).
ns	*number*	yes	The value timestamp nanoseconds.
commands	*array of objects*	no	Remote commands execution result. Note that remote command execution on an active agent is supported since Zabbix agent 7.0. Older active agents will ignore any remote commands included in the active checks server response.
id	*number*	no	Remote command identifier.
value	*string*	no	Remote command execution result if the execution was successful.
error	*string*	no	Remote command execution error message if the execution failed.
session	*string*	yes	Unique session identifier generated each time when agent is started (32 hexadecimal digits long).
host	*string*	yes	Host name.
version	*string*	yes	The agent version number.
variant	*number*	yes	The agent variant (*1* - Zabbix agent, *2* - Zabbix agent 2).

A virtual ID is assigned to each value. Value ID is a simple ascending counter, unique within one data session (identified by the session token). This ID is used to discard duplicate values that might be sent in poor connectivity environments.

Server response

The agent data response is sent by the server back to the agent after processing the agent data request.

```
{
  "response": "success",
  "info": "processed: 2; failed: 0; total: 2; seconds spent: 0.003534"
}
```

Field	Type	Mandatory	Value
response	*string*	yes	`success` | `failed`
info	*string*	yes	Item processing results.

Attention:
If sending of some values fails on the server (for example, because host or item has been disabled or deleted), agent will not retry sending of those values.

For example:

1. Agent opens a TCP connection
2. Agent sends a list of values
3. Server processes the data and sends the status back
4. TCP connection is closed

Attention:
Error message will be trimmed to 2048 symbols on server side.

Heartbeat message

Agent sends

The heartbeat message is sent by an active agent to Zabbix server/proxy every `HeartbeatFrequency` seconds (configured in the Zabbix agent or Zabbix agent 2 configuration file).

It is used to monitor the availability of active checks.

```
{
  "request": "active check heartbeat",
  "host": "Zabbix server",
  "heartbeat_freq": 60,
  "version": "8.2.0",
  "variant": 2
}
```

Field	Type	Mandatory	Value
request	*string*	yes	`active check heartbeat`
host	*string*	yes	The host name.
heartbeat_freq	*number*	yes	The agent heartbeat frequency (HeartbeatFrequency configuration parameter).
version	*string*	yes	The agent version number.
variant	*number*	yes	The agent variant (*1* - Zabbix agent, *2* - Zabbix agent 2).

Redirect response

When a host has been reassigned, the server may instruct the agent to redirect its heartbeat (and subsequent active checks) to another proxy or server instance.

```
  {
    "response": "failed",
    "redirect": {
      "revision": 2,
      "address": "192.0.2.1:10055"
    }
  }
```

Field		Type	Mandatory	Value
response		*string*	yes	`success` | `failed`
redirect		*object*	yes	Redirect instructions.
	revision	*number*	yes	Configuration revision identifier.
	address	*string*	yes	Target server/proxy address.

Older XML protocol

Note:
Zabbix will take up to 16 MB of XML Base64-encoded data, but a single decoded value should be no longer than 64 KB; otherwise, it will be truncated to 64 KB while decoding.

Running agent as root

Since Zabbix **5.0.0**, the systemd service file for Zabbix agent in official packages explicitly includes directives for `User` and `Group`. Both are set to `zabbix`.

It is no longer possible to configure which user Zabbix agent runs as via `zabbix_agentd.conf` file, because the agent will bypass this configuration and run as the user specified in the systemd service file. To run Zabbix agent as root you need to make the modifications described below.

Zabbix agent

To override the default user and group for Zabbix agent, run:

```
systemctl edit zabbix-agent
```

Then, add the following content:

```
[Service]
User=root
Group=root
```

Reload daemons and restart the zabbix-agent service:

```
systemctl daemon-reload
systemctl restart zabbix-agent
```

For **Zabbix agent** this re-enables the functionality of configuring user in the agent configuration file. Now you need to set `User=root` and `AllowRoot=1` configuration parameters in the configuration file.

Zabbix agent 2

To override the default user and group for Zabbix agent 2, run:

```
systemctl edit zabbix-agent2
```

Then, add the following content:

```
[Service]
User=root
Group=root
```

Reload daemons and restart the zabbix-agent2 service:

```
systemctl daemon-reload
systemctl restart zabbix-agent2
```

For **Zabbix agent2** this completely determines the user that it runs as. No additional modifications are required.

Configuration parameters: Zabbix agent

Overview

Configuration parameters allow customizing Zabbix agent behavior.

Specify parameter values in the Zabbix agent configuration file (`zabbix_agentd.conf`). All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Alias	Set an alias for an item key.
AllowKey	Allow the execution of those item keys that match a pattern.
AllowKeyRegexp	Allow the execution of item keys that match a regular expression-based pattern.
AllowRoot	Allow the agent to run as `root`.
BufferSend	Set the frequency of sending values from the buffer to Zabbix server/proxy.
BufferSize	Set the maximum number of values in the memory buffer.
DebugLevel	Set the debug level.
DenyKey	Deny the execution of those item keys that match a pattern.
DenyKeyRegexp	Deny the execution of item keys that match a regular expression-based pattern.
EnableRemoteCommands	Allow remote commands from Zabbix server. This parameter is deprecated; use `AllowKey`, `AllowKeyRegexp`, `DenyKey`, or `DenyKeyRegexp` instead.
HeartbeatFrequency	Set the frequency of heartbeat messages in seconds.
HostInterface	Specify the host interface.
HostInterfaceItem	Specify an item for getting the host interface.
HostMetadata	Specify the host metadata.
HostMetadataItem	Specify a Zabbix agent item for getting the host metadata.
Hostname	Specify the hostname.
HostnameItem	Specify a Zabbix agent item for getting the hostname.
Include	Specify a directory or individual files to be included in the configuration file.
ListenBacklog	Set the maximum number of pending connections in the TCP queue.
ListenIP	Specify the agent listen IP address for connections from the server/proxy.
ListenPort	Specify the agent listen port.
LoadModule	Specify the module to load at agent startup.
LoadModulePath	Specify full path to the location of agent modules.
LogFile	Specify the agent log file.
LogFileSize	Set the maximum log file size.
LogRemoteCommands	Enable logging of executed shell commands as warnings.
LogType	Specify the log output type.
MaxLinesPerSecond	Set the maximum number of new lines per second for sending to Zabbix server/proxy when processing `log` or `logrt` active checks.
PidFile	Specify the PID file.
RefreshActiveChecks	Specify the frequency of refreshing the list of active checks.
Server	Specify the IP addresses or DNS names of Zabbix servers/proxies.
ServerActive	Specify the Zabbix server/proxy address or cluster configuration to get active checks from.
SourceIP	Specify the source IP address.
StartAgents	Set the number of pre-forked instances of `zabbix_agentd` for processing passive checks.
Timeout	Specify how long to wait (in seconds) for establishing connection and exchanging data with Zabbix proxy or server.
TLSAccept	Specify the level of encryption required for incoming connections.
TLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.
TLSCertFile	Specify full pathname of the file containing the agent certificate or certificate chain, used for encrypted communications between Zabbix components.
TLSCipherAll	Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.
TLSCipherAll13	Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.
TLSCipherCert	Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for certificate-based encryption.
TLSCipherCert13	Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for certificate-based encryption.
TLSCipherPSK	Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for PSK-based encryption.
TLSCipherPSK13	Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for PSK-based encryption.
TLSConnect	Specify the level of encryption required for connections to Zabbix server or proxy.
TLSCRLFile	Specify full pathname of the file containing revoked certificates. This parameter is used for encrypted communications between Zabbix components.

Parameter	Description
TLSKeyFile	Specify full pathname of the file containing the agent private key, used for encrypted communications between Zabbix components.
TLSPSKFile	Specify full pathname of the file containing the agent pre-shared key, used for encrypted communications with Zabbix server or proxy.
TLSPSKIdentity	Specify the pre-shared key identity string, used for encrypted communications with Zabbix server or proxy.
TLSServerCertIssuer	Specify the allowed server (proxy) certificate issuer.
TLSServerCertSubject	Specify the allowed server (proxy) certificate subject.
UnsafeUserParameters	Allow all characters to be passed in arguments to user-defined parameters.
User	Drop privileges to a specific, existing user on the system.
UserParameter	Specify a user-defined parameter to monitor.
UserParameterDir	Specify the default search path for UserParameter commands.

Note that:

- The default values reflect daemon defaults, not the values in the shipped configuration files.
- Values support environment variables.
- Zabbix supports configuration files only in UTF-8 encoding without BOM.
- Comments starting with # are only supported in the beginning of the line.

Parameter details

Alias

Set an alias for an item key. It can be used to substitute a long and complex item key with a shorter and simpler one.
 Multiple `Alias` parameters may be present. Multiple parameters with the same `Alias` key are not allowed.
 Different `Alias` keys may reference the same item key.
 Aliases can be used in `HostMetadataItem` but not in `HostnameItem` parameter.

Example 1: Retrieving the ID of user `zabbix`.

```
Alias=zabbix.userid:vfs.file.regexp[/etc/passwd,"^zabbix:.:([0-9]+)",,,,\1]
```

Now the `zabbix.userid` shorthand key may be used to retrieve data.

Example 2: Getting CPU utilization with default and custom parameters.

```
Alias=cpu.util:system.cpu.util
Alias=cpu.util[*]:system.cpu.util[*]
```

This allows you to use the `cpu.util` key to get CPU utilization percentage with default parameters as well as use `cpu.util[all, idle, avg15]` to get specific data about CPU utilization.

Example 3: Running multiple low-level discovery rules processing the same discovery items.

```
Alias=vfs.fs.discovery[*]:vfs.fs.discovery
```

Now it is possible to set up several discovery rules using `vfs.fs.discovery` with different parameters for each rule, e.g., `vfs.fs.discovery[foo]`, `vfs.fs.discovery[bar]`, etc.

AllowKey

Allow the execution of those item keys that match a pattern. The key pattern is a wildcard expression that supports the * character to match any number of any characters.
 Multiple key matching rules may be defined in combination with `DenyKey`, `AllowKeyRegexp`, and `DenyKeyRegexp`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority. See also: Restricting agent checks.

AllowKeyRegexp

Allow the execution of those item keys that match a pattern. The key pattern is a regular expression.
 Multiple key matching rules may be defined in combination with `DenyKeyRegexp`, `AllowKey`, and `DenyKey`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority. See also: Restricting agent checks.

Example:

```
AllowKeyRegexp=^system\.run\[(?:/usr)?/bin/hostname(?: -[fsd])?\]$
```

AllowRoot

Allow the agent to run as `root`. If disabled and the agent is started by `root`, the agent will try to switch to user `zabbix` instead. Has no effect if started under a regular user.

Default: 0
 Values: 0 - do not allow; 1 - allow

BufferSend

Set the frequency of sending values from the buffer to Zabbix server/proxy (in seconds).

Default: 5
 Range: 1-3600

BufferSize

Set the maximum number of values in the memory buffer.
 The agent will send all collected data to the Zabbix server or proxy if the buffer is full.

Default: 100
 Range: 2-65535

DebugLevel

Set the debug level: - *0* - basic information about starting and stopping of Zabbix processes; - *1* - critical information; - *2* - error information; - *3* - warnings; - *4* - for debugging (produces lots of information); - *5* - extended debugging (produces even more information).

Default: 3
 Range: 0-5

DenyKey

Deny the execution of those item keys that match a pattern. The key pattern is a wildcard expression that supports the `*` character to match any number of any characters.
 Multiple key matching rules may be defined in combination with `AllowKey`, `AllowKeyRegexp`, and `DenyKeyRegexp`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority. See also: Restricting agent checks.

Default: `DenyKey=system.run[*]`

DenyKeyRegexp

Deny the execution of those item keys that match a pattern. The key pattern is a regular expression.
 Multiple key matching rules may be defined in combination with `AllowKeyRegexp`, `AllowKey`, and `DenyKey`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority.
 All `system.run[]` items are denied by default (`DenyKey=system.run[*]`), unless another `system.run[*]` rule is specified.

See also: Restricting agent checks.

EnableRemoteCommands

Allow remote commands from Zabbix server.
 This parameter is **deprecated**, use `AllowKey=system.run[*]`, `AllowKeyRegexp=^sys` `DenyKey=system.run[*]`, or `DenyKeyRegexp=^system\.run\[.*\]$` instead.
 It is an internal alias for `AllowKey`, `AllowKeyRegexp`, `DenyKey`, and `DenyKeyRegexp` parameters depending on value: - 0 - `DenyKey=system.run[*]` (`DenyKeyRegexp=^system\.run\[.*\]$`) - 1 - `AllowKey=system.run[*]` (`AllowKeyRegexp=^system\.run\[.*\]$`)

Default: 0
 Values: 0 - do not allow, 1 - allow

HeartbeatFrequency

Set the frequency of heartbeat messages in seconds. Used for monitoring the availability of active checks.
0 - heartbeat messages disabled.

Default: 60
 Range: 0-3600

HostInterface

Specify the host interface (IP address or DNS name) used during host autoregistration. This value will be used to populate the interface in the newly created host, and allows configuring either IP or DNS address explicitly. For more details, see Using DNS as default interface.

If not defined, the value will be acquired from `HostInterfaceItem`.

The agent will issue an error and not start if the value is over the limit of 255 characters.

Range: 0-255 characters

HostInterfaceItem

Specify an item to determine the host interface (IP address or DNS name) during host autoregistration. This value is used only if `HostInterface` is not defined. For more details, see Using DNS as default interface.

During an autoregistration request the agent will log a warning message if the value returned by the specified item is over the limit of 255 characters.

The system.run[] item is supported regardless of AllowKey, AllowKeyRegexp, DenyKey, or DenyKeyRegexp settings.

HostMetadata

Specify the metadata to identify or distinguish the host during host autoregistration (active agent). HostMetadata allows distinguishing between hosts beyond the hostname.

If not defined, the value will be acquired from HostMetadataItem.

The agent will issue an error and not start if the specified value is over the limit of 2034 bytes or a non-UTF-8 string. When a parameter expects an IP address or DNS name, values that are valid UTF-8 but are not valid IPs or DNS names will also be rejected and reported as invalid.

Multi-line metadata is not supported—the output will be truncated at the first line break.

Range: 0-2034 bytes

HostMetadataItem

Specify a Zabbix agent item to obtain host metadata. This option is only used when HostMetadata is not defined.

The HostMetadataItem value is retrieved on each autoregistration attempt and is used only at host autoregistration process (active agent). HostMetadataItem allows distinguishing between hosts beyond the hostname.

User parameters and aliases are supported. The system.run[] item is supported regardless of AllowKey, AllowKeyRegexp, DenyKey, or DenyKeyRegexp settings.

During an autoregistration request the agent will log a warning message if the value returned by the specified item is over the limit of 65535 UTF-8 code points. The value returned by the item must be a UTF-8 string; otherwise, it will be ignored. When a parameter expects an IP address or DNS name, values that are valid UTF-8 but are not valid IPs or DNS names will also be rejected and reported as invalid.

Multi-line metadata is not supported—the output will be truncated at the first line break.

Hostname

Specify unique, case-sensitive hostnames. A comma-delimited list.
 Required for active checks and must match hostnames as configured on the server.
 The value is acquired from HostnameItem if undefined.
Allowed characters: alphanumeric, '.', ' ', '_' and '-'. Maximum length: 128 characters per hostname, 2048 characters for the entire line.

Default: Set by HostnameItem

HostnameItem

Specify a Zabbix agent item for getting the hostname.
 This option is only used when Hostname is not defined.
 User parameters or aliases are not supported, but the system.run[] item is supported regardless of AllowKey, AllowKeyRegexp, DenyKey, or DenyKeyRegexp values.

Default: system.hostname

Include

Specify a directory or individual files to be included in the configuration file.
 During Zabbix installation, this parameter is set to /usr/local/etc, unless you change this during compile time.

General rules for the Include parameter:

- The path can be relative to the zabbix_agentd.conf file location.
- You can specify multiple Include parameters in one configuration file.
- Zabbix processes the Include parameter as soon as it reads it, so consider placing it at the end of the configuration file.

Rules for included files:

- Each file must be readable.
- Each file must have correct syntax. Otherwise, the component will not start.

Rules for including directories or file patterns:

- Do not assume any order of inclusion (e.g., files are not included in alphabetical order).
- Do not define the same parameter in more than one included file (e.g., to override a general setting with a specific one), since you cannot predict which value will be used.
- Watch out for backup copies that some text editors create automatically (e.g., editing include/my_specific.conf may create include/my_specific.conf.BAK), since both files will be included.

To include only relevant files from a directory, use the wildcard (*) for pattern matching.

Example:

```
Include=/absolute/path/to/config/files/*.conf
```

ListenBacklog

Set the maximum number of pending connections in the TCP queue.
 The default value is a hard-coded constant, which depends on the system.
 The maximum supported value depends on the system, too high values may be silently truncated to the 'implementation-specified maximum'.

Default: SOMAXCONN
 Range: 0 - INT_MAX

ListenIP

Specify the agent listen IP address for connections from the server/proxy. A comma-delimited list.

Default: 0.0.0.0

ListenPort

Specify the agent listen port for connections from the server/proxy.

Default: 10050
 Range: 1024-32767

LoadModule

Specify the module to load at agent startup.
 Modules are used to extend the functionality of the agent.
 The module must be located in the directory specified by LoadModulePath or the path must precede the module name. If the preceding path is absolute (starts with /) then LoadModulePath is ignored.
 Formats: - LoadModule=<module.so> - LoadModule=<path/module.so> - LoadModule=</abs_path/module.so>

It is allowed to include multiple LoadModule parameters.

LoadModulePath

Specify full path to the location of agent modules.
 The default depends on compilation options.

LogFile

Specify the agent log file.

Mandatory: Yes, if LogType is set to file; otherwise, no

LogFileSize

Set the maximum log file size (in MB).
0 - disable automatic log rotation.
Note: If the log file size limit is reached and file rotation fails, for whatever reason, the existing log file is truncated and started anew.

Default: 1
 Range: 0-1024

LogRemoteCommands

Enable logging of the executed shell commands as warnings.
 Commands will be logged only if executed remotely.
 Log entries will not be created if system.run[] is launched locally by HostMetadataItem, HostInterfaceItem or HostnameItem parameters.

Default: 0
 Values: 0 - disabled, 1 - enabled

LogType

Specify the log output type: - *file* - write log to the file specified by LogFile parameter; - *system* - write log to syslog; - *console* - write log to standard output.

Default: file

MaxLinesPerSecond

Set the maximum number of new lines per second for sending to Zabbix server/proxy when processing log or logrt active checks.
 The provided value will be overridden by the maxlines parameter, provided in the log or logrt item key.
Note: Zabbix will process 10 times more new lines than set in MaxLinesPerSecond to seek the required string in log items.

Default: 20
 Range: 1-1000

PidFile

Specify the PID file.

Default: /tmp/zabbix_agentd.pid

RefreshActiveChecks

Specify the frequency of refreshing the list of active checks, in seconds.
 Note that after failing to refresh active checks the next refresh will be attempted in 60 seconds.

Default: 5
 Range: 1-86400

Server

Specify the IP addresses, optionally in CIDR notation, or DNS names of Zabbix servers/proxies. A comma-delimited list.
 Incoming connections will be accepted only from the hosts listed here. If IPv6 support is enabled then `127.0.0.1`, `::127.0.0.1`, `::ffff:127.0.0.1` are treated equally and `::/0` will allow any IPv4 or IPv6 address. `0.0.0.0/0` can be used to allow any IPv4 address. Note that "IPv4-compatible IPv6 addresses" (`0000::/96` `prefix`) are supported but deprecated by RFC4291. Spaces are allowed.

Example:

```
Server=127.0.0.1,192.168.1.0/24,::1,2001:db8::/32,zabbix.example.com
```

Mandatory: yes, if StartAgents is not explicitly set to 0

ServerActive

Specify the Zabbix server/proxy address or cluster configuration to get active checks from.
 The server/proxy address is an IP address or DNS name and optional port separated by colon.
Cluster configuration is one or more server or proxy group member addresses separated by semicolon. Multiple Zabbix servers/clusters and Zabbix proxies can be specified, separated by comma. Unless using proxy groups, more than one Zabbix proxy should not be specified from each Zabbix server/cluster. If Zabbix proxy is specified then Zabbix server/cluster for that proxy should not be specified.
Multiple comma-delimited addresses can be provided to use several independent Zabbix servers in parallel. Spaces are allowed.
If the port is not specified, default port is used.
IPv6 addresses must be enclosed in square brackets if port for that host is specified. If port is not specified, square brackets for IPv6 addresses are optional.
If this parameter is not specified, active checks are disabled.

Example for Zabbix proxy:

```
ServerActive=127.0.0.1:10051
```

Example for Zabbix proxy group:

```
ServerActive=proxy1.example.com;proxy2.example.com;proxy3.example.com;proxy4.example.com;proxy5.example.co
```

Example for multiple servers:

```
ServerActive=127.0.0.1:20051,zabbix.domain,[::1]:30051,::1,[12fc::1]
```

Example for high availability:

```
ServerActive=zabbix.cluster.node1;zabbix.cluster.node2:20051;zabbix.cluster.node3
```

Example for high availability with two clusters and one server:

```
ServerActive=zabbix.cluster.node1;zabbix.cluster.node2:20051,zabbix.cluster2.node1;zabbix.cluster2.node2,z
```

SourceIP

Specify the source IP address for:

- outgoing connections to Zabbix server or Zabbix proxy;
- making connections while executing some items (`web.page.get`, `net.tcp.port`, etc.).

StartAgents

Set the number of pre-forked instances of `zabbix_agentd` for processing passive checks.
 If set to 0, passive checks are disabled and the agent will not listen on any TCP port.

Default: 10
 Range: 0-100

Timeout

Specify how long to wait (in seconds) for establishing connection and exchanging data with Zabbix proxy or server.

This parameter defines the duration of various communication operations, including:

- awaiting a response from the Zabbix server;
- sending requests to Zabbix server, including item configuration requests and item data in active checks;
- retrieving log data through logfile;
- sending heartbeat messages;
- maximum duration for `vfs.*` checks;

- being used by Zabbix agent modules;
- serving as a fallback in scenarios where a server or proxy older than version 7.0 sends checks without timeouts.

This timeout will **not** be used for those agent checks that have configurable timeout settings in the frontend (on global, proxy, or per-item level).

Default: 3
 Range: 1-30

TLSAccept

Specify the level of encryption required for incoming connections.
 Used for passive checks.
 Multiple values can be specified, separated by comma: - *unencrypted* - accept connections without encryption (default); - *psk* - accept connections with TLS and a pre-shared key (PSK); - *cert* - accept connections with TLS and a certificate.

Mandatory: yes, if TLS certificate or PSK parameters are defined (even for *unencrypted* connection); otherwise, no

TLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.

TLSCertFile

Specify full pathname of the file containing the agent certificate or certificate chain, used for encrypted communications with Zabbix components.

TLSCipherAll

Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.

Example:

```
TLS_AES_256_GCM_SHA384:TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256
```

TLSCipherAll13

Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.

Example for GnuTLS:

```
NONE:+VERS-TLS1.2:+ECDHE-RSA:+RSA:+ECDHE-PSK:+PSK:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+SHA1:+CURVE-ALL
```

Example for OpenSSL:

```
EECDH+aRSA+AES128:RSA+aRSA+AES128:kECDHEPSK+AES128:kPSK+AES128
```

TLSCipherCert

Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for certificate-based encryption.

Example for GnuTLS:

```
NONE:+VERS-TLS1.2:+ECDHE-RSA:+RSA:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+SHA1:+CURVE-ALL:+COMP-NULL:+SIG
```

Example for OpenSSL:

```
EECDH+aRSA+AES128:RSA+aRSA+AES128
```

TLSCipherCert13

Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for certificate-based encryption.

TLSCipherPSK

Specify the GnuTLS priority string or OpenSSL (TLS 1.2) cipher string. Override the default ciphersuite selection criteria for PSK-based encryption.

Example for GnuTLS:

```
NONE:+VERS-TLS1.2:+ECDHE-PSK:+PSK:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+SHA1:+CURVE-ALL:+COMP-NULL:+SIG
```

Example for OpenSSL:

```
kECDHEPSK+AES128:kPSK+AES128
```

TLSCipherPSK13

Specify the cipher string for OpenSSL 1.1.1 or newer in TLS 1.3. Override the default ciphersuite selection criteria for PSK-based encryption.

Example:

```
TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256
```

TLSConnect

Specify the level of encryption required for agent connections to Zabbix server or proxy.
 Used for active checks.
 Only one value can be specified: - *unencrypted* - connect without encryption (default); - *psk* - connect using TLS and a pre-shared key (PSK); - *cert* - connect using TLS and a certificate.

Mandatory: yes, if TLS certificate or PSK parameters are defined (even for *unencrypted* connection); otherwise, no

TLSCRLFile

Specify full pathname of the file containing revoked certificates. This parameter is used for encrypted communications between Zabbix components.

TLSKeyFile

Specify full pathname of the file containing the agent private key, used for encrypted communications between Zabbix components.

TLSPSKFile

Specify full pathname of the file containing the agent pre-shared key, used for encrypted communications with Zabbix server or proxy.

TLSPSKIdentity

Specify the pre-shared key identity string, used for encrypted communications with Zabbix server or proxy.

TLSServerCertIssuer

Specify the allowed server (proxy) certificate issuer.

TLSServerCertSubject

Specify the allowed server (proxy) certificate subject.

UnsafeUserParameters

Allow all characters to be passed in arguments to user-defined parameters. The following characters are not allowed: \ ' " ` * ? [ ] { } ~ $! & ; ( ) < > | # @ %
Additionally, newline characters are not allowed.

Default: 0
 Values: 0 - do not allow, 1 - allow

User

Drop privileges to a specific, existing user on the system.
Only has effect if run as `root` and `AllowRoot` is disabled.

Default: `zabbix`

UserParameter

Specify a user-defined parameter to monitor. There can be several user-defined parameters.
Format: `UserParameter=<key>,<shell command>`
Note that the shell command must not return empty string or EOL only. Shell commands may have relative paths, if the `UserParameterDir` parameter is specified.

Example:

```
UserParameter=system.test,who|wc -l
UserParameter=check_cpu,./custom_script.sh
```

UserParameterDir

Specify the default search path for `UserParameter` commands. If used, the agent will change its working directory to the one specified here before executing a command. Thereby, `UserParameter` commands can have a relative `./` prefix instead of a full path.
Only one entry is allowed.

Example:

```
UserParameterDir=/opt/myscripts
```

See also

1. Differences in the Zabbix agent configuration for active and passive checks starting from version 2.0.0

Configuration parameters: Zabbix agent (Windows)

Overview

Configuration parameters allow customizing Zabbix agent behavior on Windows.

Specify parameter values in the Zabbix agent configuration file (`zabbix_agentd.conf`). All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name).
- Parameter details.

Parameter	Description
Alias	Set an alias for an item key.
AllowKey	Allow the execution of those item keys that match a pattern.
AllowKeyRegexp	Allow the execution of item keys that match a regular expression-based pattern.
BufferSend	Set the frequency of sending values from the buffer to Zabbix server/proxy.
BufferSize	Set the maximum number of values in the memory buffer.
DebugLevel	Set the debug level.
DenyKey	Deny the execution of those item keys that match a pattern.
DenyKeyRegexp	Deny the execution of item keys that match a regular expression-based pattern.
EnableRemoteCommands	Allow remote commands from Zabbix server. This parameter is deprecated; use `AllowKey`, `AllowKeyRegexp`, `DenyKey`, or `DenyKeyRegexp` instead.
HeartbeatFrequency	Set the frequency of heartbeat messages in seconds.
HostInterface	Specify the host interface.
HostInterfaceItem	Specify an item for getting the host interface.
HostMetadata	Specify the host metadata.
HostMetadataItem	Specify a Zabbix agent item for getting the host metadata.
Hostname	Specify the hostname.
HostnameItem	Specify a Zabbix agent item for getting the hostname.
Include	Specify a directory or individual files to be included in the configuration file.
ListenBacklog	Set the maximum number of pending connections in the TCP queue.
ListenIP	Specify the agent listen IP address for connections from the server/proxy.
ListenPort	Specify the agent listen port.
LogFile	Specify the agent log file.
LogFileSize	Set the maximum log file size.
LogRemoteCommands	Enable logging of executed shell commands as warnings.
LogType	Specify the log output type.
MaxLinesPerSecond	Set the maximum number of new lines per second for sending to Zabbix server/proxy when processing `log`, `logrt`, or `eventlog` active checks.
PerfCounter	Define a new parameter <parameter_name> which is the average value for system performance counter <perf_counter_path> for the specified time period <period> (in seconds).
PerfCounterEn	Define a new parameter <parameter_name> which is the average value for system performance counter <perf_counter_path> for the specified time period <period> (in seconds). Compared to PerfCounter, the perfcounter paths must be in English.
RefreshActiveChecks	Specify the frequency of refreshing the list of active checks.
Server	Specify the IP addresses or DNS names of Zabbix servers/proxies.
ServerActive	Specify the Zabbix server/proxy address or cluster configuration to get active checks from.
SourceIP	Specify the source IP address.
StartAgents	Set the number of pre-forked instances of `zabbix_agentd` for processing passive checks.
Timeout	Specify how long to wait (in seconds) for establishing connection and exchanging data with Zabbix server or proxy.
TLSAccept	Specify the level of encryption required for incoming connections.
TLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.
TLSCertFile	Specify full pathname of the file containing the agent certificate or certificate chain, used for encrypted communications between Zabbix components.
TLSConnect	Specify the level of encryption required for connections to Zabbix server or proxy.
TLSCRLFile	Specify full pathname of the file containing revoked certificates. This parameter is used for encrypted communications between Zabbix components.

Parameter	Description
TLSKeyFile	Specify full pathname of the file containing the agent private key, used for encrypted communications between Zabbix components.
TLSPSKFile	Specify full pathname of the file containing the agent pre-shared key, used for encrypted communications with Zabbix server or proxy.
TLSPSKIdentity	Specify the pre-shared key identity string, used for encrypted communications with Zabbix server or proxy.
TLSServerCertIssuer	Specify the allowed server (proxy) certificate issuer.
TLSServerCertSubject	Specify the allowed server (proxy) certificate subject.
UnsafeUserParameters	Allow all characters to be passed in arguments to user-defined parameters.
UserParameter	Specify a user-defined parameter to monitor.
UserParameterDir	Specify the default search path for UserParameter commands.

Note that:

- The default values reflect daemon defaults, not the values in the shipped configuration files.
- Values support environment variables.
- Zabbix supports configuration files only in UTF-8 encoding without BOM.
- Comments starting with # are only supported in the beginning of the line.

Parameter details

Alias

Set an alias for an item key. It can be used to substitute a long and complex item key with a shorter and simpler one.
 Multiple `Alias` parameters may be present. Multiple parameters with the same `Alias` key are not allowed.
 Different `Alias` keys may reference the same item key.
 Aliases can be used in `HostMetadataItem` but not in `HostnameItem` or `PerfCounter` parameter.

Example 1: Retrieving the paging file usage in percentage from the server.

```
Alias=pg_usage:perf_counter[\Paging File(_Total)\% Usage]
```

Now the shorthand key `pg_usage` may be used to retrieve data.

Example 2: Getting the CPU load with default and custom parameters.

```
Alias=cpu.load:system.cpu.load
Alias=cpu.load[*]:system.cpu.load[*]
```

This allows you to use `cpu.load` key to get the CPU load with default parameters as well as use `cpu.load[percpu,avg15]` to get specific data about the CPU load.

Example 3: Running multiple low-level discovery rules processing the same discovery items.

```
Alias=vfs.fs.discovery[*]:vfs.fs.discovery
```

Now it is possible to set up several discovery rules using `vfs.fs.discovery` with different parameters for each rule, e.g., `vfs.fs.discovery[foo]`, `vfs.fs.discovery[bar]`, etc.

AllowKey

Allow the execution of those item keys that match a pattern. The key pattern is a wildcard expression that supports the * character to match any number of any characters.
 Multiple key matching rules may be defined in combination with `DenyKey`, `AllowKeyRegexp`, and `DenyKeyRegexp`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority. See also: Restricting agent checks.

AllowKeyRegexp

Allow the execution of those item keys that match a pattern. The key pattern is a regular expression.
 Multiple key matching rules may be defined in combination with `DenyKeyRegexp`, `AllowKey`, and `DenyKey`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority. See also: Restricting agent checks.

Example:

```
AllowKeyRegexp=^system\.run\[(?:/usr)?/bin/hostname(?: -[fsd])?\]$
```

BufferSend

Set the frequency of sending values from the buffer to Zabbix server/proxy (in seconds).

Default: 5
 Range: 1-3600

BufferSize

Set the maximum number of values in the memory buffer.
 The agent will send all collected data to the Zabbix server or proxy if the buffer is full.

Default: 100
 Range: 2-65535

DebugLevel

Set the debug level: - *0* - basic information about starting and stopping of Zabbix processes; - *1* - critical information; - *2* - error information; - *3* - warnings; - *4* - for debugging (produces lots of information); - *5* - extended debugging (produces even more information).

Default: 3
 Range: 0-5

DenyKey

Deny the execution of those item keys that match a pattern. The key pattern is a wildcard expression that supports the * character to match any number of any characters.
 Multiple key matching rules may be defined in combination with `AllowKey`, `AllowKeyRegexp`, and `DenyKeyRegexp`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority. See also: Restricting agent checks.

Default: `DenyKey=system.run[*]`

DenyKeyRegexp

Deny the execution of those item keys that match a pattern. The key pattern is a regular expression.
 Multiple key matching rules may be defined in combination with `AllowKeyRegexp`, `AllowKey`, and `DenyKey`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority.
 All `system.run[]` items are denied by default (`DenyKey=system.run[*]`), unless another `system.run[*]` rule is specified.

See also: Restricting agent checks.

EnableRemoteCommands

Allow remote commands from Zabbix server.
 This parameter is **deprecated**, use `AllowKey=system.run[*]`, `AllowKeyRegexp=^sys` `DenyKey=system.run[*]`, or `DenyKeyRegexp=^system\.run\[.*\]$` instead.
 It is an internal alias for `AllowKey`, `AllowKeyRegexp`, `DenyKey`, and `DenyKeyRegexp` parameters depending on value: - 0 - `DenyKey=system.run[*]` (`DenyKeyRegexp=^system\.run\[.*\]$`) - 1 - `AllowKey=system.run[*]` (`AllowKeyRegexp=^system\.run\[.*\]$`)

Default: 0
 Values: 0 - do not allow, 1 - allow

HeartbeatFrequency

Set the frequency of heartbeat messages in seconds. Used for monitoring the availability of active checks.
0 - heartbeat messages disabled.

Default: 60
 Range: 0-3600

HostInterface

Specify the host interface (IP address or DNS name) used during host autoregistration. This value will be used to populate the interface in the newly created host, and allows configuring either IP or DNS address explicitly. For more details, see Using DNS as default interface.

If not defined, the value will be acquired from HostInterfaceItem.

The agent will issue an error and not start if the value is over the limit of 255 characters.

Range: 0-255 characters

HostInterfaceItem

Specify an item to determine the host interface (IP address or DNS name) during host autoregistration. This value is used only if `HostInterface` is not defined. For more details, see Using DNS as default interface.

During an autoregistration request the agent will log a warning message if the value returned by the specified item is over the limit of 255 characters.

The `system.run[]` item is supported regardless of `AllowKey`, `AllowKeyRegexp`, `DenyKey`, or `DenyKeyRegexp` settings.

HostMetadata

Specify the metadata used to identify or distinguish the host during host autoregistration (active agent). `HostMetadata` allows distinguishing between hosts beyond the hostname.

If not defined, the value will be acquired from HostMetadataItem.

The agent will issue an error and not start if the specified value is over the limit of 2034 bytes or a non-UTF-8 string. When a parameter expects an IP address or DNS name, values that are valid UTF-8 but are not valid IPs or DNS names will also be rejected and reported as invalid.

Multi-line metadata is not supported—the output will be truncated at the first line break.

Range: 0-2034 bytes

HostMetadataItem

Specify a Zabbix agent item to obtain host metadata. This option is only used when HostMetadata is not defined.

The HostMetadataItem value is retrieved on each autoregistration attempt and is used only at host autoregistration process (active agent). HostMetadataItem allows distinguishing between hosts beyond the hostname.

User parameters, performance counters and aliases are supported. The system.run[] item is supported regardless of AllowKey, AllowKeyRegexp, DenyKey, or DenyKeyRegexp settings.

During an autoregistration request the agent will log a warning message if the value returned by the specified item is over the limit of 65535 UTF-8 code points. The value returned by the item must be a UTF-8 string; otherwise, it will be ignored. When a parameter expects an IP address or DNS name, values that are valid UTF-8 but are not valid IPs or DNS names will also be rejected and reported as invalid.

Multi-line metadata is not supported—the output will be truncated at the first line break.

Hostname

Specify unique, case-sensitive hostnames. A comma-delimited list.
 Required for active checks and must match hostnames as configured on the server. The value is acquired from HostnameItem if undefined.
Allowed characters: alphanumeric, '.', ' ', '_' and '-'. Maximum length: 128 characters per hostname, 2048 characters for the entire line.

Default: Set by HostnameItem

HostnameItem

Specify a Zabbix agent item for getting the hostname.
 This option is only used when Hostname is not defined.
 User parameters, performance counters or aliases are not supported, but the system.run[] item is supported regardless of AllowKey, AllowKeyRegexp, DenyKey, or DenyKeyRegexp values.
See also a more detailed description.

Default: system.hostname

Include

Specify a directory or individual files to be included in the configuration file (located in C:\Program Files\Zabbix Agent by default if Zabbix agent is installed using Windows MSI installer packages; located in the folder specified during installation if Zabbix agent is installed as a zip archive).

General rules for the Include parameter:

- The path can be relative to the zabbix_agentd.conf file location.
- You can specify multiple Include parameters in one configuration file.
- Zabbix processes the Include parameter as soon as it reads it, so consider placing it at the end of the configuration file.

Rules for included files:

- Each file must be readable.
- Each file must have correct syntax. Otherwise, the component will not start.

Rules for including directories or file patterns:

- Do not assume any order of inclusion (e.g., files are not included in alphabetical order).
- Do not define the same parameter in more than one included file (e.g., to override a general setting with a specific one), since you cannot predict which value will be used.
- Watch out for backup copies that some text editors create automatically (e.g., editing include/my_specific.conf may create include/my_specific.conf.BAK), since both files will be included.

To include only relevant files from a directory, use the wildcard (*) for pattern matching.

Example:

```
Include=C:\Program Files\Zabbix Agent\zabbix_agentd.d\*.conf
```

ListenBacklog

Set the maximum number of pending connections in the TCP queue.
 The default value is a hard-coded constant, which depends on the system.
 The maximum supported value depends on the system, too high values may be silently truncated to the 'implementation-specified maximum'.

Default: SOMAXCONN
 Range: 0 - INT_MAX

ListenIP

Specify the agent listen IP address for connections from the server/proxy. A comma-delimited list.

Default: 0.0.0.0

ListenPort

Specify the agent listen port for connections from the server/proxy.

Default: 10050
 Range: 1024-32767

LogFile

Specify the agent log file.

Default: c:\zabbix_agentd.log
 Mandatory: Yes, if LogType is set to file; otherwise, no

LogFileSize

Set the maximum log file size (in MB).
0 - disable automatic log rotation.
Note: If the log file size limit is reached and file rotation fails, for whatever reason, the existing log file is truncated and started anew.

Default: 1
 Range: 0-1024

LogRemoteCommands

Enable the logging of the executed shell commands as warnings.
 Commands will be logged only if executed remotely.
 Log entries will not be created if system.run[] is launched locally by HostMetadataItem, HostInterfaceItem or HostnameItem parameters.

Default: 0
 Values: 0 - disabled, 1 - enabled

LogType

Specify the log output type: - *file* - write log to the file specified by LogFile parameter; - *system* - write log to Windows Event Log; - *console* - write log to standard output.

Default: file

MaxLinesPerSecond

Set the maximum number of new lines per second for sending to Zabbix server/proxy when processing log, logrt, or eventlog active checks.
 The provided value will be overridden by the maxlines parameter, provided in the log, logrt, or eventlog item key.
Note: Zabbix will process 10 times more new lines than set in MaxLinesPerSecond to seek the required string in log items.

Default: 20
 Range: 1-1000

PerfCounter

Define a new parameter <parameter_name> which is the average value for system performance counter <perf_counter_path> for the specified time period <period> (in seconds).
Syntax: <parameter_name>,"<perf_counter_path>",<period>

For example, if you wish to receive the average number of processor interrupts per second for the last minute, you can define a new parameter "interrupts" as the following:

```
PerfCounter = interrupts,"\Processor(0)\Interrupts/sec",60
```

Please note the double quotes around the performance counter path. The parameter name (interrupts) is to be used as the item key when creating an item. Samples for calculating the average value will be taken every second.
You may run typeperf -qx to get the list of all performance counters available in Windows.

PerfCounterEn

Define a new parameter <parameter_name> which is the average value for system performance counter <perf_counter_path> for the specified time period <period> (in seconds). Compared to PerfCounter, the perfcounter paths must be in English. Supported only on **Windows Server 2008/Vista** and later.
Syntax: <parameter_name>,"<perf_counter_path>",<period>

For example, if you wish to receive the average number of processor interrupts per second for the last minute, you can define a new parameter "interrupts" as the following:

```
PerfCounterEn = interrupts,"\Processor(0)\Interrupts/sec",60
```

Please note the double quotes around the performance counter path. The parameter name (interrupts) is to be used as the item key when creating an item. Samples for calculating the average value will be taken every second.
You can find the list of English strings by viewing the following registry key: `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib\009`.

RefreshActiveChecks

Specify the frequency of refreshing the list of active checks, in seconds.
 Note that after failing to refresh active checks the next refresh will be attempted in 60 seconds.

Default: 5
 Range: 1-86400

Server

Specify the IP addresses, optionally in CIDR notation, or DNS names of Zabbix servers/proxies. A comma-delimited list.
 Incoming connections will be accepted only from the hosts listed here. If IPv6 support is enabled then `127.0.0.1`, `::127.0.0.1`, `::ffff:127.0.0.1` are treated equally and `::/0` will allow any IPv4 or IPv6 address. `0.0.0.0/0` can be used to allow any IPv4 address. Note that "IPv4-compatible IPv6 addresses" (`0000::/96 prefix`) are supported but deprecated by RFC4291. Spaces are allowed.

Example:

```
Server=127.0.0.1,192.168.1.0/24,::1,2001:db8::/32,zabbix.example.com
```

Mandatory: yes, if StartAgents is not explicitly set to 0

ServerActive

Specify the Zabbix server/proxy address or cluster configuration to get active checks from.
 The server/proxy address is an IP address or DNS name and optional port separated by colon.
The cluster configuration is one or more server or proxy group member addresses separated by semicolon. Multiple Zabbix servers/clusters and Zabbix proxies can be specified, separated by comma. Unless using proxy groups, more than one Zabbix proxy should not be specified from each Zabbix server/cluster. If Zabbix proxy is specified then Zabbix server/cluster for that proxy should not be specified.
Multiple comma-delimited addresses can be provided to use several independent Zabbix servers in parallel. Spaces are allowed.
If the port is not specified, default port is used.
IPv6 addresses must be enclosed in square brackets if port for that host is specified. If port is not specified, square brackets for IPv6 addresses are optional.
If this parameter is not specified, active checks are disabled.

Example for Zabbix proxy:

```
ServerActive=127.0.0.1:10051
```

Example for Zabbix proxy group:

```
ServerActive=proxy1.example.com;proxy2.example.com;proxy3.example.com;proxy4.example.com;proxy5.example.co
```

Example for multiple servers:

```
ServerActive=127.0.0.1:20051,zabbix.domain,[::1]:30051,::1,[12fc::1]
```

Example for high availability:

```
ServerActive=zabbix.cluster.node1;zabbix.cluster.node2:20051;zabbix.cluster.node3
```

Example for high availability with two clusters and one server:

```
ServerActive=zabbix.cluster.node1;zabbix.cluster.node2:20051,zabbix.cluster2.node1;zabbix.cluster2.node2,z
```

Range: The number of active servers listed in `ServerActive` plus the number of pre-forked instances for passive checks specified in `StartAgents` must be less than 64.

SourceIP

Specify the source IP address for:

- outgoing connections to Zabbix server or Zabbix proxy;
- making connections while executing some items (`web.page.get`, `net.tcp.port`, etc.).

StartAgents

Set the number of pre-forked instances of `zabbix_agentd` for processing passive checks.
 If set to 0, passive checks are disabled and the agent will not listen on any TCP port.

Default: 10
 Range: 0-100 (*)

Timeout

Specify how long to wait (in seconds) for establishing connection and exchanging data with Zabbix proxy or server.

This parameter defines the duration of various communication operations, including:

- awaiting a response from the Zabbix server;
- sending requests to Zabbix server, including item configuration requests and item data in active checks;
- retrieving log data through logfile or Windows event log monitoring;
- sending heartbeat messages;
- maximum duration for `vfs.*` checks;
- used by Zabbix agent modules;
- serving as a fallback in scenarios where a server or proxy older than version 7.0 sends checks without timeouts.

This timeout will **not** be used for those agent checks that have configurable timeout settings in the frontend (on global, proxy, or per-item level).

Default: 3
 Range: 1-30

TLSAccept

Specify the level of encryption required for incoming connections.
 Used for passive checks.
 Multiple values can be specified, separated by comma: - *unencrypted* - accept connections without encryption (default); - *psk* - accept connections with TLS and a pre-shared key (PSK); - *cert* - accept connections with TLS and a certificate.

Mandatory: yes, if TLS certificate or PSK parameters are defined (even for *unencrypted* connection); otherwise, no

TLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.

TLSCertFile

Specify full pathname of the file containing the agent certificate or certificate chain, used for encrypted communications with Zabbix components.

TLSConnect

Specify the level of encryption required for agent connections to Zabbix server or proxy.
 Used for active checks.
 Only one value can be specified: - *unencrypted* - connect without encryption (default); - *psk* - connect using TLS and a pre-shared key (PSK); - *cert* - connect using TLS and a certificate.

Mandatory: yes, if TLS certificate or PSK parameters are defined (even for *unencrypted* connection); otherwise, no

TLSCRLFile

Specify full pathname of the file containing revoked certificates. This parameter is used for encrypted communications between Zabbix components.

TLSKeyFile

Specify full pathname of the file containing the agent private key, used for encrypted communications between Zabbix components.

TLSPSKFile

Specify full pathname of the file containing the agent pre-shared key, used for encrypted communications with Zabbix server or proxy.

TLSPSKIdentity

Specify the pre-shared key identity string, used for encrypted communications with Zabbix server or proxy.

TLSServerCertIssuer

Specify the allowed server (proxy) certificate issuer.

TLSServerCertSubject

Specify the allowed server (proxy) certificate subject.

UnsafeUserParameters

Allow all characters to be passed in arguments to user-defined parameters. The following characters are not allowed: \ ' " ` * ? [ ] { } ~ $! & ; ( ) < > | # @ %
Additionally, newline characters are not allowed.

Default: 0
 Values: 0 - do not allow, 1 - allow

UserParameter

Specify a user-defined parameter to monitor. There can be several user-defined parameters.
Format: `UserParameter=<key>,<shell command>`
Note that the shell command must not return empty string or EOL only. Shell commands may have relative paths, if the `UserParameterDir` parameter is specified.

Example:

```
UserParameter=system.test,dir /b | find /c /v ""
UserParameter=check_cpu,.\custom\_script.bat
```

UserParameterDir

Specify the default search path for `UserParameter` commands. If used, the agent will change its working directory to the one specified here before executing a command. Thereby, `UserParameter` commands can have a relative `./` prefix instead of a full path. Only one entry is allowed.

Example:

```
UserParameterDir=C:\opt\myscripts
```

Zabbix agent on Microsoft Windows

Configuring agent

Both generations of Zabbix agents run as a Windows service. For Zabbix agent 2, replace `agentd` with `agent2` in the instructions below.

You can run a single instance of Zabbix agent or multiple instances of the agent on a Microsoft Windows host. A single instance may use either:

- the default configuration file, located in the same directory as the agent binary;
- a configuration file specified in the command line.

In case of multiple instances each agent instance must have its own configuration file (one of the instances can use the default configuration file).

An example configuration file is available in the Zabbix source archive as:

- `conf/zabbix_agentd.conf` for Zabbix agent;
- `conf/zabbix_agent2.conf` for Zabbix agent2.

If you want to install Zabbix agent/agent 2 for Windows as a service from an archive without specifying the configuration file explicitly, then, before installing the agent:

- `conf/zabbix_agentd.conf` should be copied manually to the directory where `zabbix_agentd.exe` will be installed;
- `conf/zabbix_agent2.conf` and the `conf/zabbix_agent2.d` directory should be copied manually to the directory where `zabbix_agent2.exe` will be installed.

See the Zabbix agent configuration file or Zabbix agent 2 configuration file options for details on configuring Zabbix Windows agent.

Hostname parameter

To perform active checks on a host Zabbix agent needs to have the hostname defined. Moreover, the hostname value set on the agent side should exactly match the *Host name* configured for the host in the frontend.

The hostname value on the agent side can be defined by either the `Hostname` or `HostnameItem` parameter in the agent configuration file - or the default values are used if any of these parameters are not specified.

The default value for `HostnameItem` parameter is the value returned by the `system.hostname` agent key. For Windows, it returns result of the `gethostname()` function, which queries namespace providers to determine the local host name. If no namespace provider responds, the NetBIOS name is returned.

The default value for `Hostname` is the value returned by the `HostnameItem` parameter. So, in effect, if both these parameters are unspecified, the actual hostname will be the host NetBIOS name; Zabbix agent will use NetBIOS host name to retrieve the list of active checks from Zabbix server and send results to it.

The `system.hostname` key supports two optional parameters - *type* and *transform*.

Type determines the type of the name the item should return:

- *netbios* (default) - returns the NetBIOS host name which is limited to 15 symbols and is in the UPPERCASE only;
- *host* - case-sensitive, returns the full, real Windows host name (without a domain);

- *shorthost* - returns part of the hostname before the first dot. It will return a full string if the name does not contain a dot.
- *fqdn* - returns Fully Qualified Domain Name (without the trailing dot).

Transform allows you to specify additional transformation rule for the hostname:

- *none* (default) - use the original letter case;
- *lower* - convert the text into lowercase.

So, to simplify the configuration of `zabbix_agentd.conf` file and make it unified, different approaches can be used:

1. Leave `Hostname` or `HostnameItem` parameters undefined and Zabbix agent will use NetBIOS host name as the hostname.
2. Leave `Hostname` parameter undefined and define `HostnameItem` like this:
   - `HostnameItem=system.hostname[host]` - for Zabbix agent to use the full, real (case-sensitive) Windows host name as the hostname.
   - `HostnameItem=system.hostname[shorthost,lower]` - for Zabbix agent to use only part of the hostname before the first dot, converted into lowercase.
   - `HostnameItem=system.hostname[fqdn]` - for Zabbix agent to use the Fully Qualified Domain Name as the hostname.

Host name is also used as part of Windows service name which is used for installing, starting, stopping and uninstalling the Windows service. For example, if Zabbix agent configuration file specifies `Hostname=Windows_db_server`, then the agent will be installed as a Windows service `Zabbix Agent [Windows_db_server]`. Therefore, to have a different Windows service name for each Zabbix agent instance, each instance must use a different host name.

Installing agent as Windows service

Before installing the agent, copy `conf/zabbix_agentd.conf` manually to the directory where zabbix_agentd.exe will be installed.

To install a single instance of Zabbix agent with the default configuration file:

```
zabbix_agentd.exe --install
```

Attention:
On a 64-bit system, a 64-bit Zabbix agent version is required for all checks related to running 64-bit processes to work correctly.

If you wish to use a configuration file other than the default one, you should use the following command for service installation:

```
zabbix_agentd.exe --config <your_configuration_file> --install
```

A full path to the configuration file should be specified.

Multiple instances of Zabbix agent can be installed as services like this:

```
zabbix_agentd.exe --config <configuration_file_for_instance_1> --install --multiple-agents
zabbix_agentd.exe --config <configuration_file_for_instance_2> --install --multiple-agents
...
zabbix_agentd.exe --config <configuration_file_for_instance_N> --install --multiple-agents
```

The installed service should now be visible in Control Panel.

Starting agent

To start the agent service, you can use Control Panel or do it from command line.

To start a single instance of Zabbix agent with the default configuration file:

```
zabbix_agentd.exe --start
```

To start a single instance of Zabbix agent with another configuration file:

```
zabbix_agentd.exe --config <your_configuration_file> --start
```

To start one of multiple instances of Zabbix agent:

```
zabbix_agentd.exe --config <configuration_file_for_this_instance> --start --multiple-agents
```

Stopping agent

To stop the agent service, you can use Control Panel or do it from command line.

To stop a single instance of Zabbix agent started with the default configuration file:

```
zabbix_agentd.exe --stop
```

To stop a single instance of Zabbix agent started with another configuration file:

```
zabbix_agentd.exe --config <your_configuration_file> --stop
```

To stop one of multiple instances of Zabbix agent:

```
zabbix_agentd.exe --config <configuration_file_for_this_instance> --stop --multiple-agents
```

Uninstalling agent Windows service

To uninstall a single instance of Zabbix agent using the default configuration file:

```
zabbix_agentd.exe --uninstall
```

To uninstall a single instance of Zabbix agent using a non-default configuration file:

```
zabbix_agentd.exe --config <your_configuration_file> --uninstall
```

To uninstall multiple instances of Zabbix agent from Windows services:

```
zabbix_agentd.exe --config <configuration_file_for_instance_1> --uninstall --multiple-agents
zabbix_agentd.exe --config <configuration_file_for_instance_2> --uninstall --multiple-agents
...
zabbix_agentd.exe --config <configuration_file_for_instance_N> --uninstall --multiple-agents
```

Limitations

Zabbix agent for Windows does not support non-standard Windows configurations where CPUs are distributed non-uniformly across NUMA nodes. If logical CPUs are distributed non-uniformly, then CPU performance metrics may not be available for some CPUs. For example, if there are 72 logical CPUs with 2 NUMA nodes, both nodes must have 36 CPUs each.

Zabbix agent on IBM AIX

Overview

Pre-compiled Zabbix agent binaries for IBM AIX are available for download as a `.tar.gz` archive. This page describes how to install the agent from the archive, run it manually or with the System Resource Controller (SRC), start it at system boot, and remove it from SRC.

Installing agent

The AIX `tar` command does not support the `-z` option, so decompress the archive with `gunzip`. Use the `-o` option of `tar` so that the extracted files are owned by the user who extracts them, not by the numeric user and group IDs stored in the archive. For example, to place the agent in `/opt/zabbix` (run as root; `/opt/zabbix` must not exist yet):

```
gunzip -c zabbix_agent-7.4.7-aix-7.2-powerpc-openssl.tar.gz | tar -xof -
mv zabbix_agent-7.4.7-aix-7.2-powerpc-openssl /opt/zabbix
```

In the AIX archive, the agent binary is located in `usr/sbin/` and the configuration file in `etc/zabbix/`. Before starting the agent, adjust `/opt/zabbix/etc/zabbix/zabbix_agentd.conf` as needed (at least the `Server`, `ServerActive`, and `Hostname` parameters).

If started as root, the agent switches to the `zabbix` user (or the user specified by the `User` parameter) unless `AllowRoot=1` is set in the configuration file. This user must exist on the system; otherwise, the agent does not start. If the `zabbix` user does not exist yet, create it together with a group of the same name:

```
mkgroup zabbix
mkuser pgrp=zabbix login=false rlogin=false zabbix
```

The `login=false` and `rlogin=false` attributes prevent the account from being used for logging in; the agent can still run as this user.

Starting agent manually

The agent can be started manually:

```
/opt/zabbix/usr/sbin/zabbix_agentd -c /opt/zabbix/etc/zabbix/zabbix_agentd.conf
```

Running agent with SRC

Registering agent

The agent can also be registered with the System Resource Controller (SRC). When started through SRC, the agent must run in foreground mode (the `-f` option). Otherwise, SRC loses track of the agent process and cannot stop it. For example (run as root):

```
mkssys \
  -s zabbix_agentd \
  -p /opt/zabbix/usr/sbin/zabbix_agentd \
  -a "-c /opt/zabbix/etc/zabbix/zabbix_agentd.conf -f" \
  -u zabbix \
  -G tcpip \
  -S -n 15 -f 9 \
  -R
```

In this example:

- `-s` sets the subsystem name used with the SRC commands;
- `-p` sets the full path to the agent binary;
- `-a` sets the agent command-line arguments as a single quoted string; SRC splits this string into separate arguments when it starts the agent (only one `-a` option is accepted);
- `-u` sets the user that the agent runs as; the user name is converted to a numeric user ID when the subsystem is created, so the user must already exist;
- `-G` adds the agent to the `tcpip` subsystem group, so that group commands also apply to it (for example, `stopsrc -g tcpip` also stops the agent);
- `-S -n 15 -f 9` makes SRC control the agent with signals: `stopsrc` sends `SIGTERM` (15) for a normal stop and `SIGKILL` (9) for a forced stop (`stopsrc -f`);
- `-R` makes SRC restart the agent if it terminates abnormally (for example, if one of its processes crashes); the agent is not restarted after `stopsrc`.

Since the user ID is stored in the subsystem definition, update the definition if the `zabbix` user is recreated with a different user ID:

```
chssys -s zabbix_agentd -u zabbix
```

Starting and stopping agent

The agent can be started, queried, and stopped with `startsrc`, `lssrc`, and `stopsrc`:

```
startsrc -s zabbix_agentd
lssrc -s zabbix_agentd
stopsrc -s zabbix_agentd
```

Avoid the forced stop (`stopsrc -f`). It kills only the main agent process; the other agent processes keep running and keep the listening port open, so the agent cannot be started again until they are terminated manually (for example, with `kill -9`).

Starting agent at system boot

Registering the agent with SRC does not make it start at system boot. To start the agent automatically at boot, add an entry to `/etc/inittab`:

```
mkitab "zabbix_agentd:2:once:/usr/bin/startsrc -s zabbix_agentd >/dev/console 2>&1"
```

Removing agent from SRC

To remove the agent from system startup and from SRC:

```
stopsrc -s zabbix_agentd
rmitab zabbix_agentd
rmssys -s zabbix_agentd
```

Agent 2

Overview Zabbix agent 2 is a new generation of Zabbix agent, written in Go (with some C code reused from Zabbix agent). It is designed to:

- Reduce the number of TCP connections;
- Provide improved concurrency of checks;
- Be easily extendable with plugins that provide simple checks with minimal code and support complex checks consisting of long-running scripts and standalone data gathering with periodic reporting;
- Function as a replacement for Zabbix agent, supporting all previous features.

All metrics in Zabbix agent 2 are collected by plugins.

These features are specific to plugins:

- Scheduled and flexible intervals for both passive and active checks
- Task queue management with respect to schedule and task concurrency
- Plugin-level timeouts
- Compatibility check of Zabbix agent 2 and its plugins at startup

Two types of plugins exist:

- Built-in plugins
- Loadable plugins

You can also create your own plugins.

Custom-created plugins need not be built into agent 2. You can add them as loadable plugins, thus simplifying the creation of plugins for new metrics.

Loadable plugins can be added on runtime without recompiling Zabbix agent 2.

See also:

- Agent vs agent 2 comparison
- Configuration parameters: Zabbix agent 2
- Configuration parameters: Zabbix agent 2 (Windows)
- Configuration parameters: Plugins
- Zabbix agent 2 plugin protocol

Passive and active checks

Zabbix agent 2 supports passive and active checks, similarly to Zabbix agent. Additionally, Zabbix agent 2 active checks support flexible/scheduling intervals and check concurrency within one active server.

> **Note:**
> By default, after a restart, Zabbix agent 2 will schedule the first data collection for active checks at a conditionally random time within the item's update interval to prevent spikes in resource usage. To perform active checks that do not have *Scheduling* update interval immediately after the agent restart, set the `ForceActiveChecksOnStart` parameter (global-level) or the `Plugins.<Plugin name>.System.ForceActiveChecksOnStart` plugin parameter (affects only specific plugin checks) in the configuration file. Plugin-level parameter, if set, will override the global parameter.

Check concurrency

Checks from different plugins can be executed concurrently. The number of concurrent checks within one plugin is limited by the plugin capacity setting. Each plugin may have a hardcoded capacity setting (1000 being default) that can be lowered using the `Plugins.<PluginName>.System.Capacity=N` setting in the `Plugins` configuration parameter.

Supported platforms For supported platforms, refer to the Requirements page.

Agent 2 on Unix-like systems Zabbix agent 2 on Unix-like systems is run on the host being monitored.

Installation

Zabbix agent 2 can be installed on Linux-based systems using one of the following methods:

- Zabbix packages - select the Agent 2 component (after selecting your Zabbix version, OS distribution, and OS version) and follow the instructions;
- Zabbix sources - download the source files and compile the agent by configuring with the `--enable-agent2` option.

> **Note:**
> The monitoring capabilities of Zabbix agent 2 can be extended with loadable plugins, which are available separately. For details, see Loadable plugins.

If installed as package

Zabbix agent 2 runs as a foreground process and relies on an external service manager (e.g., systemd) for background execution. Zabbix agent 2 does not have built-in daemonization support on Linux.

The agent can be started by executing:

```
systemctl start zabbix-agent2
```

To stop, restart, or check the status of Zabbix agent 2, use the following commands:

```
systemctl stop zabbix-agent2
systemctl restart zabbix-agent2
systemctl status zabbix-agent2
```

Start up manually

You can start Zabbix agent by locating the `zabbix_agent2` binary and running it directly; for example:

```
zabbix_agent2
```

Agent 2 on Windows systems Zabbix agent 2 runs as a standalone process; however, it can also be run as a Windows service.

Installation

Zabbix agent 2 can be installed on Windows using one of the following methods:

- Pre-compiled Zabbix agent binaries - download the agent MSI installer package and follow the instructions on the Windows agent installation from MSI page;
- Zabbix sources - download the source files and follow the instructions on the Building Zabbix agent 2 on Windows page.

Note:
The monitoring capabilities of Zabbix agent 2 can be extended with loadable plugins, which are available separately. For details, see Loadable plugins.

See the Zabbix agent on Microsoft Windows page for additional details on installing Zabbix agent 2 (from a ZIP archive) as a Windows service.

Options The following command-line parameters can be used with Zabbix agent 2:

Parameter	Description
UNIX and Windows agent	
`-c --config <config-file>`	Path to the configuration file. You may use this option to specify a configuration file that is not the default one.
`-f --foreground`	Run Zabbix agent in foreground (default: true).
`-p --print`	Print known items and exit. Note that to return user parameter results as well, you must specify the configuration file (if it is not in the default location).
`-t --test <item key>`	Test specified item and exit. Note that to return user parameter results as well, you must specify the configuration file (if it is not in the default location).
`-T --test-config`	Validate configuration file and exit.
`-h --help`	Print help information and exit.
`-v --verbose`	Print debugging information. Use this option with `-p` and `-t` options.
`-V --version`	Print agent version and license information.
`-R --runtime-control <option>`	Perform administrative functions. See runtime control.
Windows agent only	
`-m --multiple-agents`	Use multiple agent instances (with `-i`, `-d`, `-s`, `-x` options). To distinguish service names of instances, each service name will include the Hostname value from the specified configuration file.
`-S --startup-type <value>`	Set the startup type of the Zabbix Windows agent service. Allowed values: `automatic` - *(default)* start the service automatically at Windows startup; `delayed` - delay starting the service after the automatically started services have completed startup; `manual` - start the service manually (by a user or application); `disabled` - disable the service, so that it cannot be started by a user or application. You may use this option together with the -i option, or separately to modify the startup type of an already installed service.
`-i --install`	Install Zabbix Windows agent as service.

Parameter	Description
-d --uninstall	Uninstall Zabbix Windows agent service.
-s --start	Start Zabbix Windows agent service.
-x --stop	Stop Zabbix Windows agent service.

Specific examples of using command-line parameters:

- Print all built-in agent items with values.
- Test a user parameter with mysql.ping key defined in the specified configuration file.
- Installing a Zabbix Agent service for Windows using the default path to the configuration file.
- Modifying the startup type of an installed Zabbix Agent service for Windows using the configuration file located in the same folder as agent executable.

```
zabbix_agent2 --print
zabbix_agent2 -t "mysql.ping" -c /etc/zabbix/zabbix_agent2.conf
zabbix_agent2.exe -i
zabbix_agent2.exe -c zabbix_agent2.conf -S delayed
```

Runtime control

Runtime control provides some options for remote control.

Option	Description
log_level_increase	Increase log level.
log_level_decrease	Decrease log level.
metrics	List available metrics.
periodic_prof_enable	Enable periodic profiling. All metrics collect data in the moment and dump them, except CPU data. Periodic CPU data is collected over time into a hidden temporary file. When the interval ends, the file is published as a CPU profile and the configured number of completed files is retained. The temporary file does not count toward ProfilerMaxFilesPerProfile.
periodic_prof_disable	Disable periodic profiling.
periodic_prof_execute	Write profiles after a 5-second sample. When periodic profiling is enabled, this command finishes the current CPU interval, writes a separate 5-second CPU profile, and starts a new periodic CPU file. Explicitly requested profiles are rotated according to ProfilerMaxFilesPerProfile when the command completes. While this command is collecting profiles, all profiler runtime commands are rejected with "profiler is busy".
periodic_prof_set_interval	Set profiling interval.
version	Display agent version.
userparameter_reload	Reload values of the UserParameter and Include options from the current configuration file.
help	Display help information on runtime control.

Examples:

- Increasing log level for agent 2.
- Print runtime control options.

```
zabbix_agent2 -R log_level_increase
zabbix_agent2 -R help
```

Configuration file The default path to the configuration file depends on the operating system and Zabbix installation method:

- On UNIX, the default is:
  - For installation from sources: /usr/local/etc/zabbix_agent2.conf, or as set by compile-time variables --sysconfdir or --prefix.
  - For installation from packages: /etc/zabbix/zabbix_agent2.conf.
- On Windows, the default is C:\Program Files\Zabbix Agent 2\zabbix_agent2.conf.

If a single instance of the agent is running on a host, it can use the default configuration file or a configuration file specified on the command line. In case of multiple instances, each agent instance must have its own configuration file (one of the instances can use the default configuration file).

The configuration parameters of Zabbix agent 2 mostly match those of Zabbix agent, with some exceptions:

Parameters		Description
Zabbix agent 2-specific configuration parameters		
	ControlSocket	The runtime control socket path. Zabbix agent 2 uses a control socket for runtime commands.
	EnablePersistentBuffer, PersistentBufferFile, PersistentBufferPeriod	These parameters are used to configure persistent storage on Zabbix agent 2 for active items.
	ForceActiveChecksOnStart	Determines whether the agent should perform active checks immediately after restart or spread evenly over time.
	Plugins	Plugins may have their own parameters, in the format Plugins.<Plugin name>.<Parameter>=<value>. A common plugin parameter is System.Capacity, setting the limit of checks that can be executed at the same time.
	StatusPort	The port Zabbix agent 2 will be listening on for HTTP status request and display of a list of configured plugins and some internal parameters.
Zabbix agent-specific configuration parameters not supported on Zabbix agent 2		
	AllowRoot, User	Not supported on Zabbix agent 2 because it does not support daemonization.
	EnableRemoteCommands	Deprecated; the AllowKey, AllowKeyRegexp, DenyKey, and DenyKeyRegexp parameters are used instead in both agents.
	ListenBacklog	Used in Zabbix agent to define the maximum number of pending connections in the TCP queue. Not supported on Zabbix agent 2.
	LoadModule, LoadModulePath	Loadable modules are not supported on Zabbix agent 2.
	MaxLinesPerSecond, LogRemoteCommands	In Zabbix agent 2, Plugins.Log.MaxLinesPerSecond and Plugins.SystemRun.LogRemoteCommands are used instead.

Parameters	Description
StartAgents	This parameter is used in Zabbix agent to increase passive check concurrency or disable them. In Zabbix agent 2, the concurrency is configured at a plugin level and can be limited by a capacity setting. Passive checks are disabled in Zabbix Agent 2 if the Server parameter is not specified.

For more details, see the configuration file options for Zabbix agent2.

Exit codes Zabbix agent 2 can also be compiled with older OpenSSL versions (1.0.1, 1.0.2).

In this case Zabbix provides mutexes for locking in OpenSSL. If a mutex lock or unlock fails then an error message is printed to the standard error stream (STDERR) and Agent 2 exits with return code 2 or 3, respectively.

Agent vs agent 2 comparison

This section describes the differences between the Zabbix agent and the Zabbix agent 2.

Parameter	Zabbix agent	Zabbix agent 2
Programming language	C	Go with some parts in C
Daemonization	yes	by systemd only (yes on Windows)
Supported extensions	Custom loadable modules in C.	Custom plugins in Go.
Requirements		
Supported platforms	Linux, IBM AIX, FreeBSD, NetBSD, OpenBSD, HP-UX, Mac OS X, Solaris: 9, 10, 11, Windows: all desktop and server versions since XP	Linux, Windows: all desktop and server versions, on which a supported Go version can be installed.
Supported crypto libraries	GnuTLS 3.1.18 and newer. OpenSSL 1.0.1, 1.0.2, 1.1.0, 1.1.1, 3.0.x. LibreSSL - tested with versions 2.7.4, 2.8.2 (certain limitations apply, see the Encryption page for details).	Linux: OpenSSL 1.0.1 and later. MS Windows: OpenSSL 1.1.1 or later. The OpenSSL library must have PSK support enabled. LibreSSL is not supported.
Monitoring processes		
Processes	A separate active check process for each server/proxy record.	Single process with automatically created threads. The maximum number of threads is determined by the GOMAXPROCS environment variable.
Metrics	**UNIX**: see a list of supported items. **Windows**: see a list of additional Windows-specific items.	**UNIX**: All metrics supported by Zabbix agent. Additionally, the agent 2 provides Zabbix-native monitoring solution for: Docker, Memcached, MySQL, PostgreSQL, Redis, systemd, and other monitoring targets - see a full list of agent 2 specific items. **Windows**: All metrics supported by Zabbix agent, and also net.tcp.service* checks of HTTPS, LDAP. Additionally, the agent 2 provides Zabbix-native monitoring solution for: PostgreSQL, Redis.
Concurrency	Active checks for single server are executed sequentially.	Checks from different plugins or multiple checks within one plugin can be executed concurrently.
Third-party traps	no	yes
Additional features		
Persistent storage	no	yes
Persistent files for log*[] metrics	yes (only on Unix)	no

Parameter	Zabbix agent	Zabbix agent 2
Log data upload	Can be performed during log gathering to free the buffer.	Log gathering is stopped when the buffer is full, therefore the `BufferSize` parameter must be at least `MaxLinesPerSecond` x 2.
Changes user at runtime	yes (Unix-like systems only)	no (controlled by systemd)
User-configurable ciphersuites	yes	no

See also:

- *Zabbix processes description*: Zabbix agent, Zabbix agent 2
- *Configuration parameters*: Zabbix agent UNIX / Windows, Zabbix agent 2 UNIX / Windows

Built-in plugins

These plugins are built into Zabbix agent 2 and are available out-of-the-box. Click the plugin name to go to the plugin repository with additional information.

Plugin name	Description	Supported item keys	Comments
Agent	Metrics of the Zabbix agent being used.	agent.hostname, agent.ping, agent.version	Supported keys have the same parameters as Zabbix agent keys.
CPU	System CPU monitoring (number of CPUs/CPU cores, discovered CPUs, utilization percentage).	system.cpu.discovery, system.cpu.num, system.cpu.util	Supported keys have the same parameters as Zabbix agent keys.
Docker	Monitoring of Docker containers.	docker.container_info, docker.container_stats, docker.containers, docker.containers.discovery, docker.data_usage, docker.images, docker.images.discovery, docker.info, docker.ping	See also: Configuration parameters
File	File metrics collection.	vfs.file.cksum, vfs.file.contents, vfs.file.exists, vfs.file.md5sum, vfs.file.regexp, vfs.file.regmatch, vfs.file.size, vfs.file.time	Supported keys have the same parameters as Zabbix agent keys.
Kernel	Kernel monitoring.	kernel.maxfiles, kernel.maxproc	Supported keys have the same parameters as Zabbix agent keys.
Log	Log file monitoring.	log, log.count, logrt, logrt.count	Supported keys have the same parameters as Zabbix agent keys. See also: Plugin configuration parameters (Unix/Windows)
Memcached	Memcached server monitoring.	memcached.ping, memcached.stats	
Modbus	Reads Modbus data.	modbus.get	Supported keys have the same parameters as Zabbix agent keys.

Plugin name	Description	Supported item keys	Comments
MQTT	Receives published values of MQTT topics.	mqtt.get	To configure an encrypted connection to the MQTT broker, specify TLS parameters in the agent configuration file as a named session or default parameters. Currently TLS parameters cannot be passed as item key parameters.
MySQL	Monitoring of MySQL and its forks.	mysql.custom.query, mysql.db.discovery, mysql.db.size, mysql.get_status_variables, mysql.ping, mysql.replication.discovery, mysql.replication.get_slave_status, mysql.version	To configure an encrypted connection to the database, specify TLS parameters in the agent configuration file as a named session or default parameters. Currently TLS parameters cannot be passed as item key parameters.
NetIf	Monitoring of network interfaces.	net.if.collisions, net.if.discovery, net.if.get, net.if.in, net.if.out, net.if.total	Supported keys have the same parameters as Zabbix agent keys.
Oracle	Oracle Database monitoring.	oracle.diskgroups.stats, oracle.diskgroups.discovery, oracle.archive.info, oracle.archive.discovery, oracle.cdb.info, oracle.custom.query, oracle.datafiles.stats, oracle.db.discovery, oracle.fra.stats, oracle.instance.info, oracle.pdb.info, oracle.pdb.discovery, oracle.pga.stats, oracle.ping, oracle.proc.stats, oracle.redolog.info, oracle.sga.stats, oracle.sessions.stats, oracle.sys.metrics, oracle.sys.params, oracle.ts.stats, oracle.ts.discovery, oracle.user.info, oracle.version	Install the Oracle Instant Client before using the plugin.
Proc	Process CPU utilization percentage.	proc.cpu.util	Supported key has the same parameters as Zabbix agent key.
Redis	Redis server monitoring.	redis.config, redis.info, redis.ping, redis.slowlog.count	To configure an encrypted connection to Redis, specify TLS parameters in the agent configuration file as a named session or default parameters. TLS parameters cannot be passed as item key parameters. Note that an incorrect or otherwise invalid TLS configuration can prevent Zabbix agent 2 from starting, so verify certificate files, permissions and paths before enabling TLS.

Plugin name	Description	Supported item keys	Comments

Smart	S.M.A.R.T. monitoring.	smart.attribute.discovery, smart.disk.discovery, smart.disk.get	The minimum required smartctl version is 7.1.
			Sudo/root access rights to smartctl are required for the user executing Zabbix agent 2. The plugin uses only the following commands:
			`/usr/sbin/smartctl -a *`
			`/usr/sbin/smartctl --scan *`
			`/usr/sbin/smartctl -j -V`
			
			Supported keys can be used with Zabbix agent 2 only on Linux/Windows, both as a passive and active check.
			See also:
			Configuration parameters

SW	Listing of installed packages.	system.sw.packages, system.sw.packages.get	The supported keys have the same parameters as Zabbix agent key.
Swap	Swap space size in bytes/percentage.	system.swap.size	Supported key has the same parameters as Zabbix agent key.
SystemRun	Runs specified command.	system.run	Supported key has the same parameters as Zabbix agent key. See also: Plugin configuration parameters (Unix/Windows)
Systemd	Monitoring of systemd services.	systemd.unit.discovery, systemd.unit.get, systemd.unit.info	
TCP	TCP connection availability check.	net.tcp.port	Supported key has the same parameters as Zabbix agent key.
UDP	Monitoring of the UDP services availability and performance.	net.udp.service, net.udp.service.perf	Supported keys have the same parameters as Zabbix agent keys.
Uname	Retrieval of information about the system.	system.hostname, system.sw.arch, system.uname	Supported keys have the same parameters as Zabbix agent keys.
Uptime	System uptime metrics collection.	system.uptime	Supported key has the same parameters as Zabbix agent key.
VFSDev	VFS metrics collection.	vfs.dev.discovery, vfs.dev.get, vfs.dev.read, vfs.dev.write	Supported keys have the same parameters as Zabbix agent keys.
WebCertificate	Monitoring of TLS/SSL website certificates.	web.certificate.get	
WebPage	Web page monitoring.	web.page.get, web.page.perf, web.page.regexp	Supported keys have the same parameters as Zabbix agent keys.
ZabbixAsync	Asynchronous metrics collection.	net.tcp.listen, net.udp.listen, sensor, system.boottime, system.cpu.intr, system.cpu.load, system.cpu.switches, system.hw.cpu, system.hw.macaddr, system.localtime, system.sw.os, system.swap.in, system.swap.out, vfs.fs.discovery	Supported keys have the same parameters as Zabbix agent keys.

Plugin name	Description	Supported item keys	Comments
ZabbixStats	Zabbix server/proxy internal metrics or number of delayed items in a queue.	zabbix.stats	Supported keys have the same parameters as Zabbix agent keys.
ZabbixSync	Synchronous metrics collection.	net.dns, net.dns.record, net.tcp.service, net.tcp.service.perf, proc.mem, proc.num, system.hw.chassis, system.hw.devices, system.sw.packages, system.users.num, vfs.dir.count, vfs.dir.size, vfs.fs.get, vfs.fs.inode, vfs.fs.size, vm.memory.size.	Supported keys have the same parameters as Zabbix agent keys.

Loadable plugins

Loadable plugins are not available out-of-the-box for Zabbix agent 2 and need to be installed separately:

- On Linux, you can use packages (e.g., install Ember+ on Ubuntu with `apt install zabbix-agent2-plugin-ember-plus`) or build plugins
- On Windows, you can install from MSI or build plugins

Attention:
Before installing a plugin, please check its README file. It may contain specific requirements and installation instructions.

Click the plugin name to go to the plugin repository, which contains the plugin README file with additional information.

Plugin name	Description	Supported item keys	Comments
Ceph	Ceph monitoring.	ceph.df.details, ceph.osd.stats, ceph.osd.discovery, ceph.osd.dump, ceph.ping, ceph.pool.discovery, ceph.status	See also Ceph plugin configuration parameters.
Ember+	Monitoring of Ember+.	ember.get	See also Ember+ plugin configuration parameters.

Plugin name	Description	Supported item keys	Comments
MongoDB	Monitoring of MongoDB servers and clusters (document-based, distributed database).	mongodb.cfg.discovery, mongodb.collection.stats, mongodb.collections.discovery, mongodb.collections.usage, mongodb.connpool.stats, mongodb.db.stats, mongodb.db.discovery, mongodb.jumbo_chunks.count, mongodb.oplog.stats, mongodb.ping, mongodb.rs.config, mongodb.rs.status, mongodb.server.status, mongodb.sh.discovery, mongodb.version	To configure encrypted connections to the database, specify TLS parameters in the agent configuration file as a named session parameters. Currently TLS parameters cannot be passed as item key parameters. See also MongoDB plugin configuration parameters.
MSSQL	Monitoring of MSSQL database.	mssql.availability.group.get, mssql.custom.query, mssql.db.get, mssql.job.status.get, mssql.last.backup.get, mssql.local.db.get, mssql.mirroring.get, mssql.nonlocal.db.get, mssql.perfcounter.get, mssql.ping, mssql.quorum.get, mssql.quorum.member.get, mssql.replica.get, mssql.version	To configure an encrypted connection to the database, specify TLS parameters in the agent configuration file as a named session or default parameters. Currently TLS parameters cannot be passed as item key parameters. See also MSSQL plugin configuration parameters.
NVIDIA GPU	Monitoring of NVIDIA GPU.	nvml.device.count, nvml.device.decoder.utilization, nvml.device.ecc.mode, nvml.device.encoder.stats.get, nvml.device.encoder.utilization, nvml.device.energy.consumption, nvml.device.errors.memory, nvml.device.errors.register, nvml.device.fan.speed.avg, nvml.device.get, nvml.device.graphics.frequency, nvml.device.memory.bar1.get, nvml.device.memory.fb.get, nvml.device.memory.frequency, nvml.device.pci.utilization, nvml.device.performance.state, nvml.device.power.limit, nvml.device.power.usage, nvml.device.serial, nvml.device.sm.frequency, nvml.device.temperature, nvml.device.utilization, nvml.device.video.frequency, nvml.system.driver.version, nvml.version	See also NVIDIA GPU plugin configuration parameters.

Plugin name	Description	Supported item keys	Comments
PostgreSQL	Monitoring of PostgreSQL and its forks.	pgsql.autovacuum.count, pgsql.archive, pgsql.bgwriter, pgsql.cache.hit, pgsql.connections, pgsql.custom.query, pgsql.dbstat, pgsql.dbstat.sum, pgsql.db.age, pgsql.db.bloating_tables, pgsql.db.discovery, pgsql.db.size, pgsql.locks, pgsql.oldest.xid, pgsql.ping, pgsql.queries, pgsql.replication.count, pgsql.replication.process, pgsql.replication.process.discovery, pgsql.replication.recovery_role, pgsql.replication.status, pgsql.replication_lag.b, pgsql.replication_lag.sec, pgsql.uptime, pgsql.version, pgsql.wal.stat	To configure encrypted connections to the database, specify TLS parameters in the agent configuration file as a named session or default parameters. Currently TLS parameters cannot be passed as item key parameters. See also PostgreSQL plugin configuration parameters.

Note:
Loadable plugins, when launched with:
 - `-V --version` - print plugin version and license information;
 - `-h --help` - print help information.
 - `-t, --test <item key>` - launch plugin for testing (plugin config ignored).

Configuration parameters: Zabbix agent 2

Overview

Configuration parameters allow customizing Zabbix agent 2 behavior.

Zabbix agent 2 is a new generation of Zabbix agent and may be used in place of Zabbix agent.

Specify parameter values in the Zabbix agent 2 configuration file (`zabbix_agent2.conf`). All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Alias	Set an alias for an item key.
AllowKey	Allow the execution of those item keys that match a pattern.
AllowKeyRegexp	Allow the execution of item keys that match a regular expression-based pattern.
BufferSend	Set the frequency of sending values from the buffer to Zabbix server/proxy.
BufferSize	Set the maximum number of values in the memory buffer.
ControlSocket	Specify the control socket for sending runtime commands with the '-R' option.
DebugLevel	Set the debug level.
DenyKey	Deny the execution of those item keys that match a pattern.
DenyKeyRegexp	Deny the execution of item keys that match a regular expression-based pattern.
EnablePersistentBuffer	Enable the usage of local persistent storage for active items.
EnableProfiler	Enable periodic Go profiling on agent startup.
ForceActiveChecksOnStart	Perform active checks immediately after the restart for the first received configuration.
HeartbeatFrequency	Set the frequency of heartbeat messages in seconds.

Parameter	Description
HostInterface	Specify the host interface.
HostInterfaceItem	Specify an item for getting the host interface.
HostMetadata	Specify the host metadata.
HostMetadataItem	Specify a Zabbix agent item for getting the host metadata.
Hostname	Specify the hostname.
HostnameItem	Specify a Zabbix agent item for getting the hostname.
Include	Specify a directory or individual files to be included in the configuration file.
ListenIP	Specify the agent listen IP address for connections from the server/proxy.
ListenPort	Specify the agent listen port.
LogFile	Specify the agent log file.
LogFileSize	Set the maximum log file size.
LogType	Specify the log output type.
PersistentBufferFile	Specify the file for keeping the SQLite database.
PersistentBufferPeriod	Set the time period for storing data when there is no connection to Zabbix server or proxy.
PidFile	Specify the PID file.
Plugins.<PluginName>.System.Capacity	Set the capacity of checks per plugin that can be executed at the same time.
Plugins.Log.MaxLinesPerSecond	Set the maximum number of new lines per second for sending to Zabbix server/proxy when processing `log` or `logrt` active checks.
Plugins.SystemRun.LogRemoteCommands	Enable the logging of the executed shell commands as warnings.
PluginSocket	Specify path to the UNIX socket for loadable plugin communications.
PluginTimeout	Set the timeout for connections to loadable plugins.
ProfilerDir	The directory where profiling data will be written.
ProfilerMaxFilesPerProfile	The maximum number of files to keep for each profile type.
ProfilerInterval	The profiling data write interval.
RefreshActiveChecks	Specify the frequency of refreshing the list of active checks.
Server	Specify the IP addresses or DNS names of Zabbix servers/proxies.
ServerActive	Specify the Zabbix server/proxy address or cluster configuration to get active checks from.
SourceIP	Specify the source IP address.
StatusPort	Specify port for listening to HTTP status requests (`http://localhost:<port>/status`).
Timeout	Specify how long to wait (in seconds) for establishing connection and exchanging data with Zabbix server or proxy.
TLSAccept	Specify the level of encryption required for incoming connections.
TLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.
TLSCertFile	Specify full pathname of the file containing the agent certificate or certificate chain, used for encrypted communications between Zabbix components.
TLSCipherAll	Specify the OpenSSL cipher string (TLS 1.2). Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.
TLSCipherAll13	Specify the OpenSSL cipher string (TLS 1.3) for OpenSSL 1.1.1 or newer. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.
TLSCipherCert	Specify the OpenSSL cipher string (TLS 1.2). Override the default ciphersuite selection criteria for certificate-based encryption.
TLSCipherCert13	Specify the OpenSSL cipher string (TLS 1.3) for OpenSSL 1.1.1 or newer. Override the default ciphersuite selection criteria for certificate-based encryption.
TLSCipherPSK	Specify the OpenSSL cipher string (TLS 1.2). Override the default ciphersuite selection criteria for PSK-based encryption.
TLSCipherPSK13	Specify the OpenSSL cipher string (TLS 1.3) for OpenSSL 1.1.1 or newer. Override the default ciphersuite selection criteria for PSK-based encryption.
TLSConnect	Specify the level of encryption required for connections to Zabbix server or proxy.
TLSCRLFile	Specify full pathname of the file containing revoked certificates. This parameter is used for encrypted communications between Zabbix components.
TLSKeyFile	Specify full pathname of the file containing the agent private key, used for encrypted communications between Zabbix components.
TLSPSKFile	Specify full pathname of the file containing the agent pre-shared key, used for encrypted communications with Zabbix server or proxy.
TLSPSKIdentity	Specify the pre-shared key identity string, used for encrypted communications with Zabbix server or proxy.
TLSServerCertIssuer	Specify the allowed server (proxy) certificate issuer.
TLSServerCertSubject	Specify the allowed server (proxy) certificate subject.
UnsafeUserParameters	Allow all characters to be passed in arguments to user-defined parameters.
UserParameter	Specify a user-defined parameter to monitor.

Parameter	Description
UserParameterDir	Specify the default search path for UserParameter commands.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files.
- Values support environment variables.
- Zabbix supports configuration files only in UTF-8 encoding without BOM.
- Comments starting with # are only supported in the beginning of the line.

Parameter details

Alias

Set an alias for an item key. It can be used to substitute a long and complex item key with a shorter and simpler one.
 Multiple `Alias` parameters may be present. Multiple parameters with the same `Alias` key are not allowed.
 Different `Alias` keys may reference the same item key.
 Aliases can be used in `HostMetadataItem` but not in `HostnameItem` parameter.

Example 1: Retrieving the ID of user `zabbix`.

```
Alias=zabbix.userid:vfs.file.regexp[/etc/passwd,"^zabbix:.:([0-9]+)",,,,\1]
```

Now the `zabbix.userid` shorthand key may be used to retrieve data.

Example 2: Getting CPU utilization with default and custom parameters.

```
Alias=cpu.util:system.cpu.util
Alias=cpu.util[*]:system.cpu.util[*]
```

This allows you to use the `cpu.util` key to get CPU utilization percentage with default parameters as well as use `cpu.util[all, idle, avg15]` to get specific data about CPU utilization.

Example 3: Running multiple low-level discovery rules processing the same discovery items.

```
Alias=vfs.fs.discovery[*]:vfs.fs.discovery
```

Now it is possible to set up several discovery rules using `vfs.fs.discovery` with different parameters for each rule, e.g., `vfs.fs.discovery[foo]`, `vfs.fs.discovery[bar]`, etc.

AllowKey

Allow the execution of those item keys that match a pattern. The key pattern is a wildcard expression that supports the * character to match any number of any characters.
Multiple key matching rules may be defined in combination with `DenyKey`. Multiple key matching rules may be defined in combination with `DenyKey`, `AllowKeyRegexp`, and `DenyKeyRegexp`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority. See also: Restricting agent checks.

AllowKeyRegexp

Allow the execution of those item keys that match a pattern. The key pattern is a regular expression.
 Multiple key matching rules may be defined in combination with `DenyKeyRegexp`, `AllowKey`, and `DenyKey`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority. See also: Restricting agent checks.

Example:

```
AllowKeyRegexp=^system\.run\[(?:/usr)?/bin/hostname(?: -[fsd])?\]$
```

BufferSend

Set the frequency of sending values from the buffer to Zabbix server/proxy (in seconds).
 Note that if the buffer is full, the data will be sent sooner.

Default: 5
 Range: 1-3600

BufferSize

Set the maximum number of values in the memory buffer.
 The agent will send all collected data to the Zabbix server or proxy if the buffer is full. This parameter should only be used if persistent buffer is disabled (`EnablePersistentBuffer=0`).

Default: 1000
 Range: 2-65535

ControlSocket

Specify the control socket for sending runtime commands with the -R option.

Default: /tmp/agent.sock

DebugLevel

Set the debug level: - *0* - basic information about starting and stopping of Zabbix processes; - *1* - critical information; - *2* - error information; - *3* - warnings; - *4* - for debugging (produces lots of information); - *5* - extended debugging (produces even more information).

Default: 3
 Range: 0-5

DenyKey

Deny the execution of those item keys that match a pattern.
 The key pattern is a wildcard expression that supports the * character to match any number of any characters.
Multiple key matching rules may be defined in combination with AllowKey. Multiple key matching rules may be defined in combination with AllowKey, AllowKeyRegexp, and DenyKeyRegexp. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority. See also: Restricting agent checks.

Default: DenyKey=system.run[*]

DenyKeyRegexp

Deny the execution of those item keys that match a pattern. The key pattern is a regular expression.
 Multiple key matching rules may be defined in combination with AllowKeyRegexp, AllowKey, and DenyKey. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority.
 All system.run[] items are denied by default (DenyKey=system.run[*]), unless another system.run[*] rule is specified.

See also: Restricting agent checks.

Example:

```
DenyKeyRegexp=^system\.run\[.*(?:rm -rf|poweroff|reboot).*\]$
```

EnablePersistentBuffer

Enable the usage of local persistent storage for active items.
 If persistent storage is disabled, the memory buffer will be used.

Default: 0
 Values: 0 - disabled, 1 - enabled

EnableProfiler

Enable periodic Go profiling on agent startup. Profiling can also be enabled or disabled at runtime using the -R option. While periodic_prof_execute is collecting profiles, all profiler runtime commands are rejected as busy.

Default: 0
 Values: 0 - disabled, 1 - enabled

ForceActiveChecksOnStart

Perform active checks immediately after the restart for the first received configuration.
 Also available as a per-plugin configuration parameter, for example: Plugins.Uptime.System.ForceActiveChecksOnStart=1

Default: 0
 Values: 0 - disabled, 1 - enabled

HeartbeatFrequency

Set the frequency of heartbeat messages in seconds.
 Used for monitoring the availability of active checks.
0 - heartbeat messages disabled.

Default: 60
 Range: 0-3600

HostInterface

Specify the host interface (IP address or DNS name) used during host autoregistration.
 This value will be used to populate the interface in the newly created host, and allows configuring either IP or DNS address explicitly.
 For more details, see Using DNS as default interface.

If not defined, the value will be acquired from HostInterfaceItem.

The agent will issue an error and not start if the value is over the limit of 255 characters.

Range: 0-255 characters

HostInterfaceItem

Specify an item to determine the host interface (IP address or DNS name) during host autoregistration.
 This value is used only if HostInterface is not defined.
 For more details, see Using DNS as default interface.

During an autoregistration request the agent will log a warning message if the value returned by the specified item is over the limit of 255 characters.

The system.run[] item is supported regardless of AllowKey, AllowKeyRegexp, DenyKey, or DenyKeyRegexp settings.

HostMetadata

Specify the metadata used to identify or distinguish the host during host autoregistration.
 HostMetadata allows distinguishing between hosts beyond the hostname.

If not defined, the value will be acquired from HostMetadataItem.

The agent will issue an error and not start if the specified value is over the limit of 2034 bytes or a non-UTF-8 string. When a parameter expects an IP address or DNS name, values that are valid UTF-8 but are not valid IPs or DNS names will also be rejected and reported as invalid.

Multi-line metadata is not supported—the output will be truncated at the first line break.

Range: 0-2034 bytes

HostMetadataItem

Specify an item to obtain host metadata.
 This option is only used when HostMetadata is not defined.

The HostMetadataItem value is retrieved on each autoregistration attempt and is used only at host autoregistration process. HostMetadataItem allows distinguishing between hosts beyond the hostname.

User parameters and aliases are supported. The system.run[] item is supported regardless of AllowKey, AllowKeyRegexp, DenyKey, or DenyKeyRegexp settings.

During an autoregistration request the agent will log a warning message if the value returned by the specified item is over the limit of 65535 UTF-8 code points. The value returned by the item must be a UTF-8 string; otherwise, it will be ignored. When a parameter expects an IP address or DNS name, values that are valid UTF-8 but are not valid IPs or DNS names will also be rejected and reported as invalid.

Multi-line metadata is not supported—the output will be truncated at the first line break.

Hostname

Specify unique, case-sensitive hostnames. A comma-delimited list.
 Required for active checks and must match hostnames as configured on the server.
 The value is acquired from HostnameItem if undefined.
Allowed characters: alphanumeric, '.', ' ', '_' and '-'. Maximum length: 128 characters per hostname, 2048 characters for the entire line.

Default: Set by HostnameItem

HostnameItem

Specify an item for getting the hostname.
 This option is only used when Hostname is not defined.
 User parameters or aliases are not supported, but the system.run[] item is supported regardless of AllowKey, AllowKeyRegexp, DenyKey, or DenyKeyRegexp values.

Default: system.hostname

Include

Specify a directory or individual files to be included in the configuration file.
 During Zabbix installation, this parameter is set to /usr/local/etc, unless you change this during compile time.

General rules for the Include parameter:

- The path can be relative to the zabbix_agent2.conf file location.
- You can specify multiple Include parameters in one configuration file.
- Zabbix processes the Include parameter as soon as it reads it, so consider placing it at the end of the configuration file.

Rules for included files:

- Each file must be readable.
- Each file must have correct syntax. Otherwise, the component will not start.

Rules for including directories or file patterns:

- Do not assume any order of inclusion (e.g., files are not included in alphabetical order).
- Do not define the same parameter in more than one included file (e.g., to override a general setting with a specific one), since you cannot predict which value will be used.
- Watch out for backup copies that some text editors create automatically (e.g., editing include/my_specific.conf may create include/my_specific.conf.BAK), since both files will be included.

To include only relevant files from a directory, use the wildcard (*) for pattern matching.

Example:

```
Include=/absolute/path/to/config/files/*.conf
```

ListenIP

Specify the agent listen IP address for connections from the server/proxy. A comma-delimited list.
 The first IP address is sent to the Zabbix server, if connecting to it, to retrieve the list of active checks.

Default: 0.0.0.0

ListenPort

Specify the agent listen port for connections from the server/proxy.

Default: 10050
 Range: 1024-32767

LogFile

Specify the agent log file.

Default: /tmp/zabbix_agent2.log
 Mandatory: Yes, if LogType is set to file; otherwise, no

LogFileSize

Set the maximum log file size (in MB).
0 - disable automatic log rotation.
Note: If the log file size limit is reached and file rotation fails, for whatever reason, the existing log file is truncated and started anew.

Default: 1
 Range: 0-1024

LogType

Specify the log output type: - *file* - write log to the file specified by LogFile parameter; - *system* - write log to syslog; - *console* - write log to standard output.

Default: file

PersistentBufferFile

Specify the file for keeping the SQLite database. Must be a full filename.
 This parameter is only used if persistent buffer is enabled (EnablePersistentBuffer=1).

PersistentBufferPeriod

Set the time period for storing data when there is no connection to Zabbix server or proxy. Older data will be lost.
 Log data will be preserved. This parameter is only used if persistent buffer is enabled (EnablePersistentBuffer=1).

Default: 1h
 Range: 1m-365d

PidFile

Specify the PID file.

Default: /tmp/zabbix_agent2.pid

Plugins.<PluginName>.System.Capacity

Set the limit of checks per <PluginName> plugin that can be executed at the same time.

Default: 1000 Range: 1-1000

Plugins.Log.MaxLinesPerSecond

Set the maximum number of new lines per second for sending to Zabbix server/proxy when processing log or logrt active checks.
 The provided value will be overridden by the maxlines parameter, provided in the log and logrt item key.
Note: Zabbix will process 10 times more new lines than set in MaxLinesPerSecond to seek the required string in log items.

Default: 20
 Range: 1-1000

Plugins.SystemRun.LogRemoteCommands

Enable the logging of the executed shell commands as warnings.
 The commands will be logged only if executed remotely.
 Log entries will not be created if system.run[] is launched locally by the HostMetadataItem, HostInterfaceItem or HostnameItem parameters.

Default: 0
 Values: 0 - disabled, 1 - enabled

PluginSocket

Specify path to the UNIX socket for loadable plugin communications.

Default: `/tmp/agent.plugin.sock`

PluginTimeout

Specify timeout for connections to loadable plugins, in seconds.

Default: `Timeout<br>` Range: 1-30

ProfilerDir

The directory where profiling data will be written.

Default: `/tmp/zabbix_agent2_profiles<br>`

ProfilerMaxFilesPerProfile

The maximum number of files to keep for each profile type. The limit is applied to completed files after each profile dump. The active CPU profile is collected in a hidden temporary file and does not count toward the limit. Files created by `periodic_prof_execute` are also subject to the limit.

Default: 10
 Range: 1-100

ProfilerInterval

The profiling data write interval, in seconds.

Default: 3600
 Range: 1-86400

RefreshActiveChecks

Specify the frequency of refreshing the list of active checks, in seconds.
 Note that after failing to refresh active checks the next refresh will be attempted in 60 seconds.

Default: 5
 Range: 1-86400

Server

Specify the IP addresses, optionally in CIDR notation, or DNS names of Zabbix servers/proxies. A comma-delimited list.
 Incoming connections will be accepted only from the hosts listed here. If IPv6 support is enabled then `127.0.0.1`, `::127.0.0.1`, `::ffff:127.0.0.1` are treated equally and `::/0` will allow any IPv4 or IPv6 address. `0.0.0.0/0` can be used to allow any IPv4 address. Spaces are allowed. If this parameter is not specified, disables passive checks and the agent will not listen on any TCP port.

Example:

```
Server=127.0.0.1,192.168.1.0/24,::1,2001:db8::/32,zabbix.example.com
```

ServerActive

Specify the Zabbix server/proxy address or cluster configuration to get active checks from.
 The server/proxy address is an IP address or DNS name and optional port separated by colon.
The cluster configuration is one or more server or proxy group member addresses separated by semicolon. Multiple Zabbix servers/clusters and Zabbix proxies can be specified, separated by comma. Unless using proxy groups, more than one Zabbix proxy should not be specified from each Zabbix server/cluster. If a Zabbix proxy is specified then Zabbix server/cluster for that proxy should not be specified.
Multiple comma-delimited addresses can be provided to use several independent Zabbix servers in parallel. Spaces are allowed.
If the port is not specified, default port is used.
IPv6 addresses must be enclosed in square brackets if port for that host is specified. If port is not specified, square brackets for IPv6 addresses are optional.
If this parameter is not specified, active checks are disabled.

Example for Zabbix proxy:

```
ServerActive=127.0.0.1:10051
```

Example for Zabbix proxy group:

```
ServerActive=proxy1.example.com;proxy2.example.com;proxy3.example.com;proxy4.example.com;proxy5.example.co
```

Example for multiple servers:

```
ServerActive=127.0.0.1:20051,zabbix.domain,[::1]:30051,::1,[12fc::1]
```

Example for high availability:

```
ServerActive=zabbix.cluster.node1;zabbix.cluster.node2:20051;zabbix.cluster.node3
```

Example for high availability with two clusters and one server:

```
ServerActive=zabbix.cluster.node1;zabbix.cluster.node2:20051,zabbix.cluster2.node1;zabbix.cluster2.node2,z
```

SourceIP

Specify the source IP address for:

- outgoing connections to Zabbix server or Zabbix proxy.
- making connections while executing some items (`web.page.get`, `net.tcp.port`, etc.).

StatusPort

Specify port for listening to HTTP status requests (`http://localhost:<port>/status`).

Range: 1024-32767

Timeout

Specify how long to wait (in seconds) for establishing connection and exchanging data with Zabbix proxy or server.

This parameter defines the duration of various communication operations, including:

- awaiting a response from the Zabbix server;
- sending requests to Zabbix server, including item configuration requests and item data in active checks;
- retrieving log data through logfile;
- sending heartbeat messages;
- maximum duration for `vfs.*` checks;
- serving as a fallback in scenarios where a server or proxy older than version 7.0 sends checks without timeouts.

This timeout will **not** be used for those agent checks that have configurable timeout settings in the frontend (on global, proxy, or per-item level).

Default: 3
 Range: 1-30

TLSAccept

Specify the level of encryption required for incoming connections.
 Used for passive checks.
 Multiple values can be specified, separated by comma: - *unencrypted* - accept connections without encryption (default); - *psk* - accept connections with TLS and a pre-shared key (PSK); - *cert* - accept connections with TLS and a certificate.

Mandatory: yes, if TLS certificate or PSK parameters are defined (even for *unencrypted* connection); otherwise, no

TLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.

TLSCertFile

Specify full pathname of the file containing the agent certificate or certificate chain, used for encrypted communications with Zabbix components.

TLSCipherAll

Specify the OpenSSL cipher string (TLS 1.2). Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.

Example:

```
EECDH+aRSA+AES128:RSA+aRSA+AES128:kECDHEPSK+AES128:kPSK+AES128
```

TLSCipherAll13

Specify the OpenSSL cipher string (TLS 1.3) for OpenSSL 1.1.1 or newer. Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.

Example:

```
TLS_AES_256_GCM_SHA384:TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256
```

TLSCipherCert

Specify the OpenSSL cipher string (TLS 1.2). Override the default ciphersuite selection criteria for certificate-based encryption.

Example:

```
EECDH+aRSA+AES128:RSA+aRSA+AES128
```

Note that this parameter cannot be used together with `TLSAccept=cert,psk`; for certificate connection (`TLSConnect=cert`) use `TLSCipherAll` instead.

TLSCipherCert13

Specify the OpenSSL cipher string (TLS 1.3) for OpenSSL 1.1.1 or newer. Override the default ciphersuite selection criteria for certificate-based encryption.

Note that this parameter cannot be used together with `TLSAccept=cert,psk`; for certificate connection (`TLSConnect=cert`) use `TLSCipherAll13` instead.

TLSCipherPSK

Specify the OpenSSL cipher string (TLS 1.2). Override the default ciphersuite selection criteria for PSK-based encryption.

Example:

```
kECDHEPSK+AES128:kPSK+AES128
```

TLSCipherPSK13

Specify the OpenSSL cipher string (TLS 1.3) for OpenSSL 1.1.1 or newer. Override the default ciphersuite selection criteria for PSK-based encryption.

Example:

```
TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256
```

TLSConnect

Specify the level of encryption required for agent connections to Zabbix server or proxy.
 Used for active checks.
 Only one value can be specified: - *unencrypted* - connect without encryption (default); - *psk* - connect using TLS and a pre-shared key (PSK); - *cert* - connect using TLS and a certificate.

Mandatory: yes, if TLS certificate or PSK parameters are defined (even for *unencrypted* connection); otherwise, no

TLSCRLFile

Specify full pathname of the file containing revoked certificates. This parameter is used for encrypted communications between Zabbix components.

TLSKeyFile

Specify full pathname of the file containing the agent private key, used for encrypted communications between Zabbix components.

TLSPSKFile

Specify full pathname of the file containing the agent pre-shared key, used for encrypted communications with Zabbix server or proxy.

TLSPSKIdentity

Specify the pre-shared key identity string, used for encrypted communications with Zabbix server or proxy.

TLSServerCertIssuer

Specify the allowed server (proxy) certificate issuer.

TLSServerCertSubject

Specify the allowed server (proxy) certificate subject.

UnsafeUserParameters

Allow all characters to be passed in arguments to user-defined parameters. The following characters are not allowed: \ ' " ` * ? [ ] { } ~ $! & ; ( ) < > | # @ %
Additionally, newline characters are not allowed.

Default: 0
 Values: 0 - do not allow, 1 - allow

UserParameter

Specify a user-defined parameter to monitor. There can be several user-defined parameters.
Format: `UserParameter=<key>,<shell command>`
Note that the shell command must not return empty string or EOL only. Shell commands may have relative paths, if the `UserParameterDir` parameter is specified.

Example:

```
UserParameter=system.test,who|wc -l
UserParameter=check_cpu,./custom_script.sh
```

UserParameterDir

Specify the default search path for `UserParameter` commands. If used, the agent will change its working directory to the one specified here before executing a command. Thereby, UserParameter commands can have a relative `./` prefix instead of a full path.
Only one entry is allowed.

Example:

```
UserParameterDir=/opt/myscripts
```

Configuration parameters: Zabbix agent 2 (Windows)

Overview

Configuration parameters allow customizing Zabbix agent 2 behavior on Windows.

Specify parameter values in the Zabbix agent 2 configuration file (`zabbix_agent2.conf`). All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Alias	Set an alias for an item key.
AllowKey	Allow the execution of those item keys that match a pattern.
AllowKeyRegexp	Allow the execution of item keys that match a regular expression-based pattern.
BufferSend	Set the frequency of sending values from the buffer to Zabbix server.
BufferSize	Set the maximum number of values in the memory buffer.
ControlSocket	Specify the control socket, used to send runtime commands with the `-R` option.
DebugLevel	Set the debug level.
DenyKey	Deny the execution of those item keys that match a pattern.
DenyKeyRegexp	Deny the execution of item keys that match a regular expression-based pattern.
EnablePersistentBuffer	Enable the usage of local persistent storage for active items.
EnableProfiler	Enable periodic Go profiling on agent startup.
ForceActiveChecksOnStart	Perform active checks immediately after the restart for the first received configuration.
HeartbeatFrequency	Set the frequency of heartbeat messages in seconds.
HostInterface	Specify the host interface.
HostInterfaceItem	Specify an item for getting the host interface.
HostMetadata	Specify the host metadata.
HostMetadataItem	Specify a Zabbix agent item used for getting the host metadata.
Hostname	Specify the hostname.
HostnameItem	Specify a Zabbix agent item for getting the hostname.
Include	Specify a directory or individual files to be included in the configuration file.
ListenIP	Specify the agent listen IP address for connections from the server/proxy.
ListenPort	Specify the agent listen port.
LogFile	Specify the agent log file.
LogFileSize	Set the maximum log file size.
LogType	Specify the log output type.
PerfCounter	Specify a new parameter <parameter_name> which is the average value for system performance counter <perf_counter_path> for the specified time period <period> (in seconds).
PerfCounterEn	Specify a new parameter <parameter_name> which is the average value for system performance counter <perf_counter_path> for the specified time period <period> (in seconds). Compared to PerfCounter, the perfcounter paths must be in English.
PersistentBufferFile	Specify the file where Zabbix agent 2 should keep the SQLite database.
PersistentBufferPeriod	Set the duration of storing data when there is no connection to the server or proxy.
Plugins.<PluginName>.System.Capacity	Set the limit for checks per plugin that can be executed at the same time.
Plugins.Log.MaxLinesPerSecond	Set the maximum number of new lines the agent will send per second to Zabbix server or proxy when processing `log`, `logrt`, and `eventlog` active checks.
Plugins.SystemRun.LogRemoteCommands	Enable logging of the executed shell commands as warnings.
PluginSocket	Specify the path to the Windows named pipe for loadable plugin communications.
PluginTimeout	Set the timeout for connections with loadable plugins, in seconds.
ProfilerDir	The directory where profiling data will be written.
ProfilerMaxFilesPerProfile	The maximum number of files to keep for each profile type.

Parameter	Description
ProfilerInterval	The profiling data write interval.
RefreshActiveChecks	Specify the frequency of refreshing the list of active checks.
Server	Specify the IP addresses or DNS names of Zabbix servers/proxies.
ServerActive	Specify the Zabbix server/proxy address or cluster configuration to get active checks from.
SourceIP	Specify the source IP address.
StatusPort	Set the port that the agent listens on for HTTP status requests (`http://localhost:<port>/status`).
Timeout	Specify how long to wait (in seconds) for establishing connection and exchanging data with Zabbix server or proxy.
TLSAccept	Specify the level of encryption required for incoming connections.
TLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.
TLSCertFile	Specify full pathname of the file containing the agent certificate or certificate chain, used for encrypted communications between Zabbix components.
TLSCipherAll13	Specify the OpenSSL cipher string (TLS 1.3). Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.
TLSCipherCert13	Specify the OpenSSL cipher string (TLS 1.3). Override the default ciphersuite selection criteria for certificate-based encryption.
TLSCipherPSK13	Specify the OpenSSL cipher string (TLS 1.3). Override the default ciphersuite selection criteria for PSK-based encryption.
TLSConnect	Specify the level of encryption required for connections to Zabbix server or proxy.
TLSCRLFile	Specify full pathname of the file containing revoked certificates. This parameter is used for encrypted communications between Zabbix components.
TLSKeyFile	Specify full pathname of the file containing the agent private key, used for encrypted communications between Zabbix components.
TLSPSKFile	Specify full pathname of the file containing the agent pre-shared key, used for encrypted communications with Zabbix server or proxy.
TLSPSKIdentity	Specify the pre-shared key identity string, used for encrypted communications with Zabbix server or proxy.
TLSServerCertIssuer	Specify the allowed server (proxy) certificate issuer.
TLSServerCertSubject	Specify the allowed server (proxy) certificate subject.
UnsafeUserParameters	Allow all characters to be passed in arguments to user-defined parameters.
UserParameter	Specify a user-defined parameter to monitor.
UserParameterDir	Specify the default search path for `UserParameter` commands.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files;
- Values support environment variables;
- Zabbix supports configuration files only in UTF-8 encoding without BOM;
- Comments starting with # are only supported in the beginning of the line.

Parameter details

Alias

Set an alias for an item key. It can be used to substitute a long and complex item key with a shorter and simpler one.
 Multiple `Alias` parameters may be present. Multiple parameters with the same `Alias` key are not allowed.
 Different `Alias` keys may reference the same item key.
 Aliases can be used in `HostMetadataItem` but not in the `HostnameItem` parameter.

Example 1: Retrieving the paging file usage in percentage from the server.

```
Alias=pg_usage:perf_counter[\Paging File(_Total)\% Usage]
```

Now the shorthand key `pg_usage` may be used to retrieve data.

Example 2: Getting the CPU load with default and custom parameters.

```
Alias=cpu.load:system.cpu.load
Alias=cpu.load[*]:system.cpu.load[*]
```

This allows you to use `cpu.load` key to get the CPU load with default parameters as well as use `cpu.load[percpu,avg15]` to get specific data about the CPU load.

Example 3: Running multiple low-level discovery rules processing the same discovery items.

```
Alias=vfs.fs.discovery[*]:vfs.fs.discovery
```

Now it is possible to set up several discovery rules using `vfs.fs.discovery` with different parameters for each rule, e.g., `vfs.fs.discovery[foo]`, `vfs.fs.discovery[bar]`, etc.

AllowKey

Allow the execution of those item keys that match a pattern. The key pattern is a wildcard expression that supports the * character to match any number of any characters.
 Multiple key matching rules may be defined in combination with `DenyKey`, `AllowKeyRegexp`, and `DenyKeyRegexp`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority. See also: Restricting agent checks.

AllowKeyRegexp

Allow the execution of those item keys that match a pattern. The key pattern is a regular expression.
 Multiple key matching rules may be defined in combination with `DenyKeyRegexp`, `AllowKey`, and `DenyKey`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority. See also: Restricting agent checks.

Example:

```
AllowKeyRegexp=^system\.run\[(?:/usr)?/bin/hostname(?: -[fsd])?\]$
```

BufferSend

Set the frequency of sending values from the buffer to Zabbix server/proxy (in seconds).
 Note that if the buffer is full, the data will be sent sooner.

Default: 5
 Range: 1-3600

BufferSize

Set the maximum number of values in the memory buffer.
 The agent will send all collected data to the Zabbix server or proxy if the buffer is full.
 This parameter should only be used if persistent buffer is disabled (`EnablePersistentBuffer=0`).

Default: 1000
 Range: 2-65535

ControlSocket

Specify the control socket, used to send runtime commands with the `-R` option.

Default: `\\.\pipe\agent.sock`

DebugLevel

Set the debug level: - *0* - basic information about starting and stopping of Zabbix processes; - *1* - critical information; - *2* - error information; - *3* - warnings; - *4* - for debugging (produces lots of information); - *5* - extended debugging (produces even more information).

Default: 3
 Range: 0-5

DenyKey

Deny the execution of those item keys that match a pattern. The key pattern is a wildcard expression that supports the * character to match any number of any characters.
 Multiple key matching rules may be defined in combination with `AllowKey`, `AllowKeyRegexp`, and `DenyKeyRegexp`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority.

See also: Restricting agent checks.

Default: `DenyKey=system.run[*]`

DenyKeyRegexp

Deny the execution of those item keys that match a pattern. The key pattern is a regular expression.
 Multiple key matching rules may be defined in combination with `AllowKeyRegexp`, `AllowKey`, and `DenyKey`. The parameters are processed one by one according to their appearance order. The order determines which one matching rule will have priority.
 All `system.run[]` items are denied by default (`DenyKey=system.run[*]`), unless another `system.run[*]` rule is specified.

See also: Restricting agent checks.

EnablePersistentBuffer

Enable the usage of local persistent storage for active items. If persistent storage is disabled, the memory buffer will be used.

Default: 0
 Values: 0 - disabled, 1 - enabled

EnableProfiler

Enable periodic Go profiling on agent startup. Profiling can also be enabled or disabled at runtime using the -R option. While `periodic_prof_execute` is collecting profiles, all profiler runtime commands are rejected as busy.

Default: 0
 Values: 0 - disabled, 1 - enabled

ForceActiveChecksOnStart

Perform active checks immediately after the restart for the first received configuration. Also available as a per-plugin configuration parameter, for example: `Plugins.Uptime.System.ForceActiveChecksOnStart=1`

Default: 0
 Values: 0 - disabled, 1 - enabled

HeartbeatFrequency

Set the frequency of heartbeat messages in seconds. Used for monitoring the availability of active checks.
 0 - heartbeat messages disabled.

Default: 60
 Range: 0-3600

HostInterface

Specify the host interface (IP address or DNS name) used during host autoregistration. This value will be used to populate the interface in the newly created host, and allows configuring either IP or DNS address explicitly. For more details, see Using DNS as default interface.

If not defined, the value will be acquired from HostInterfaceItem.

The agent will issue an error and not start if the value is over the limit of 255 characters.

Range: 0-255 characters

HostInterfaceItem

Specify an item to determine the host interface (IP address or DNS name) during host autoregistration. This value is used only if `HostInterface` is not defined. For more details, see Using DNS as default interface.

During an autoregistration request the agent will log a warning message if the value returned by the specified item is over the limit of 255 characters.

The `system.run[]` item is supported regardless of `AllowKey`, `AllowKeyRegexp`, `DenyKey`, or `DenyKeyRegexp` settings.

HostMetadata

Specify the metadata used to identify or distinguish the host during host autoregistration (active agent). `HostMetadata` allows distinguishing between hosts beyond the hostname.

If not defined, the value will be acquired from `HostMetadataItem`.

The agent will issue an error and not start if the specified value is over the limit of 2034 bytes or a non-UTF-8 string. When a parameter expects an IP address or DNS name, values that are valid UTF-8 but are not valid IPs or DNS names will also be rejected and reported as invalid.

Multi-line metadata is not supported—the output will be truncated at the first line break.

Range: 0-2034 bytes

HostMetadataItem

Specify a Zabbix agent item to obtain host metadata. This option is only used when `HostMetadata` is not defined.

The `HostMetadataItem` value is retrieved on each autoregistration attempt and is used only at host autoregistration process (active agent). `HostMetadataItem` allows distinguishing between hosts beyond the hostname.

User parameters and aliases are supported. The `system.run[]` item is supported regardless of `AllowKey`, `AllowKeyRegexp`, `DenyKey`, or `DenyKeyRegexp` settings.

During an autoregistration request the agent will log a warning message if the value returned by the specified item is over the limit of 65535 UTF-8 code points. The value returned by the item must be a UTF-8 string; otherwise, it will be ignored. When a parameter expects an IP address or DNS name, values that are valid UTF-8 but are not valid IPs or DNS names will also be rejected and reported as invalid.

Multi-line metadata is not supported—the output will be truncated at the first line break.

Hostname

Specify unique, case-sensitive hostnames. A comma-delimited list.
 Required for active checks and must match hostnames as configured on the server. The value is acquired from HostnameItem if undefined.
 Allowed characters: alphanumeric, '.', ' ', '_' and '-'. Maximum length: 128 characters per hostname, 2048 characters for the entire line.

Default: Set by HostnameItem

HostnameItem

Specify an item for getting the hostname. This option is only used when `Hostname` is not defined. User parameters or aliases are not supported, but the `system.run[]` item is supported regardless of `AllowKey`, `AllowKeyRegexp`, `DenyKey`, or `DenyKeyRegexp` values.

Default: `system.hostname`

Include

Specify a directory or individual files to be included in the configuration file (located in `C:\Program Files\Zabbix Agent 2` by default if Zabbix agent is installed using Windows MSI installer packages; located in the folder specified during installation if Zabbix agent is installed as a zip archive). During Zabbix installation, this parameter is set to `/usr/local/etc`, unless you change this during compile time.

General rules for the `Include` parameter:

- The path can be relative to the `zabbix_agent2.conf` file location (e.g., `Include=.\zabbix_agent2.d\plugins.d\*.conf`).
- You can specify multiple `Include` parameters in one configuration file.
- Zabbix processes the `Include` parameter as soon as it reads it, so consider placing it at the end of the configuration file.

Rules for included files:

- Each file must be readable.
- Each file must have correct syntax. Otherwise, the component will not start.

Rules for including directories or file patterns:

- Do not assume any order of inclusion (e.g., files are not included in alphabetical order).
- Do not define the same parameter in more than one included file (e.g., to override a general setting with a specific one), since you cannot predict which value will be used.
- Watch out for backup copies that some text editors create automatically (e.g., editing `include/my_specific.conf` may create `include/my_specific.conf.BAK`), since both files will be included.

To include only relevant files from a directory, use the wildcard (*) for pattern matching.

Example:

```
Include=C:\Program Files\Zabbix Agent2\zabbix_agent2.d\*.conf
```

ListenIP

Specify the agent listen IP address for connections from the server/proxy. A comma-delimited list.
 The first IP address is sent to the Zabbix server, if connecting to it, to retrieve the list of active checks.

Default: `0.0.0.0`

ListenPort

Specify the agent listen port for connections from the server/proxy.

Default: 10050
 Range: 1024-32767

LogFile

Specify the agent log file.

Default: `c:\zabbix_agent2.log`
 Mandatory: Yes, if `LogType` is set to `file`; otherwise, no.

LogFileSize

Set the maximum log file size (in MB).
 0 - disable automatic log rotation.
 Note: If the log file size limit is reached and file rotation fails, for whatever reason, the existing log file is truncated and started anew.

Default: 1
 Range: 0-1024

LogType

Specify the log output type: - *file* - write log to the file specified by `LogFile` parameter; - *console* - write log to standard output.

Default: `file`

PerfCounter

Define a new parameter <parameter_name> which is the average value for system performance counter <perf_counter_path> for the specified time period <period> (in seconds).
 Syntax: <parameter_name>,"<perf_counter_path>",<period>

For example, if you wish to receive the average number of processor interrupts per second for the last minute, you can define a new parameter "interrupts" as the following:

```
PerfCounter = interrupts,"\Processor(0)\Interrupts/sec",60
```

Please note the double quotes around the performance counter path. The parameter name (interrupts) is to be used as the item key when creating an item. Samples for calculating the average value will be taken every second.
 You may run `typeperf -qx` to get the list of all performance counters available in Windows.

PerfCounterEn

Define a new parameter <parameter_name> which is the average value for system performance counter <perf_counter_path> for the specified time period <period> (in seconds). Compared to PerfCounter, the perfcounter paths must be in English. Supported only on **Windows Server 2008/Vista** and later.
 Syntax: <parameter_name>,"<perf_counter_path>",<period>

For example, if you wish to receive the average number of processor interrupts per second for the last minute, you can define a new parameter "interrupts" as the following:

```
PerfCounterEn = interrupts,"\Processor(0)\Interrupts/sec",60
```

Please note the double quotes around the performance counter path. The parameter name (interrupts) is to be used as the item key when creating an item. Samples for calculating the average value will be taken every second.
 You can find the list of English strings by viewing the following registry key: `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib\009`.

PersistentBufferFile

Specify the file where Zabbix agent 2 should keep the SQLite database. Must be a full filename.
 This parameter is only used if persistent buffer is enabled (`EnablePersistentBuffer=1`).

PersistentBufferPeriod

Set the duration of storing data when there is no connection to the server or proxy. Older data will be lost. Log data will be preserved.
 This parameter is only used if persistent buffer is enabled (`EnablePersistentBuffer=1`).

Default: `1h`
 Range: 1m-365d

Plugins.<PluginName>.System.Capacity

Set the limit of checks per `<PluginName>` plugin that can be executed at the same time.

Default: 1000 Range: 1-1000

Plugins.Log.MaxLinesPerSecond

Set the maximum number of new lines the agent will send per second to Zabbix server or proxy when processing `log`, `logrt`, and `eventlog` active checks. The provided value will be overridden by the `maxlines` parameter, provided in the `log`, `logrt` or `eventlog` item key.
 Note: Zabbix will process 10 times more new lines than set in `MaxLinesPerSecond` to seek the required string in log items.

Default: 20
 Range: 1-1000

Plugins.SystemRun.LogRemoteCommands

Enable the logging of the executed shell commands as warnings. The commands will be logged only if executed remotely. Log entries will not be created if `system.run[]` is launched locally by the `HostMetadataItem`, `HostInterfaceItem` or `HostnameItem` parameters.

Default: 0
 Values: 0 - disabled, 1 - enabled

PluginSocket

Specify the path to the Windows named pipe for loadable plugin communications.

Default: `\\.\pipe\agent.plugin.sock`

PluginTimeout

Set the timeout for connections with loadable plugins, in seconds.

Default: `Timeout`
 Range: 1-30

ProfilerDir

The directory where profiling data will be written.

Default: `c:\zabbix_pprof<br>`

ProfilerMaxFilesPerProfile

The maximum number of files to keep for each profile type. The limit is applied to completed files after each profile dump. The active CPU profile is collected in a hidden temporary file and does not count toward the limit. Files created by `periodic_prof_execute` are also subject to the limit.

Default: 10
 Range: 1-100

ProfilerInterval

The profiling data write interval, in seconds.

Default: 3600
 Range: 1-86400

RefreshActiveChecks

Specify the frequency of refreshing the list of active checks, in seconds.
 Note that after failing to refresh active checks the next refresh will be attempted in 60 seconds.

Default: 5
 Range: 1-86400

Server

Specify the IP addresses, optionally in CIDR notation, or DNS names of Zabbix servers/proxies. A comma-delimited list.
 If the parameter is not specified, passive checks are disabled and the agent will not listen on any TCP port.
 Incoming connections will be accepted only from the hosts listed here. If IPv6 support is enabled then `127.0.0.1`, `::127.0.0.1`, `::ffff:127.0.0.1` are treated equally and `::/0` will allow any IPv4 or IPv6 address. `0.0.0.0/0` can be used to allow any IPv4 address. Spaces are allowed.

Example:

```
Server=127.0.0.1,192.168.1.0/24,::1,2001:db8::/32,zabbix.example.com
```

ServerActive

Specify the Zabbix server/proxy address or cluster configuration to get active checks from.
 If the parameter is not specified, active checks are disabled.

The server/proxy address is an IP address or DNS name and optional port separated by colon. If the port is not specified, default port is used.
 IPv6 addresses must be enclosed in square brackets if port for that host is specified. If port is not specified, square brackets for IPv6 addresses are optional.

The cluster configuration is one or more server or proxy group member addresses separated by semicolon.

Multiple Zabbix servers/clusters and Zabbix proxies can be specified, separated by comma.

Unless using proxy groups, more than one Zabbix proxy should not be specified from each Zabbix server/cluster. If a Zabbix proxy is specified then Zabbix server/cluster for that proxy should not be specified.

Multiple comma-delimited addresses can be provided to use several independent Zabbix servers in parallel. Spaces are allowed.

Example for Zabbix proxy:

```
ServerActive=127.0.0.1:10051
```

Example for Zabbix proxy group:

```
ServerActive=proxy1.example.com;proxy2.example.com;proxy3.example.com;proxy4.example.com;proxy5.example.co
```

Example for multiple servers:

```
ServerActive=127.0.0.1:20051,zabbix.domain,\[::1\]:30051,::1,\[12fc::1\]
```

Example for high availability:

```
ServerActive=zabbix.cluster.node1;zabbix.cluster.node2:20051;zabbix.cluster.node3
```

Example for high availability with two clusters and one server:

```
ServerActive=zabbix.cluster.node1;zabbix.cluster.node2:20051,zabbix.cluster2.node1;zabbix.cluster2.node2,z
```

SourceIP

Specify the source IP address for:

- outgoing connections to Zabbix server or Zabbix proxy;
- making connections while executing some items (`web.page.get`, `net.tcp.port`, etc.).

StatusPort

Set the port that the agent listens on for HTTP status requests (`http://localhost:<port>/status`).

Range: 1024-32767

Timeout

Specify how long to wait (in seconds) for establishing connection and exchanging data with Zabbix proxy or server.

This parameter defines the duration of various communication operations, including:

- awaiting a response from the Zabbix server;
- sending requests to Zabbix server, including item configuration requests and item data in active checks;
- retrieving log data through logfile or Windows event log monitoring;
- sending heartbeat messages;
- maximum duration for `vfs.*` checks;
- serving as a fallback in scenarios where a server or proxy older than version 7.0 sends checks without timeouts.

This timeout will **not** be used for those agent checks that have configurable timeout settings in the frontend (on global, proxy, or per-item level).

Default: 3
 Range: 1-30

TLSAccept

Specify the level of encryption required for incoming connections.
 Used for passive checks.

Multiple values can be specified, separated by comma:

- *unencrypted* - accept connections without encryption (default);
- *psk* - accept connections with TLS and a pre-shared key (PSK);
- *cert* - accept connections with TLS and a certificate.

Mandatory: yes, if TLS certificate or PSK parameters are defined (even for *unencrypted* connection); otherwise, no

TLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.

TLSCertFile

Specify full pathname of the file containing the agent certificate or certificate chain, used for encrypted communications with Zabbix components.

TLSCipherAll13

Specify the OpenSSL cipher string (TLS 1.3). Override the default ciphersuite selection criteria for certificate- and PSK-based encryption.

Example:

```
TLS_AES_256_GCM_SHA384:TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256
```

TLSCipherCert13

Specify the OpenSSL cipher string (TLS 1.3). Override the default ciphersuite selection criteria for certificate-based encryption.

Note that this parameter cannot be used together with `TLSAccept=cert,psk`; for certificate connection (`TLSConnect=cert`) use `TLSCipherAll13` instead.

TLSCipherPSK13

Specify the OpenSSL cipher string (TLS 1.3). Override the default ciphersuite selection criteria for PSK-based encryption.

Example:

```
TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256
```

TLSConnect

Specify the level of encryption required for agent connections to Zabbix server or proxy.
 Used for active checks.
 Only one value can be specified: - *unencrypted* - connect without encryption (default); - *psk* - connect using TLS and a pre-shared key (PSK); - *cert* - connect using TLS and a certificate.

Mandatory: yes, if TLS certificate or PSK parameters are defined (even for *unencrypted* connection); otherwise, no.

TLSCRLFile

Specify full pathname of the file containing revoked certificates. This parameter is used for encrypted communications between Zabbix components.

TLSKeyFile

Specify full pathname of the file containing the agent private key, used for encrypted communications between Zabbix components.

TLSPSKFile

Specify full pathname of the file containing the agent pre-shared key, used for encrypted communications with Zabbix server or proxy.

TLSPSKIdentity

Specify the pre-shared key identity string, used for encrypted communications with Zabbix server or proxy.

TLSServerCertIssuer

Specify the allowed server (proxy) certificate issuer.

TLSServerCertSubject

Specify the allowed server (proxy) certificate subject.

UnsafeUserParameters

Allow all characters to be passed in arguments to user-defined parameters. The following characters are not allowed: \ ' " ` * ? [ ] { } ~ $! & ; ( ) < > | # @ %
 Additionally, newline characters are not allowed.

Default: 0
 Values: 0 - do not allow, 1 - allow

UserParameter

Specify a user-defined parameter to monitor. There can be several user-defined parameters.
 Format: `UserParameter=<key>,<shell command>`
 Note that the shell command must not return empty string or EOL only. Shell commands may have relative paths, if the `UserParameterDir` parameter is specified.

Example:

```
UserParameter=system.test,dir /b | find /c /v ""
UserParameter=check_cpu,.\custom\_script.bat
```

UserParameterDir

Specify the default search path for `UserParameter` commands. If used, the agent will change its working directory to the one specified here before executing a command. Thereby, UserParameter commands can have a relative `.\` prefix instead of a full path.
Only one entry is allowed.

Example:

```
UserParameterDir=C:\opt\myscripts
```

Configuration parameters: Plugins

See configuration parameters for Zabbix agent 2 plugins:

- Ceph plugin
- Docker plugin
- Ember+ plugin
- Memcached plugin
- Modbus plugin
- MongoDB plugin
- MQTT plugin
- MSSQL plugin
- MySQL plugin

- NVIDIA GPU plugin
- Oracle plugin
- PostgreSQL plugin
- Redis plugin
- S.M.A.R.T. plugin

Plugin configuration principles

This page describes plugin configuration principles and best practices.

All plugins are configured using `Plugins.*` parameter, which can either be part of the Zabbix agent 2 configuration file or a plugin's configuration file. If a plugin uses a separate configuration file, path to this file should be specified in the `Include` parameter of Zabbix agent 2 configuration file.

A typical plugin parameter has the following structure:

`Plugins.<PluginName>.<Parameter>=<Value>`

Additionally, there are two specific groups of parameters:

- `Plugins.<PluginName>.Default.<Parameter>=<Value>` used for defining default parameter values.
- `Plugins.<PluginName>.<SessionName>.<Parameter>=<Value>` used for defining separate sets of parameters for different monitoring targets via named sessions.

All parameter names should adhere to the following requirements:

- It is recommended to capitalize the names of your plugins.
- The parameter should be capitalized.
- Special characters are not allowed.
- Nesting isn't limited by a maximum level.
- The number of parameters is not limited.

For example, to perform active checks that do not have *Scheduling* update interval immediately after the agent restart only for the Uptime plugin, set `Plugins.Uptime.System.ForceActiveChecksOnStart=1` in the configuration file. Similarly, to set custom limit for concurrent checks for the CPU plugin, set the `Plugins.CPU.System.Capacity=N` in the configuration file.

Default values

You can set default values for the connection-related parameters (URI, username, password, etc.) in the configuration file in the format:

```
Plugins.<PluginName>.Default.<Parameter>=<Value>
```

For example, `Plugins.Mysql.Default.Username=zabbix`, `Plugins.MongoDB.Default.Uri=tcp://127.0.0.1:27017`, etc.

If a value for such parameter is not provided in an item key or in the named session parameters, the plugin will use the default value. If a default parameter is also undefined, hardcoded defaults will be used.

> **Note:**
> If an item key does not have any parameters, Zabbix agent 2 will attempt to collect the metric using values defined in the default parameters section.

Named sessions

Named sessions represent an additional level of plugin parameters and can be used to specify separate sets of authentication parameters for each of the instances being monitored. Each named session parameter should have the following structure:

```
Plugins.<PluginName>.Sessions.<SessionName>.<Parameter>=<Value>
```

A session name can be used as a connString item key parameter instead of specifying a URI, username, and/or password separately.

In item keys, the first parameter can be either a connString or a URI. If the first key parameter doesn't match any session name, it will be treated as a URI. Note that passing embedded URI credentials in the item key is not supported, use named session parameters instead.

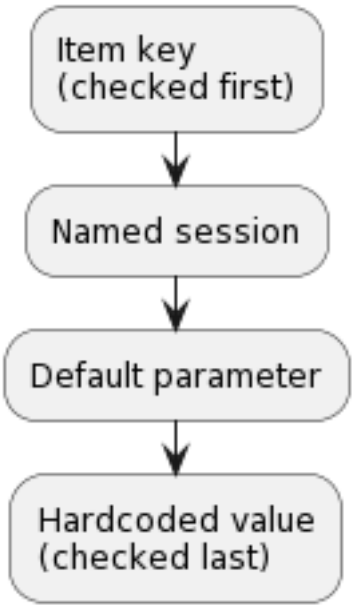
The list of available named session parameters depends on the plugin.

It is possible to override session parameters by specifying new values in the item key parameters (see example).

If a parameter is not defined for the named session, Zabbix agent 2 will use the value defined in the default plugin parameter.

Parameter priority

Zabbix agent 2 plugins search for connection-related parameter values in the following order:



1. The first item key parameter is compared to session names. If no match is found, it is treated as an actual value; in this case, step 3 will be skipped. If a match is found, the parameter value (usually, a URI) must be defined in the named session.
2. Other parameters will be taken from the item key if defined.
3. If an item key parameter (for example, password) is empty, plugin will look for the corresponding named session parameter.
4. If the session parameter is also not specified, the value defined in the corresponding default parameter will be used.
5. If all else fails, the plugin will use the hardcoded default value.

Example 1

Monitoring of two instances ”MySQL1” and ”MySQL2”.

Configuration parameters:

```
Plugins.Mysql.Sessions.MySQL1.Uri=tcp://127.0.0.1:3306
Plugins.Mysql.Sessions.MySQL1.User=mysql1_user
Plugins.Mysql.Sessions.MySQL1.Password=unique_password
Plugins.Mysql.Sessions.MySQL2.Uri=tcp://192.0.2.1:3306
Plugins.Mysql.Sessions.MySQL2.User=mysql2_user
Plugins.Mysql.Sessions.MySQL2.Password=different_password
```

As a result of this configuration, each session name may be used as a connString in an item key, e.g., `mysql.ping[MySQL1]` or `mysql.ping[MySQL2]`.

Example 2

Providing some of the parameters in the item key.

Configuration parameters:

```
Plugins.PostgreSQL.Sessions.Session1.Uri=tcp://192.0.2.234:5432
Plugins.PostgreSQL.Sessions.Session1.User=old_username
Plugins.PostgreSQL.Sessions.Session1.Password=session_password
```

Item key: `pgsql.ping[session1,new_username,,postgres]`

As a result of this configuration, the agent will connect to PostgreSQL using the following parameters:

- URI from session parameter: `192.0.2.234:5432`
- Username from the item key: `new_username`
- Password from session parameter (since it is omitted in the item key): `session_password`
- Database name from the item key: `postgres`

Example 3

Collecting a metric using default configuration parameters.

Configuration parameters:

```
Plugins.PostgreSQL.Default.Uri=tcp://192.0.2.234:5432
Plugins.PostgreSQL.Default.User=zabbix
Plugins.PostgreSQL.Default.Password=password
```

Item key: `pgsql.ping[,,,postgres]`

As a result of this configuration, the agent will connect to PostgreSQL using the parameters:

- Default URI: `192.0.2.234:5432`
- Default username: `zabbix`
- Default password: `password`
- Database name from the item key: `postgres`

Connections

Some plugins support gathering metrics from multiple instances simultaneously. Both local and remote instances can be monitored. TCP and Unix-socket connections are supported.

It is recommended to configure plugins to keep connections to instances in an open state. The benefits are reduced network congestion, latency, and CPU and memory usage due to the lower number of connections. The client library takes care of this.

> **Note:**
> Time period for which unused connections should remain open can be determined by `Plugins.<PluginName>.KeepAlive` parameter. Example: `Plugins.Memcached.KeepAlive`

Ceph plugin

Overview

Configuration parameters allow customizing the Ceph plugin behavior.

It is recommended to specify them in their own configuration file (e.g., `ceph.conf`), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.Ceph.Default.ApiKey	Specify the default API key for connecting to Ceph.
Plugins.Ceph.Default.Mode	Specify the default mode for connecting to Ceph.
Plugins.Ceph.Default.URI	Specify the default URI for connecting to Ceph.
Plugins.Ceph.Default.User	Specify the default user name for connecting to Ceph.
Plugins.Ceph.InsecureSkipVerify	Specify whether an HTTP client should verify the server's certificate chain and host name.
Plugins.Ceph.KeepAlive	Set the maximum waiting time before unused plugin connections are closed.
Plugins.Ceph.Sessions.<SessionName>.ApiKey	Specify the named session API key for connecting to Ceph.
Plugins.Ceph.Sessions.<SessionName>.Mode	Specify the named session mode for connecting to Ceph.
Plugins.Ceph.Sessions.<SessionName>.URI	Specify the named session URI for connecting to Ceph.
Plugins.Ceph.Sessions.<SessionName>.User	Specify the named session user name for connecting to Ceph.
Plugins.Ceph.Timeout	Set the request execution timeout.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files;
- Values support environment variables;
- Zabbix supports configuration files only in UTF-8 encoding without BOM;
- Comments starting with "#" are only supported at the beginning of the line.

Parameter details

Plugins.Ceph.Default.ApiKey

Specify the default API key for connecting to Ceph.
 Used if no value is specified in an item key or named session.

Plugins.Ceph.Default.Mode

Specify the default mode for connecting to Ceph.
 The native mode is only supported on Linux and starting with Ceph 16. The restful mode will not work with Ceph version 20 or newer.
 Note that the user credential set differs for each mode and they are not compatible with each other.

Default: `restful` (deprecated)
 Range: native/restful

Plugins.Ceph.Default.URI

Specify the default URI for connecting to Ceph.
 Used if no value is specified in an item key or named session.

Should not include embedded credentials (they will be ignored).
Must match the URI format.
Only `https` scheme is supported; a scheme can be omitted.
A port can be omitted (default=8003).

Default: `https://localhost:8003`

Examples:

```
https://127.0.0.1:8003
localhost
```

Plugins.Ceph.Default.User

Specify the default user name for connecting to Ceph.
 Used if no value is specified in an item key or named session.

Plugins.Ceph.InsecureSkipVerify

Specify whether an HTTP client should verify the server's certificate chain and host name.
 If *true*, TLS accepts any certificate presented by the server and any host name in that certificate. In this mode, TLS is susceptible to man-in-the-middle attacks (use only for testing).

Default: `false`
 Range: false/true

Plugins.Ceph.KeepAlive

Set the maximum waiting time (in seconds) before unused plugin connections are closed.

Default: 300
 Range: 60-900

Plugins.Ceph.Sessions.<SessionName>.ApiKey

Specify the named session API key for connecting to Ceph.
 <SessionName> - define name of a session for using in item keys.

Plugins.Ceph.Sessions.<SessionName>.Mode

Specify the named session mode for connecting to Ceph.
 <SessionName> - define name of a session for using in item keys.
 The native mode is only supported on Linux and starting with Ceph 16. The restful mode will not work with Ceph version 20 or newer.
 Note that the user credential set differs for each mode and they are not compatible with each other.

Default: `restful` (deprecated)
 Range: native/restful

Plugins.Ceph.Sessions.<SessionName>.URI

Specify the named session URI for connecting to Ceph.
 <SessionName> - define name of a session for using in item keys.
Should not include embedded credentials (they will be ignored).
Must match the URI format.
 Only `https` scheme is supported; a scheme can be omitted.
A port can be omitted (default=8003).

Default: `https://localhost:8003`

Examples:

```
https://127.0.0.1:8003
localhost
```

Plugins.Ceph.Sessions.<SessionName>.User

Specify the named session user name for connecting to Ceph.
 <SessionName> - define name of a session for using in item keys.

Plugins.Ceph.Timeout

Set the request execution timeout (duration to wait for a request to complete before shutting it down) in seconds.

Default: `global timeout`
 Range: 1-30

See also

- Instructions for configuring plugins
- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)

Docker plugin

Overview

Configuration parameters allow customizing the Docker plugin behavior.

It is recommended to specify them in their own configuration file (e.g., `docker.conf`), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.Docker.Endpoint	Specify the Docker daemon UNIX-socket location..
Plugins.Docker.Timeout	Specify the request execution timeout.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files;
- Values support environment variables;
- Zabbix supports configuration files only in UTF-8 encoding without BOM;
- Comments starting with "#" are only supported at the beginning of the line.

Parameter details

Plugins.Docker.Endpoint

Specify the Docker daemon UNIX-socket location.
 Must contain a scheme (only `unix://` is supported).

Default: `unix:///var/run/docker.sock`

Plugins.Docker.Timeout

Specify the request execution timeout (the duration, in seconds, to wait for a request to complete before shutting it down).

Default: `global timeout`
 Range: 1–30

See also

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

Ember+ plugin

Overview

Configuration parameters allow customizing the Ember+ plugin behavior.

It is recommended to specify them in their own configuration file (e.g., `emberplus.conf`), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

The Ember+ plugin is a loadable plugin and is available and fully described in the Ember+ plugin repository.

This plugin is currently only available to be built from the source (for both Unix and Windows).

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.EmberPlus.Default.URL	Specify the default URI for connecting to Ember+.
Plugins.EmberPlus.KeepAlive	Set the maximum waiting time before unused plugin connections are closed.
Plugins.EmberPlus.Sessions.<SessionName>.URL	Specify the named session URI for connecting to Ember+.
Plugins.EmberPlus.System.Path	Specify path to the Ember+ plugin executable.

Parameter	Description
Plugins.EmberPlus.Timeout	Set the request execution timeout.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files.
- Values support environment variables;
- Zabbix supports configuration files only in UTF-8 encoding without BOM.
- Comments starting with "#" are only supported at the beginning of the line.

Parameter details

Plugins.EmberPlus.Default.Uri

Specify the default URI for connecting to Ember+.
 The only supported schema is `tcp://`. A schema can be omitted. Embedded credentials will be ignored.

Default: `tcp://localhost:9998`

Plugins.EmberPlus.KeepAlive

Set the maximum time of waiting (in seconds) before unused plugin connections are closed.

Default: 300
 Range: 60–900

Plugins.EmberPlus.Sessions.<SessionName>.Uri

Specify the named session URI for connecting to Ember+.
 The only supported schema is `tcp://`. A schema can be omitted. Embedded credentials will be ignored.
 <SessionName> - define name of a session for using in item keys.

Default: `tcp://localhost:9998`

Plugins.EmberPlus.System.Path

Specify path to the Ember+ plugin executable.

Example: `Plugins.EmberPlus.System.Path=/usr/libexec/zabbix/zabbix-agent2-plugin-ember-plus`

Plugins.EmberPlus.Timeout

Set the duration, in seconds, to wait for a server to respond when first connecting and on follow-up operations in the session.

Default: `global timeout`
 Range: 1–30

Options

Parameter	Description
-V --version	Print the plugin version and license information.
-h --help	Print help information (shorthand).
-t, --test <item key>	Launch plugin for testing (plugin config ignored).

See also:

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

Memcached plugin

Overview

Configuration parameters allow customizing the Memcached plugin behavior.

It is recommended to specify them in their own configuration file (e.g., `memcached.conf`), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.Memcached.Default.Password	Specify the default password for connecting to Memcached.
Plugins.Memcached.Default.Uri	Specify the default URI for connecting to Memcached.
Plugins.Memcached.Default.User	Specify the default username for connecting to Memcached.
Plugins.Memcached.KeepAlive	Set the maximum waiting time before unused plugin connections are closed.
Plugins.Memcached.Sessions.<SessionName>.Password	Specify the named session password for connecting to Memcached.
Plugins.Memcached.Sessions.<SessionName>.Uri	Specify the named session URI for connecting to Memcached.
Plugins.Memcached.Sessions.<SessionName>.User	Specify the named session username for connecting to Memcached.
Plugins.Memcached.Timeout	Set the request execution timeout.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files;
- Values support environment variables;
- Zabbix supports configuration files only in UTF-8 encoding without BOM;
- Comments starting with "#" are only supported at the beginning of the line.

Parameter details

Plugins.Memcached.Default.Password

Specify the default password for connecting to Memcached.
 Used if no value is specified in an item key or named session.

Plugins.Memcached.Default.Uri

Specify the default URI for connecting to Memcached.
 Used if no value is specified in an item key or named session.
 Embedded credentials will be ignored.
 Must match the URI format. Supported schemes: `tcp`, `unix`.
 A scheme can be omitted.
 A port can be omitted (default=11211).

Default: `tcp://localhost:11211`

Examples:

```
tcp://localhost:11211
localhost
unix:/var/run/memcached.sock
```

Plugins.Memcached.Default.User

Specify the default username for connecting to Memcached.
 Used if no value is specified in an item key or named session.

Plugins.Memcached.KeepAlive

Set the maximum time of waiting (in seconds) before unused plugin connections are closed.

Default: 300
 Range: 60–900

Plugins.Memcached.Sessions.<SessionName>.Password

Specify the named session password for connecting to Memcached.
 <SessionName> - define name of a session for using in item keys.

Plugins.Memcached.Sessions.<SessionName>.Uri

Specify the named session URI for connecting to Memcached.
 <SessionName> - define name of a session for using in item keys.
 Embedded credentials will be ignored.
 Must match the URI format. Supported schemes: `tcp`, `unix`.
 A scheme can be omitted.
 A port can be omitted (default=11211).

Examples:

```
tcp://localhost:11211
localhost
unix:/var/run/memcached.sock
```

Plugins.Memcached.Sessions.<SessionName>.User

Specify the named session username for connecting to Memcached.
 <SessionName> - define name of a session for using in item keys.

Plugins.Memcached.Timeout

Set the request execution timeout (the duration, in seconds, to wait for a request to complete before shutting it down).

Default: `global timeout`
 Range: 1–30

See also:

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

Modbus plugin

Overview

Configuration parameters allow customizing the Modbus plugin behavior.

It is recommended to specify them in their own configuration file (e.g., `modbus.conf`), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.Modbus.Sessions.<SessionName>.Endpoint	Specify the named session endpoint for connecting to Modbus.
Plugins.Modbus.Sessions.<SessionName>.SlaveID	Specify the named session slave ID for connecting to Modbus.
Plugins.Modbus.Sessions.<SessionName>.Timeout	Set the named session timeout in seconds.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files;
- Values support environment variables;
- Zabbix supports configuration files only in UTF-8 encoding without BOM;
- Comments starting with "#" are only supported at the beginning of the line.

Parameter details

Plugins.Modbus.Sessions.<SessionName>.Endpoint

Specify the named session endpoint for connecting to Modbus.
 The connection string consists of a protocol scheme, a host address and a port or serial port name and attributes.
 <SessionName> - define name of a session for using in item keys.

Plugins.Modbus.Sessions.<SessionName>.SlaveID

Specify the named session slave ID for connecting to Modbus.
 <SessionName> - define name of a session for using in item keys.

Example: `Plugins.Modbus.Sessions.MB1.SlaveID=20`

Note that this named session parameter is checked only if the value provided in the item key slave ID parameter is empty.

Plugins.Modbus.Sessions.<SessionName>.Timeout

Set the named session timeout in seconds.
 <SessionName> - define name of a session for using in item keys.

Example: `Plugins.Modbus.Sessions.MB1.Timeout=2`

If you need to set the request execution timeout (how long to wait for a request to complete before shutting it down), use the item configuration form.

See also

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

MongoDB plugin

Overview

Configuration parameters allow customizing the MongoDB plugin behavior.

It is recommended to specify them in their own configuration file (e.g., `mongodb.conf`), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

The MongoDB plugin is a loadable plugin and is available and fully described in the MongoDB plugin repository.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

When connecting to MongoDB, the plugin checks its configuration in a specific order, which determines which value to use. The order, from highest to lowest priority, is as follows:

1. Parameters specified in the connection URI (e.g., query parameters defined after ? in `Plugins.MongoDB.Sessions.<session_name` have the highest priority).
2. Parameters in the Zabbix agent 2 item key (e.g., mongodb.collection.stats).
3. Parameters defined in a named session (e.g., `Plugins.MongoDB.Sessions.<SessionName>.User`, `Plugins.MongoDB.Sessions`
4. Default values (e.g., `Plugins.MongoDB.Default.User`, `Plugins.MongoDB.Default.Password`).

For example:

- If you specify the database for authenticating a connection to the MongoDB server in the `*.<SessionName>.Uri` parameter (e.g., `Plugins.MongoDB.Sessions.<session_name>.Uri=mongodb://user:password@127.0.0.1:27017/?authSource=a` then the plugin will consider the `*.<SessionName>.AuthSource` parameter to be set to `admin`, even if the `*.<SessionName>.AuthSource` parameter is set to a different value.
- If you specify the URI in the `*.<SessionName>.Uri` parameter (e.g., `Plugins.MongoDB.Sessions.<session_name>.Uri=mongo` but provide the username and the password in the Zabbix agent 2 item key, the plugin will use the username and password from the item key, even if `Plugins.MongoDB.Sessions.<SessionName>.User` and `Plugins.MongoDB.Sessions.<SessionNam` are configured.

Attention:
If the `mongodb+srv://` scheme is used, the URI may retrieve parameters from a DNS TXT record. Parameters obtained from this TXT record override all other configurations, including URI parameters, Zabbix Agent 2 item key parameters, and session parameters.
 Only one TXT record may exist for the hostname used in the connection string.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.MongoDB.Default.AuthMechanism	Specify the default authentication mechanism for connecting to MongoDB server.
Plugins.MongoDB.Default.AuthSource	Specify the default database for authenticating connection to MongoDB server.
Plugins.MongoDB.Default.Disable	Disable automatic discovery of MongoDB replica set nodes.
Plugins.MongoDB.Default.Password	Specify the default password.
Plugins.MongoDB.Default.ReplicaSet	Specify the default replica set name, used for automatic discovery of MongoDB replica set nodes.
Plugins.MongoDB.Default.ReadPreference	Specify the default order for connecting to MongoDB replica set nodes or fallback.
Plugins.MongoDB.Default.TLSConnect	Specify the level of encryption required for communications between Zabbix agent 2 and monitored databases.
Plugins.MongoDB.Default.TLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification.
Plugins.MongoDB.Default.TLSCertFile	Specify full pathname of the file containing the agent certificate or certificate chain.
Plugins.MongoDB.Default.TLSKeyFile	Specify full pathname of the file containing the database private key.
Plugins.MongoDB.Default.Uri	Specify the default URI.
Plugins.MongoDB.Default.User	Specify the default username.
Plugins.MongoDB.KeepAlive	Set the maximum time of waiting before unused plugin connections are closed.
Plugins.MongoDB.Sessions.<SessionName>.AuthMechanism	Specify the named session authentication mechanism for connecting to MongoDB server.
Plugins.MongoDB.Sessions.<SessionName>.AuthSource	Specify the named session database for authenticating connection to MongoDB server.
Plugins.MongoDB.Sessions.<SessionName>.Disable	Disable automatic discovery of MongoDB replica set nodes for named session.
Plugins.MongoDB.Sessions.<SessionName>.Password	Specify the named session password.
Plugins.MongoDB.Sessions.<SessionName>.ReplicaSet	Specify the named session replica set name, used for automatic discovery of MongoDB replica set nodes.
Plugins.MongoDB.Sessions.<SessionName>.ReadPreference	Specify the named session order for connecting to MongoDB replica set nodes or fallback.
Plugins.MongoDB.Sessions.<SessionName>.TLSConnect	Specify the level of encryption required for communications between Zabbix agent 2 and monitored databases for named session.
Plugins.MongoDB.Sessions.<SessionName>.TLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification for named session.
Plugins.MongoDB.Sessions.<SessionName>.TLSCertFile	Specify full pathname of the file containing the agent certificate or certificate chain for named session.
Plugins.MongoDB.Sessions.<SessionName>.TLSKeyFile	Specify full pathname of the file containing the database private key for named session.
Plugins.MongoDB.Sessions.<SessionName>.Uri	Specify the named session URI.

Parameter	Description
Plugins.MongoDB.Sessions.<SessionName>.User	Specify the named session username.
Plugins.MongoDB.System.Path	Specify path to the plugin executable.
Plugins.MongoDB.Timeout	Set the request execution timeout.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files.
- Values support environment variables.
- Zabbix supports configuration files only in UTF-8 encoding without BOM.
- Comments starting with "#" are only supported at the beginning of the line.

Parameter details

Plugins.MongoDB.Default.AuthMechanism

Specify the default authentication mechanism for connecting to MongoDB server.
 <SessionName> - define name of a session for using in item keys.

Range: `SCRAM-SHA-1`, `SCRAM-SHA-256`, or `MONGODB-X509`

Plugins.MongoDB.Default.AuthSource

Specify the default database for authenticating connection to MongoDB server.
 <SessionName> - define name of a session for using in item keys.

Default: `admin`

If `MONGODB-X509` authentication mechanism is used, this parameter is automatically set to `$external`; modifying it to any other value will cause an error when launching.

Plugins.MongoDB.Default.Direct

Disable automatic discovery of MongoDB replica set nodes.
 <SessionName> - define name of a session for using in item keys.

Default: `false`
 Range: `true` - connect only to the specified server, ignoring other members of the replica set, `false` - automatically discover all nodes in the replica set.

If the `tcp://` (deprecated) scheme is used, the parameter is set to `true` by default.

Plugins.MongoDB.Default.Password

Specify the default password.
 <SessionName> - define name of a session for using in item keys.

Plugins.MongoDB.Default.ReplicaSet

Specify the default replica set name, used for automatic discovery of MongoDB replica set nodes.
 <SessionName> - define name of a session for using in item keys.
 If this parameter is used with the `mongodb+srv://` scheme, it will behave as if multiple URLs have been provided; a DNS server may have a TXT record that can override this (or any other) parameter.
 Unavailable if the `Plugins.MongoDB.Sessions.<SessionName>.Direct` parameter is set to `true`; the agent will fail to start.

Plugins.MongoDB.Default.ReadPreference

Specify the default order for connecting to MongoDB replica set nodes or fallback.
 <SessionName> - define name of a session for using in item keys.
 Unavailable if the `Plugins.MongoDB.Sessions.<SessionName>.Direct` parameter is set to `true`; the agent will fail to start.

Range: `primary`, `primaryPreferred`, `secondary`, `secondaryPreferred`, `nearest`, `PrimaryMode`, `PrimaryPreferredMode`, `SecondaryMode`, `SecondaryPreferredMode`, or `NearestMode`

The values `primary`, `primaryPreferred`, `secondary`, `secondaryPreferred`, and `nearest` are preferred. Values ending in `Mode` are supported as legacy compatibility aliases.

Plugins.MongoDB.Default.TLSConnect

Specify the level of encryption required for communications between Zabbix agent 2 and monitored databases.
 <SessionName> - define name of a session for using in item keys.
 Supported since plugin version 1.2.1.

Range: `required` - require TLS connection, `verify_ca` - verify certificates, `verify_full` - verify certificates and IP address

Plugins.MongoDB.Default.TLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix agent 2 and monitored databases.
 <SessionName> - define name of a session for using in item keys.

Mandatory: `yes`, if `Plugins.MongoDB.Sessions.<SessionName>.TLSConnect` is set to `verify_ca` or `verify_full`

Plugins.MongoDB.Default.TLSCertFile

Specify full pathname of the file containing the agent certificate or certificate chain, used for encrypted communications between Zabbix agent 2 and monitored databases.
 <SessionName> - define name of a session for using in item keys.

Mandatory: `yes`, if `Plugins.MongoDB.Sessions.<SessionName>.TLSKeyFile` is specified

Plugins.MongoDB.Default.TLSKeyFile

Specify full pathname of the file containing the database private key, used for encrypted communications between Zabbix agent 2 and monitored databases.
 <SessionName> - define name of a session for using in item keys.

Mandatory: `yes`, if `Plugins.MongoDB.Sessions.<SessionName>.TLSCertFile` is specified

Plugins.MongoDB.Default.Uri

Specify the default URI.
 <SessionName> - define name of a session for using in item keys.
 Must match the URI format.
 Supported schemes: `mongodb://` and `mongodb+srv://`; the `tcp://` scheme is deprecated and kept for backward compatibility with existing configurations.
 A port can be omitted (default=27017).
Examples: “

Default: `tcp://127.0.0.1:27017`

Examples:

```
mongodb://127.0.0.1:27017
mongodb+srv://example.com
localhost
```

Plugins.MongoDB.Default.User

Specify the default username.
 <SessionName> - define name of a session for using in item keys.

Plugins.MongoDB.KeepAlive

Set the maximum time of waiting (in seconds) before unused plugin connections are closed.

Default: 300
 Range: 60–900

Plugins.MongoDB.Sessions.<SessionName>.AuthMechanism

Specify the named session authentication mechanism for connecting to MongoDB server.
 <SessionName> - define name of a session for using in item keys.

Range: `SCRAM-SHA-1`, `SCRAM-SHA-256`, or `MONGODB-X509`

Plugins.MongoDB.Sessions.<SessionName>.AuthSource

Specify the named session database for authenticating connection to MongoDB server.
 <SessionName> - define name of a session for using in item keys.

Default: `admin`

If `MONGODB-X509` authentication mechanism is used, this parameter is automatically set to `$external`; modifying it to any other value will cause an error when launching.

Plugins.MongoDB.Sessions.<SessionName>.Direct

Disable automatic discovery of MongoDB replica set nodes for named session.
 <SessionName> - define name of a session for using in item keys.

Default: `false`
 Range: `true` - connect only to the specified server, ignoring other members of the replica set, `false` - automatically discover all nodes in the replica set.

If the `tcp://` (deprecated) scheme is used, the parameter is set to `true` by default.

Plugins.MongoDB.Sessions.<SessionName>.Password

Specify the named session password.
 <SessionName> - define name of a session for using in item keys.

Plugins.MongoDB.Sessions.<SessionName>.ReplicaSet

Specify the named session replica set name, used for automatic discovery of MongoDB replica set nodes.
 <SessionName> - define name of a session for using in item keys.
 If this parameter is used with the `mongodb+srv://` scheme, it will behave

as if multiple URLs have been provided; a DNS server may have a TXT record that can override this (or any other) parameter.
 Unavailable if the `Plugins.MongoDB.Sessions.<SessionName>.Direct` parameter is set to `true`; the agent will fail to start.

Plugins.MongoDB.Sessions.<SessionName>.ReadPreference

Specify the named session order for connecting to MongoDB replica set nodes or fallback.
 <SessionName> - define name of a session for using in item keys.
 Unavailable if the `Plugins.MongoDB.Sessions.<SessionName>.Direct` parameter is set to `true`; the agent will fail to start.

Range: `primary`, `primaryPreferred`, `secondary`, `secondaryPreferred`, `nearest`, `PrimaryMode`, `PrimaryPreferredMode`, `SecondaryMode`, `SecondaryPreferredMode`, or `NearestMode`

The values `primary`, `primaryPreferred`, `secondary`, `secondaryPreferred`, and `nearest` are preferred. Values ending in `Mode` are supported as legacy compatibility aliases.

Plugins.MongoDB.Sessions.<SessionName>.TLSConnect

Specify the level of encryption required for communications between Zabbix agent 2 and monitored databases for named session.
 <SessionName> - define name of a session for using in item keys.
 Supported since plugin version 1.2.1.

Range: `required` - require TLS connection, `verify_ca` - verify certificates, `verify_full` - verify certificates and IP address

Plugins.MongoDB.Sessions.<SessionName>.TLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix agent 2 and monitored databases for named session.
 <SessionName> - define name of a session for using in item keys.

Mandatory: `yes`, if `Plugins.MongoDB.Sessions.<SessionName>.TLSConnect` is set to `verify_ca` or `verify_full`

Plugins.MongoDB.Sessions.<SessionName>.TLSCertFile

Specify full pathname of the file containing the agent certificate or certificate chain, used for encrypted communications between Zabbix agent 2 and monitored databases for named session.
 <SessionName> - define name of a session for using in item keys.

Mandatory: `yes`, if `Plugins.MongoDB.Sessions.<SessionName>.TLSKeyFile` is specified

Plugins.MongoDB.Sessions.<SessionName>.TLSKeyFile

Specify full pathname of the file containing the database private key, used for encrypted communications between Zabbix agent 2 and monitored databases for named session.
 <SessionName> - define name of a session for using in item keys.

Mandatory: `yes`, if `Plugins.MongoDB.Sessions.<SessionName>.TLSCertFile` is specified

Plugins.MongoDB.Sessions.<SessionName>.Uri

Specify the named session URI.
 <SessionName> - define name of a session for using in item keys.
 Must match the URI format.
 Supported schemes: `mongodb://` and `mongodb+srv://`; the `tcp://` scheme is deprecated and kept for backward compatibility with existing configurations.
 A port can be omitted (default=27017).
Examples: “

Default: `tcp://127.0.0.1:27017`

Examples:

```
mongodb://127.0.0.1:27017
mongodb+srv://example.com
localhost
```

Plugins.MongoDB.Sessions.<SessionName>.User

Specify the named session username.
 <SessionName> - define name of a session for using in item keys.

Plugins.MongoDB.System.Path

Specify path to the MongoDB plugin executable.

Example usage: `Plugins.MongoDB.System.Path=/usr/libexec/zabbix/zabbix-agent2-plugin-mongodb`

Plugins.MongoDB.Timeout

Set the request execution timeout (the duration, in seconds, to wait for a request to complete before shutting it down).

Default: `global timeout`
 Range: 1-30

Options

Parameter	Description
-V --version	Print the plugin version and license information.
-h --help	Print help information (shorthand).
-t, --test <item key>	Launch plugin for testing (plugin config ignored).

See also

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

MQTT plugin

Overview

Configuration parameters allow customizing the MQTT plugin behavior.

It is recommended to specify them in their own configuration file (e.g., `mqtt.conf`), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.MQTT.Default.Password	Specify the default password for connecting to MQTT.
Plugins.MQTT.Default.TLSCAFile	Specify full pathname of the file containing top-level CA(s) certificates for peer certificate verification.
Plugins.MQTT.Default.TLSCertFile	Specify full pathname of the file containing the agent certificate or certificate chain.
Plugins.MQTT.Default.TLSKeyFile	Specify full pathname of the file containing the MQTT private key.
Plugins.MQTT.Default.Topic	Specify the default topic for MQTT subscription.
Plugins.MQTT.Default.Url	Specify the default MQTT broker connection string.
Plugins.MQTT.Default.User	Specify the default username for connecting to MQTT.
Plugins.MQTT.Sessions.<Session name>.Password	Specify the named session password for connecting to MQTT.
Plugins.MQTT.Sessions.<Session name>.TLSCAFile	Specify full pathname of the file containing top-level CA(s) certificates for peer certificate verification for named session.
Plugins.MQTT.Sessions.<Session name>.TLSCertFile	Specify full pathname of the file containing the agent certificate or certificate chain for named session.
Plugins.MQTT.Sessions.<Session name>.TLSKeyFile	Specify full pathname of the file containing the MQTT private key for named session.
Plugins.MQTT.Sessions.<Session name>.Topic	Specify the named session topic for MQTT subscription.
Plugins.MQTT.Sessions.<Session name>.Url	Specify the named session MQTT broker connection string.
Plugins.MQTT.Sessions.<Session name>.User	Specify the named session username for connecting to MQTT.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files;
- Values support environment variables;
- Zabbix supports configuration files only in UTF-8 encoding without BOM;
- Comments starting with ”#” are only supported at the beginning of the line.

Parameter details

Plugins.MQTT.Default.Password

Specify the default password for connecting to MQTT.
 Used if no value is specified in an item key or named session.

Plugins.MQTT.Default.TLSCAFile

Specify full pathname of the file containing top-level CA(s) certificates for peer certificate verification for encrypted communications between Zabbix agent 2 and MQTT broker.
 Used if no value is specified in a named session.

Plugins.MQTT.Default.TLSCertFile

Specify full pathname of the file containing the agent certificate or certificate chain for encrypted communications between Zabbix agent 2 and MQTT broker.
 Used if no value is specified in a named session.

Plugins.MQTT.Default.TLSKeyFile

Specify full pathname of the file containing the MQTT private key for encrypted communications between Zabbix agent 2 and MQTT broker.
 Used if no value is specified in a named session.

Plugins.MQTT.Default.Topic

Specify the default topic for MQTT subscription.
 Used if no value is specified in an item key or named session.
 The topic may contain wildcards (”+”,”#”).

Examples:

```
path/to/file
path/to/#
path/+/topic
```

Plugins.MQTT.Default.Url

Specify the default MQTT broker connection string.
 Used if no value is specified in an item key or named session.
 Should not include query parameters.
 Must match the URL format.
 Supported schemes: `tcp` (default), `ws`, `tls`
 A scheme can be omitted.
 A port can be omitted (default=1883)

Default: `tcp://localhost:1883`

Examples:

```
tcp://host:1883
localhost
ws://host:8080
```

Plugins.MQTT.Default.User

Specify the default username for connecting to MQTT.
 Used if no value is specified in an item key or named session.

Plugins.MQTT.Sessions.<SessionName>.Password

Specify the named session password for connecting to MQTT.
 <SessionName> - define name of a session for using in item keys.

Plugins.MQTT.Sessions.<SessionName>.TLSCAFile

Specify full pathname of the file containing top-level CA(s) certificates for peer certificate verification for named session, used for encrypted communications between Zabbix agent 2 and MQTT broker.
 <SessionName> - define name of a session for using in item keys.

Plugins.MQTT.Sessions.<SessionName>.TLSCertFile

Specify full pathname of the file containing the agent certificate or certificate chain for named session, used for encrypted communications between Zabbix agent 2 and MQTT broker.
 <SessionName> - define name of a session for using in item keys.

Plugins.MQTT.Sessions.<SessionName>.TLSKeyFile

Specify full pathname of the file containing the MQTT private key for named session, used for encrypted communications between Zabbix agent 2 and MQTT broker.
 <SessionName> - define name of a session for using in item keys.

Plugins.MQTT.Sessions.<SessionName>.Topic

Specify the named session topic for MQTT subscription.
 <SessionName> - define name of a session for using in item keys.
 The topic may contain wildcards (”+”,”#”).

Examples:

```
path/to/file
path/to/#
path/+/topic
```

Plugins.MQTT.Sessions.<SessionName>.Url

Specify the named session MQTT broker connection string.
 <SessionName> - define name of a session for using in item keys.
 Should not include query parameters.
 Must match the URL format.
 Supported schemes: `tcp` (default), `ws`, `tls`
 A scheme can be omitted.
 A port can be omitted (default=1883)

Default: `tcp://localhost:1883`

Examples:

```
tcp://host:1883
localhost
ws://host:8080
```

Plugins.MQTT.Sessions.<SessionName>.User

Specify the named session username for connecting to MQTT.
 <SessionName> - define name of a session for using in item keys.

If you need to set the request execution timeout (how long to wait for a request to complete before shutting it down), use the item configuration form.

See also

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

MSSQL plugin

Overview

Configuration parameters allow customizing the MSSQL plugin behavior.

It is recommended to specify them in their own configuration file (e.g., `mssql.conf`), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

The MSSQL plugin is a loadable plugin and is available and fully described in the MSSQL plugin repository.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.MSSQL.CustomQueriesDir	Specify path to the directory containing user-defined `.sql` files with custom queries that the plugin can execute.
Plugins.MSSQL.CustomQueriesEnabled	Enable execution of the `mssql.custom.query` item key.
Plugins.MSSQL.Default.CACertPath	Specify the default file path to the public key certificate of the certificate authority (CA) that issued the certificate of the MSSQL server.
Plugins.MSSQL.Default.Database	Specify the default database name.
Plugins.MSSQL.Default.Encrypt	Specify the level of default encryption.
Plugins.MSSQL.Default.HostnameInCertificate	Specify the common name (CN) of the MSSQL server certificate by default.
Plugins.MSSQL.Default.Password	Specify the default password.
Plugins.MSSQL.Default.TLSMinVersion	Specify the minimum TLS version by default.
Plugins.MSSQL.Default.TrustServerCertificate	Trust the server certificate without validating it by default.
Plugins.MSSQL.Default.Uri	Specify the default URI.
Plugins.MSSQL.Default.User	Specify the default username.
Plugins.MSSQL.KeepAlive	Specify the maximum time of waiting (in seconds) before unused plugin connections are closed.
Plugins.MSSQL.Sessions.<SessionName>.CACertPath	Specify the named session file path to the public key certificate of the certificate authority (CA) that issued the certificate of the MSSQL server.
Plugins.MSSQL.Sessions.<SessionName>.Database	Specify the named session database name.
Plugins.MSSQL.Sessions.<SessionName>.Encrypt	Specify the level of named session encryption.
Plugins.MSSQL.Sessions.<SessionName>.HostnameInCertificate	Specify the common name (CN) of the MSSQL server certificate for named session.
Plugins.MSSQL.Sessions.<SessionName>.Password	Specify the named session password.
Plugins.MSSQL.Sessions.<SessionName>.TLSMinVersion	Specify the named session minimum TLS version.
Plugins.MSSQL.Sessions.<SessionName>.TrustServerCertificate	Trust the named session server certificate without validating it.
Plugins.MSSQL.Sessions.<SessionName>.Uri	Specify the named session URI.
Plugins.MSSQL.Sessions.<SessionName>.User	Specify the named session username.
Plugins.MSSQL.System.Path	Specify path to the MSSQL plugin executable.
Plugins.MSSQL.Timeout	Set the duration, in seconds, to wait for a server to respond when first connecting and on follow-up operations in the session.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files.

- Values support environment variables;
- Zabbix supports configuration files only in UTF-8 encoding without BOM.
- Comments starting with "#" are only supported at the beginning of the line.

Parameter details

Plugins.MSSQL.CustomQueriesDir

Specify path to the directory containing user-defined `.sql` files with custom queries that the plugin can execute.
 The plugin loads all available .sql files in the configured directory at startup. This means that any changes to the custom query files will not be reflected until the plugin is restarted. The plugin is started and stopped together with Zabbix agent 2.

Defaults:
`/usr/local/share/zabbix/custom-queries/mssql` - for Unix systems,
`*:\Program Files\Zabbix Agent 2\Custom Queries\MSSQL` - for Windows systems (where `*` is the drive name from the `ProgramFiles` environment variable)

Plugins.MSSQL.CustomQueriesEnabled

Enable execution of the `mssql.custom.query` item key.
 If disabled, no queries will be loaded from the custom query directory path.

Default: `false`
 Range: `true/false`

Plugins.MSSQL.Default.CACertPath

Specify the default file path to the public key certificate of the certificate authority (CA) that issued the certificate of the MSSQL server.
 The certificate must be in PEM format.

Plugins.MSSQL.Default.Database

Specify the default database name.

Plugins.MSSQL.Default.Encrypt

Specify the level of default encryption.

Range:
true - data sending between plugin and server is encrypted;
false - data sending between plugin and server is not encrypted beyond the login packet;
strict - data sending between plugin and server is encrypted E2E using TDS8;
disable - data sending between plugin and server is not encrypted.

Plugins.MSSQL.Default.HostNameInCertificate

Specify the common name (CN) of the MSSQL server certificate by default.

Plugins.MSSQL.Default.Password

Specify the default password for a protected MSSQL server.

Plugins.MSSQL.Default.TLSMinVersion

Specify the minimum TLS version by default.

Range: `1.0`, `1.1`, `1.2`, `1.3`

Plugins.MSSQL.Default.TrustServerCertificate

Trust the server certificate without validating it by default.

Range: `true`, `false`

Plugins.MSSQL.Default.Uri

Specify the default URI for connecting to MSSQL server.
 Embedded credentials will be ignored.
 The only supported schema is `sqlserver://`.
 The schema can be omitted.

Default: `sqlserver://localhost:1433`

Plugins.MSSQL.Default.User

Specify the default username for a protected MSSQL server.

Plugins.MSSQL.KeepAlive

Specify the maximum time of waiting (in seconds) before unused plugin connections are closed.

Default: 300
 Range: 60–900

Plugins.MSSQL.Sessions.<SessionName>.CACertPath

Specify the named session file path to the public key certificate of the certificate authority (CA) that issued the certificate of the MSSQL server.
 <SessionName> - define name of a session for using in item keys.
 The certificate must be in PEM format.

Plugins.MSSQL.Sessions.<SessionName>.Database

Specify the named session database name.
 <SessionName> - define name of a session for using in item keys.

Plugins.MSSQL.Sessions.<SessionName>.Encrypt

Specify the level of named session encryption.
 <SessionName> - define name of a session for using in item keys.

Range:
true - data sending between plugin and server is encrypted;
false - data sending between plugin and server is not encrypted beyond the login packet;
strict - data sending between plugin and server is encrypted E2E using TDS8;
disable - data sending between plugin and server is not encrypted.

Plugins.MSSQL.Sessions.<SessionName>.HostNameInCertificate

Specify the common name (CN) of the MSSQL server certificate for named session.
 <SessionName> - define name of a session for using in item keys.

Plugins.MSSQL.Sessions.<SessionName>.Password

Specify the named session password for a protected MSSQL server.
 <SessionName> - define name of a session for using in item keys.

Plugins.MSSQL.Sessions.<SessionName>.TLSMinVersion

Specify the named session minimum TLS version.
 <SessionName> - define name of a session for using in item keys.

Range: `1.0`, `1.1`, `1.2`, `1.3`

Plugins.MSSQL.Sessions.<SessionName>.TrustServerCertificate

Trust the named session server certificate without validating it.
 <SessionName> - define name of a session for using in item keys.

Range: `true`, `false`

Plugins.MSSQL.Sessions.<SessionName>.Uri

Specify the named session URI for connecting to MSSQL server.
 <SessionName> - define name of a session for using in item keys.
 Embedded credentials will be ignored.
 The only supported schema is `sqlserver://`.
 The schema can be omitted.

Default: `sqlserver://localhost:1433`

Plugins.MSSQL.Sessions.<SessionName>.User

Specify the named session username for a protected MSSQL server.
 <SessionName> - define name of a session for using in item keys.

Plugins.MSSQL.System.Path

Specify path to the MSSQL plugin executable.
 Global setting for the MSSQL plugin. Applied to all connections.

Example: `Plugins.MSSQL.System.Path=/usr/libexec/zabbix/zabbix-agent2-plugin-mssql`

Plugins.MSSQL.Timeout

Set the duration, in seconds, to wait for a server to respond when first connecting and on follow-up operations in the session.

Default: `global timeout`
 Range: 1-30

Options

Parameter	Description
-V --version	Print the plugin version and license information.
-h --help	Print help information (shorthand).
-t, --test <item key>	Launch plugin for testing (plugin config ignored).

See also

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

MySQL plugin

Overview

Configuration parameters allow customizing the MySQL plugin behavior.

It is recommended to specify them in their own configuration file (e.g., `mysql.conf`), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.Mysql.CallTimeout	Specify the maximum amount of time in seconds to wait for a request to be done.
Plugins.Mysql.CustomQueriesPath	Specify full pathname of the directory containing `.sql` files with custom queries.
Plugins.Mysql.CustomQueriesEnabled	Enable execution of the `mysql.custom.query` item key.
Plugins.Mysql.Default.Password	Specify the default password.
Plugins.Mysql.Default.TLSConnect	Specify the default level of encryption for communications between Zabbix agent 2 and monitored databases.
Plugins.Mysql.Default.TLSCAFile	Specify full pathname of the file containing top-level CA(s) certificates for peer certificate verification.
Plugins.Mysql.Default.TLSCertFile	Specify full pathname of the file containing the agent certificate or certificate chain.
Plugins.Mysql.Default.TLSKeyFile	Specify full pathname of the file containing the database private key.
Plugins.Mysql.Default.Uri	Specify the default URI.
Plugins.Mysql.Default.User	Specify the default username.
Plugins.Mysql.KeepAlive	Specify the maximum time of waiting (in seconds) before unused plugin connections are closed.
Plugins.Mysql.Sessions.<SessionName>.Password	Specify the named session password.
Plugins.Mysql.Sessions.<SessionName>.TLSConnect	Specify the named session level of encryption for communications between Zabbix agent 2 and monitored databases.
Plugins.Mysql.Sessions.<SessionName>.TLSCAFile	Specify full pathname of the file containing top-level CA(s) certificates for named session peer certificate verification.
Plugins.Mysql.Sessions.<SessionName>.TLSCertFile	Specify full pathname of the file containing the named session agent certificate or certificate chain.
Plugins.Mysql.Sessions.<SessionName>.TLSKeyFile	Specify full pathname of the file containing the name dsession database private key.
Plugins.Mysql.Sessions.<SessionName>.Uri	Specify the named session URI.
Plugins.Mysql.Sessions.<SessionName>.User	Specify the named session username.
Plugins.Mysql.Timeout	Set the maximum time in seconds for waiting when a connection has to be established.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files;
- Values support environment variables;
- Zabbix supports configuration files only in UTF-8 encoding without BOM;
- Comments starting with ”#” are only supported at the beginning of the line.

Parameter details

Plugins.Mysql.CallTimeout

Specify the maximum amount of time in seconds to wait for a request to be done.

Default: `global timeout`
 Range: 1-30

Plugins.Mysql.CustomQueriesPath

Specify full pathname of the directory containing `.sql` files with custom queries.

Defaults:
`/usr/local/share/zabbix/custom-queries/mysql` - for Unix systems,
`*:\Program Files\Zabbix Agent 2\Custom Queries\Mysql` - for Windows systems (where `*` is the drive name from the `ProgramFiles` environment variable)

Plugins.Mysql.CustomQueriesEnabled

Enable execution of the `mysql.custom.query` item key.
 If disabled, no queries will be loaded from the custom query directory path.

Default: `false`
 Range: `true/false`

Plugins.Mysql.Default.Password

Specify the default password for a protected MySQL server.

Plugins.Mysql.Default.TLSConnect

Specify the default level of encryption for communications between Zabbix agent 2 and monitored databases.

Range:
required - require TLS connection;
verify_ca - verify certificates;
verify_full - verify certificates and IP address

Plugins.Mysql.Default.TLSCAFile

Specify full pathname of the file containing top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix agent 2 and monitored databases.

Mandatory: `yes`, if `Plugins.Mysql.Default.TLSConnect` is set to `verify_ca` or `verify_full`

Plugins.Mysql.Default.TLSCertFile

Specify full pathname of the file containing the agent certificate or certificate chain, used for encrypted communications between Zabbix agent 2 and monitored databases.

Mandatory: `yes`, if `Plugins.Mysql.Default.TLSKeyFile` is specified

Plugins.Mysql.Default.TLSKeyFile

Specify full pathname of the file containing the database private key, used for encrypted communications between Zabbix agent 2 and monitored databases.

Mandatory: `yes`, if `Plugins.Mysql.Default.TLSCertFile` is specified

Plugins.Mysql.Default.Uri

Specify the default URI for connecting to MySQL server.
 Embedded credentials will be ignored.
 Supported schemes: `tcp`, `unix`
 The schema can be omitted.
 The port can be omitted (default=3306).

Examples:

```
tcp://localhost:3306
localhost
unix:/var/run/mysql.sock
```

Plugins.Mysql.Default.User

Specify the default username for a protected MySQL server.

Plugins.Mysql.KeepAlive

Specify the maximum time of waiting (in seconds) before unused plugin connections are closed.

Default: 300
 Range: 60−900

Plugins.Mysql.Sessions.<SessionName>.Password

Specify the named session password for a protected MySQL server.
 <SessionName> - define name of a session for using in item keys.

Plugins.Mysql.Sessions.<SessionName>.TLSConnect

Specify the level of named-session encryption for communications between Zabbix agent 2 and monitored databases.
 <SessionName> - define name of a session for using in item keys.

Range:
required - require TLS connection;
verify_ca - verify certificates;
verify_full - verify certificates and IP address

Plugins.Mysql.Sessions.<SessionName>.TLSCAFile

Specify full pathname of the file containing top-level CA(s) certificates for named-session peer certificate verification, used for encrypted communications between Zabbix agent 2 and monitored databases.
 <SessionName> - define name of a session for using in item keys.

Mandatory: `yes`, if `Plugins.Mysql.Sessions.<SessionName>.TLSConnect` is set to `verify_ca` or `verify_full`

Plugins.Mysql.Sessions.<SessionName>.TLSCertFile

Specify full pathname of the file containing the named-session agent certificate or certificate chain, used for encrypted communications between Zabbix agent 2 and monitored databases.
 <SessionName> - define name of a session for using in item keys.

Mandatory: yes, if Plugins.Mysql.Sessions.<SessionName>.TLSKeyFile is specified

Plugins.Mysql.Sessions.<SessionName>.TLSKeyFile

Specify full pathname of the file containing the named-session database private key, used for encrypted communications between Zabbix agent 2 and monitored databases.
 <SessionName> - define name of a session for using in item keys.

Mandatory: yes, if Plugins.Mysql.Sessions.<SessionName>.TLSCertFile is specified

Plugins.Mysql.Sessions.<SessionName>.Uri

Specify the named session URI for connecting to MySQL server.
 <SessionName> - define name of a session for using in item keys.
 Embedded credentials will be ignored.
 Supported schemes: tcp, unix
 The schema can be omitted.
 The port can be omitted (default=3306).

Examples:

```
tcp://localhost:3306
localhost
unix:/var/run/mysql.sock
```

Plugins.Mysql.Sessions.<SessionName>.User

Specify the named session username for a protected MySQL server.
 <SessionName> - define name of a session for using in item keys.

Plugins.Mysql.Timeout

Set the maximum time in seconds for waiting when a connection has to be established.

Default: global timeout
 Range: 1–30

See also

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

NVIDIA GPU plugin

Overview

Configuration parameters allow customizing the NVIDIA GPU plugin behavior.

It is recommended to specify them in their own configuration file (e.g., nvidiagpu.conf), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

The NVIDIA GPU plugin is a loadable plugin and is available and fully described in the NVIDIA GPU plugin repository.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.NVIDIA.System.Path	Specify path to the NVIDIA GPU plugin executable.
Plugins.NVIDIA.Timeout	Specify the maximum time (in seconds) to wait for a server response during connection attempts and subsequent operations in the session.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files;
- Values support environment variables;
- Zabbix supports configuration files only in UTF-8 encoding without BOM;
- Comments starting with ”#” are only supported at the beginning of the line.

Parameter details

Plugins.NVIDIA.System.Path

Specify path to the NVIDIA GPU plugin executable.

Example:

```
Plugins.NVIDIA.System.Path=/path/to/executable/nvidia
```

Plugins.NVIDIA.Timeout

Specify the maximum time (in seconds) to wait for a server response during connection attempts and subsequent operations in the session.
 The global item-type timeout or individual item timeout will override this value if greater.

Default: `global timeout`
 Range: 1–30

See also

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

Oracle plugin

Overview

Configuration parameters allow customizing the Oracle plugin behavior.

It is recommended to specify them in their own configuration file (e.g., `oracle.conf`), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.Oracle.CallTimeout	The maximum wait time in seconds for a request to be completed. This parameter currently has no effect; the item timeout setting is used instead.
Plugins.Oracle.ConnectTimeout	Specify the maximum wait time in seconds for a connection to be established.
Plugins.Oracle.CustomQueriesPath	Specify full pathname of the directory containing `.sql` files with custom queries.
Plugins.Oracle.CustomQueriesEnabled	Enable execution of the `oracle.custom.query` item key.
Plugins.Oracle.Default.Password	Specify the default password.
Plugins.Oracle.Default.Service	Specify the default service name.
Plugins.Oracle.Default.URI	Specify the default URI.
Plugins.Oracle.Default.User	Specify the default username.
Plugins.Oracle.KeepAlive	Set the maximum time of waiting (in seconds) before unused plugin connections are closed.
Plugins.Oracle.ResolveTNS	Specify how to interpret the connection string (ConnString) for connecting to the Oracle server.
Plugins.Oracle.Sessions.<SessionName>.Password	Specify the named session password.
Plugins.Oracle.Sessions.<SessionName>.Service	Specify the named session service name.
Plugins.Oracle.Sessions.<SessionName>.URI	Specify the named session URI.
Plugins.Oracle.Sessions.<SessionName>.User	Specify the named session username.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files.
- Values support environment variables.
- Zabbix supports configuration files only in UTF-8 encoding without BOM.
- Comments starting with ”#” are only supported at the beginning of the line.

Parameter details

Plugins.Oracle.CallTimeout

The maximum wait time in seconds for a request to be completed.
 This parameter currently has no effect. The maximum wait time is determined by the item timeout instead.

Default: `Timeout`[1]
 Range: 1–30

Plugins.Oracle.ConnectTimeout

Specify the maximum wait time in seconds for a connection to be established.

Default: `Timeout`[1]
 Range: 1–30

Plugins.Oracle.CustomQueriesPath

Specify full pathname of the directory containing `.sql` files with custom queries.

Defaults:
`/usr/local/share/zabbix/custom-queries/oracle` - for Unix systems,
`*:\Program Files\Zabbix Agent 2\Custom Queries\Oracle` - for Windows systems (where `*` is the drive name from the `ProgramFiles` environment variable)

Example:

```
/etc/zabbix/oracle/sql
```

Plugins.Oracle.CustomQueriesEnabled

Enable execution of the `oracle.custom.query` item key.
 If disabled, no queries will be loaded from the custom query directory path.

Default: `false`
 Range: `true/false`

Plugins.Oracle.Default.Password

Specify the default password for connecting to Oracle.
 Used if no value is specified in an item key or named session.

Plugins.Oracle.Default.Service

Specify the default service name for connecting to Oracle (SID is not supported).
 Used if no value is specified in an item key or named session.

Plugins.Oracle.Default.Uri

Specify the default URI for connecting to Oracle.
 Used if no value is specified in an item key or named session.
 Embedded credentials will be ignored.
 Must match the URI format.
 Supported schemes: `tcp`, `tcps`
 The schema can be omitted.
 The port can be omitted (default=1521).
 It is also possible to specify the TNS key or TNS value as the connection string. The TNS value must be composed without whitespaces.

Examples:

```
tcp://127.0.0.1:1521
localhost
zbx_tns_example (TNS key)
(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=localhost)(PORT=1521))(CONNECT_DATA=(SERVICE_NAME=xe))) (TNS val
```

Plugins.Oracle.Default.User

Specify the default username for connecting to Oracle.
 Used if no value is specified in an item key or named session.

Plugins.Oracle.KeepAlive

Set the maximum time of waiting (in seconds) before unused plugin connections are closed.

Default: 300
 Range: 60–900

Plugins.Oracle.ResolveTNS

Specify how to interpret the connection string (ConnString) for connecting to the Oracle server.
 If it is set to `false` the default scheme and port will be added to ConnString (if absent).
 If the option is set to `true`, the default scheme and port will be omitted (unless explicitly specified in ConnString), and ConnString will be passed to the Oracle client as is. If the Oracle client finds ConnString in the `tnsnames.ora` file, the connection description found will be used to connect to the Oracle server.

Default: `true`
 Range: `true/false`

Plugins.Oracle.Sessions.<SessionName>.Password

Specify the named session password for connecting to Oracle.
 <SessionName> - define name of a session for using in item keys.

Plugins.Oracle.Sessions.<SessionName>.Service

Specify the named session service name for connecting to Oracle (SID is not supported).
 <SessionName> - define name of a session for using in item keys.

Plugins.Oracle.Sessions.<SessionName>.Uri

Specify the named session URI for connecting to Oracle.
 <SessionName> - define name of a session for using in item keys.
 Embedded credentials will be ignored.
 Must match the URI format.
 Supported schemes: `tcp`, `tcps`
 The schema can be omitted.
 The port can be omitted (default=1521).
 It is also possible to specify the TNS key or TNS value as the connection string. The TNS value must be composed without whitespaces.

Examples:

```
tcp://127.0.0.1:1521
localhost
zbx_tns_example (TNS key)
(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=localhost)(PORT=1521))(CONNECT_DATA=(SERVICE_NAME=xe))) (TNS val
```

Plugins.Oracle.Sessions.<SessionName>.User

Specify the named session username for connecting to Oracle.
 <SessionName> - define name of a session for using in item keys.

Footnotes

[1] The timeout is defined by the `Timeout` option in the Zabbix agent 2 (UNIX) or Zabbix agent 2 (Windows) configuration file (`zabbix_agent2.conf`).

See also

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

PostgreSQL plugin

Overview

Configuration parameters allow customizing the PostgreSQL plugin behavior.

It is recommended to specify them in their own configuration file (e.g., `postgresql.conf`), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

The PostgreSQL plugin is a loadable plugin and is available and fully described in the PostgreSQL plugin repository.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.PostgreSQL.CallTimeout	Set the maximum wait time in seconds for a request to be completed.
Plugins.PostgreSQL.CustomQueriesPath	Specify full pathname of the directory containing `.sql` files with custom queries.
Plugins.PostgreSQL.CustomQueriesEnabled	Enable execution of the `postgresql.custom.query` item key.
Plugins.PostgreSQL.Default.CacheMode	Specify the default cache mode for the PostgreSQL connection.
Plugins.PostgreSQL.Default.Database	Specify the default database name.
Plugins.PostgreSQL.Default.Password	Specify the default password.
Plugins.PostgreSQL.Default.TLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificate for peer certificate verification.
Plugins.PostgreSQL.Default.TLSCertFile	Specify full pathname of the file containing the PostgreSQL certificate or certificate chain.
Plugins.PostgreSQL.Default.TLSConnect	Specify the default level of encryption for communications between Zabbix agent 2 and monitored databases.
Plugins.PostgreSQL.Default.TLSKeyFile	Specify full pathname of the file containing the PostgreSQL private key.
Plugins.PostgreSQL.Default.Uri	Specify the default URI.
Plugins.PostgreSQL.Default.User	Specify the default username.
Plugins.PostgreSQL.KeepAlive	Set the maximum time of waiting (in seconds) before unused plugin connections are closed.
Plugins.PostgreSQL.Sessions.<SessionName>.CacheMode	Specify the named session cache mode for the PostgreSQL connection.
Plugins.PostgreSQL.Sessions.<SessionName>.Database	Specify the named session database name.
Plugins.PostgreSQL.Sessions.<SessionName>.Password	Specify the named session password.
Plugins.PostgreSQL.Sessions.<SessionName>.TLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificate for peer certificate verification.
Plugins.PostgreSQL.Sessions.<SessionName>.TLSCertFile	Specify full pathname of the file containing the PostgreSQL certificate or certificate chain.
Plugins.PostgreSQL.Sessions.<SessionName>.TLSConnect	Specify the named session level of encryption for communications between Zabbix agent 2 and monitored databases.
Plugins.PostgreSQL.Sessions.<SessionName>.TLSKeyFile	Specify full pathname of the file containing the PostgreSQL private key.
Plugins.PostgreSQL.Sessions.<SessionName>.Uri	Specify the named session URI.
Plugins.PostgreSQL.Sessions.<SessionName>.User	Specify the named session username.

|Plugins.PostgreSQL.System.Path|Specify path to the PostgreSQL plugin executable.| |Plugins.PostgreSQL.Timeout|Set the request execution timeout.|

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files.
- Values support environment variables;
- Zabbix supports configuration files only in UTF-8 encoding without BOM.
- Comments starting with "#" are only supported at the beginning of the line.

Parameter details

Plugins.PostgreSQL.CallTimeout

Set the maximum wait time in seconds for a request to be completed.

Default: `global timeout`
 Range: 1-30

Plugins.PostgreSQL.CustomQueriesPath

Specify full pathname of the directory containing `.sql` files with custom queries.

Defaults:
`/usr/local/share/zabbix/custom-queries/postgreSQL` - for Unix systems,
`*:\Program Files\Zabbix Agent 2\Custom Queries\PostgreSQL` - for Windows systems (where `*` is the drive name from the `ProgramFiles` environment variable)

Example:

```
/etc/zabbix/oracle/sql
```

Plugins.PostgreSQL.CustomQueriesEnabled

Enable execution of the `postgresql.custom.query` item key.
 If disabled, no queries will be loaded from the custom query directory path.

Default: `false`
 Range: `true/false`

Plugins.PostgreSQL.Default.CacheMode

Specify the default cache mode for the PostgreSQL connection.

Default: `prepare`
 Range:
`prepare` - create prepared statements on the PostgreSQL server;
`describe` - use the anonymous prepared statement to describe a statement without creating a statement on the server.

Note that `describe` is primarily useful when the environment does not allow prepared statements such as when running a connection pooler like PgBouncer.

Plugins.PostgreSQL.Default.Database

Specify the default database for connecting to PostgreSQL.
 Used if no value is specified in an item key or named session.

Plugins.PostgreSQL.Default.Password

Specify the default password for connecting to PostgreSQL.
 Used if no value is specified in an item key or named session.

Plugins.PostgreSQL.Default.TLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificate for peer certificate verification for encrypted communications between Zabbix agent 2 and monitored databases.
 Used if no value is specified in a named session.

Mandatory: `yes`, if `Plugins.PostgreSQL.Default.TLSConnect` is set to `verify_ca` or `verify_full`

Plugins.PostgreSQL.Default.TLSCertFile

Specify full pathname of the file containing the PostgreSQL certificate or certificate chain for encrypted communications between Zabbix agent 2 and monitored databases.
 Used if no value is specified in a named session.

Mandatory: `yes`, if `Plugins.PostgreSQL.Default.TLSKeyFile` is specified

Plugins.PostgreSQL.Default.TLSConnect

Specify the default level of encryption for communications between Zabbix agent 2 and monitored databases.
 Used if no value is specified in a named session.

Range:
`required` - connect using TLS as transport mode without identity checks;
`verify_ca` - connect using TLS and verify certificate;
`verify_full` - connect using TLS, verify certificate and verify that database identity (CN) specified by DBHost matches its certificate.

Undefined encryption type means unencrypted connection.

Plugins.PostgreSQL.Default.TLSKeyFile

Specify full pathname of the file containing the PostgreSQL private key for encrypted communications between Zabbix agent 2 and monitored databases.
 Used if no value is specified in a named session.

Mandatory: `yes`, if `Plugins.PostgreSQL.Default.TLSCertFile` is specified

Plugins.PostgreSQL.Default.Uri

Specify the default URI for connecting to PostgreSQL.
 Used if no value is specified in an item key or named session.
 Embedded credentials will be ignored.
 Must match the URI format.
 Supported schemes: `tcp`, `unix`

Examples:

```
tcp://127.0.0.1:5432
tcp://localhost
unix:/var/run/postgresql/.s.PGSQL.5432
```

Plugins.PostgreSQL.Default.User

Specify the default username for connecting to PostgreSQL.
 Used if no value is specified in an item key or named session.

Plugins.PostgreSQL.KeepAlive

Set the maximum time of waiting (in seconds) before unused plugin connections are closed.

Plugins.PostgreSQL.Sessions.<SessionName>.CacheMode

Specify the named session cache mode for the PostgreSQL connection.
 <SessionName> - define name of a session for using in item keys.

Default: `prepare`
 Range:
`prepare` - create prepared statements on the PostgreSQL server;
`describe` - use the anonymous prepared statement to describe a statement without creating a statement on the server.

Note that `describe` is primarily useful when the environment does not allow prepared statements such as when running a connection pooler like PgBouncer.

Plugins.PostgreSQL.Sessions.<SessionName>.Database

Specify the named session database for connecting to PostgreSQL.
 <SessionName> - define name of a session for using in item keys.

Plugins.PostgreSQL.Sessions.<SessionName>.Password

Specify the named session password for connecting to PostgreSQL.
 <SessionName> - define name of a session for using in item keys.
 Must match the password format.

Plugins.PostgreSQL.Sessions.<SessionName>.TLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificate for peer certificate verification for encrypted communications between Zabbix agent 2 and monitored databases.
 <SessionName> - define name of a session for using in item keys.

Mandatory: `yes`, if `Plugins.PostgreSQL.Sessions.<SessionName>.TLSConnect` is set to `verify_ca` or `verify_full`

Plugins.PostgreSQL.Sessions.<SessionName>.TLSCertFile

Specify full pathname of the file containing the PostgreSQL certificate or certificate chain for encrypted communications between Zabbix agent 2 and monitored databases.
 <SessionName> - define name of a session for using in item keys.

Mandatory: `yes`, if `Plugins.PostgreSQL.Sessions.<SessionName>.TLSKeyFile` is specified

Plugins.PostgreSQL.Sessions.<SessionName>.TLSConnect

Specify the named session level of encryption for communications between Zabbix agent 2 and monitored databases.
 <SessionName> - define name of a session for using in item keys.

Range:
`required` - connect using TLS as transport mode without identity checks;
`verify_ca` - connect using TLS and verify certificate;
`verify_full` - connect using TLS, verify certificate and verify that database identity (CN) specified by DBHost matches its certificate.

Undefined encryption type means unencrypted connection.

Plugins.PostgreSQL.Sessions.<SessionName>.TLSKeyFile

Specify full pathname of the file containing the PostgreSQL private key for encrypted communications between Zabbix agent 2 and monitored databases.
 <SessionName> - define name of a session for using in item keys.

Mandatory: yes, if Plugins.PostgreSQL.Sessions.<SessionName>.TLSCertFile is specified.

Plugins.PostgreSQL.Sessions.<SessionName>.Uri

Specify the named session URI for connecting to PostgreSQL.
 <SessionName> - define name of a session for using in item keys.
 Embedded credentials will be ignored.
 Must match the URI format.
 Supported schemes: tcp, unix

Examples:

```
tcp://127.0.0.1:5432
tcp://localhost
unix:/var/run/postgresql/.s.PGSQL.5432
```

Plugins.PostgreSQL.Sessions.<SessionName>.User

Specify the named session username for connecting to PostgreSQL.
 <SessionName> - define name of a session for using in item keys.

Plugins.PostgreSQL.System.Path

Specify path to the PostgreSQL plugin executable.

Mandatory: yes

Example:

```
Plugins.PostgreSQL.System.Path=/usr/libexec/zabbix/zabbix-agent2-plugin-postgresql
```

Plugins.PostgreSQL.Timeout

Set the request execution timeout (the duration, in seconds, to wait for a request to complete before shutting it down).

Default: global timeout
 Range: 1-30

Options

Parameter	Description
-V --version	Print the plugin version and license information.
-h --help	Print help information (shorthand).
-t, --test <item key>	Launch plugin for testing (plugin config ignored).

See also

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

Redis plugin

Overview

Configuration parameters allow customizing the Redis plugin behavior.

It is recommended to specify them in their own configuration file (e.g., redis.conf), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.Redis.Default.Password	Specify the default password for connecting to Redis.
Plugins.Redis.Default.TLSConnect	Specify the level of encryption required for Zabbix agent 2 communications.
Plugins.Redis.Default.TLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificates.
Plugins.Redis.Default.TLSCertFile	Specify full pathname of the file containing the agent certificate.
Plugins.Redis.Default.TLSKeyFile	Specify full pathname of the file containing the agent private key.
Plugins.Redis.Default.Uri	Specify the default URI for connecting to Redis.
Plugins.Redis.Default.User	Specify the default username for connecting to Redis.

Parameter	Description
Plugins.Redis.KeepAlive	Set the maximum time of waiting (in seconds) before unused plugin connections are closed.
Plugins.Redis.Sessions.<SessionName>.Password	Specify the named session password for connecting to Redis.
Plugins.Redis.Sessions.<SessionName>.TLSConnect	Specify the level of encryption required for Zabbix agent 2 communications for named session.
Plugins.Redis.Sessions.<SessionName>.TLSCAFile	Specify full pathname of the file containing the top-level CA(s) certificates for named session.
Plugins.Redis.Sessions.<SessionName>.TLSCertFile	Specify full pathname of the file containing the agent certificate for named session.
Plugins.Redis.Sessions.<SessionName>.TLSKeyFile	Specify full pathname of the file containing the agent private key for named session.
Plugins.Redis.Sessions.<SessionName>.Uri	Specify the named session URI for connecting to Redis.
Plugins.Redis.Sessions.<SessionName>.User	Specify the named session username for connecting to Redis.
Plugins.Redis.Timeout	Set the request execution timeout.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files.
- Values support environment variables.
- Zabbix supports configuration files only in UTF-8 encoding without BOM.
- Comments starting with "#" are only supported at the beginning of the line.

Parameter details

Plugins.Redis.Default.Password

Specify the default password for connecting to Redis.
 Used if no value is specified in an item key or named session.

Plugins.Redis.Default.TLSConnect

Specify the level of encryption required for Zabbix agent 2 communications.
 An empty or unset value is treated as no TLS.

Range: `required` - encrypted, unverified (only for testing), `verify_ca` - encrypted, server certificate, `verify_full` - encrypted, server certificate and server name verified with SAN

Plugins.Redis.Default.TLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificates.

Plugins.Redis.Default.TLSCertFile

Specify full pathname of the file containing the agent certificate.

Plugins.Redis.Default.TLSKeyFile

Specify full pathname of the file containing the agent private key.

Plugins.Redis.Default.Uri

Specify the default URI for connecting to Redis.
 Used if no value is specified in an item key or named session.
 Embedded credentials will be ignored. Must match the URI format.
 Supported schemes: `tcp`, `unix`
 A scheme can be omitted.
 A port can be omitted (default=6379).

Default: `tcp://localhost:6379`

Examples:

```
tcp://localhost:6379
localhost
unix:/var/run/redis.sock
```

Plugins.Redis.Default.User

Specify the default username for connecting to Redis.
 Used if no value is specified in an item key or named session.

Default: `default`

Plugins.Redis.KeepAlive

Set the maximum time of waiting (in seconds) before unused plugin connections are closed.

Default: 300
 Range: 60–900

Plugins.Redis.Sessions.<SessionName>.Password

Specify the named session password for connecting to Redis.
 <SessionName> - define the session name to be used in item keys.

Plugins.Redis.Sessions.<SessionName>.TLSConnect

Specify the level of encryption required for Zabbix agent 2 communications for named session.
 <SessionName> - define the session name to be used in item keys.
 An empty or unset value is treated as no TLS.

Range: `required` - encrypted, unverified (only for testing), `verify_ca` - encrypted, server certificate, `verify_full` - encrypted, server certificate and server name verified with SAN

|Plugins.Redis.Sessions.<SessionName>.TLSCAFile|no
(yes, if TLSConnect is `verify_ca` or `verify_full`)| | |Full pathname of a file containing the top-level CA(s) certificates for Redis server certificate verification.
<SessionName> - define the session name to be used in item keys.| |Plugins.Redis.Sessions.<SessionName>.TLSCertFile|no| | |Full pathname of a file containing the agent's certificate for client authentication.
<SessionName> - define the session name to be used in item keys.| |Plugins.Redis.Sessions.<SessionName>.TLSKeyFile|no| | |Full pathname of a file containing the agent's private key for client authentication.
<SessionName> - define the session name to be used in item keys.|

Plugins.Redis.Sessions.<SessionName>.TLSCAFile

Specify full pathname of the file containing the top-level CA(s) certificates for named session.
 <SessionName> - define the session name to be used in item keys.

Mandatory: `yes`, if `Plugins.Redis.Sessions.<SessionName>.TLSConnect` is set to `verify_ca` or `verify_full`

Plugins.Redis.Sessions.<SessionName>.TLSCertFile

Specify full pathname of the file containing the agent certificate for named session.
 <SessionName> - define the session name to be used in item keys.

Plugins.Redis.Sessions.<SessionName>.TLSKeyFile

Specify full pathname of the file containing the agent private key for named session.
 <SessionName> - define the session name to be used in item keys.

Plugins.Redis.Sessions.<SessionName>.Uri

Specify the named session URI for connecting to Redis.
 <SessionName> - define the session name to be used in item keys.
 Embedded credentials will be ignored. Must match the URI format.
 Supported schemes: `tcp`, `unix`
 A scheme can be omitted.
 A port can be omitted (default=6379).

Default: `tcp://localhost:6379`

Examples:

```
tcp://localhost:6379
localhost
unix:/var/run/redis.sock
```

Plugins.Redis.Sessions.<SessionName>.User

Specify the named session username for connecting to Redis.
 <SessionName> - define the session name to be used in item keys.

Plugins.Redis.Timeout

Set the request execution timeout (the duration, in seconds, to wait for a request to complete before shutting it down).

Default: `global timeout`
 Range: 1–30

See also

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

SMART plugin

Overview

Configuration parameters allow customizing the SMART plugin behavior.

It is recommended to specify them in their own configuration file (e.g., `smart.conf`), rather than in the agent 2 configuration file. Then use the Include directive to include this file in the Zabbix agent 2 configuration.

All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
Plugins.Smart.Path	Specify path to the smartctl executable.
Plugins.Smart.Timeout	Set the request execution timeout (the duration, in seconds, to wait for a request to complete before shutting it down).

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files;
- Values support environment variables;
- The path to the `smartctl` executable must be set either by adding its directory to the system's `PATH` environment variable or by configuring `Plugins.Smart.Path`; this applies to both Linux and Windows;
- Zabbix supports configuration files only in UTF-8 encoding without BOM;
- Comments starting with "`#`" are only supported at the beginning of the line.

Parameter details

Plugins.Smart.Path

Specify path to the smartctl executable.

Default: `smartctl`

Plugins.Smart.Timeout

Set the request execution timeout (the duration, in seconds, to wait for a request to complete before shutting it down).

Default: `global timeout<br>` Range: 1–30

See also

- Description of general Zabbix agent 2 configuration parameters: Zabbix agent 2 (UNIX) / Zabbix agent 2 (Windows)
- Instructions for configuring plugins

Zabbix agent 2 plugin protocol

Zabbix agent 2 protocol is based on code, size and data model.

Code

Type	Size	Comments
Byte	4	Payload type, currently only JSON is supported.

Size

Type	Size	Comments
Byte	4	Size of the current payload in bytes.

Payload data

Type	Size	Comments
Byte	Defined by the *Size* field	JSON formatted data.

Payload data definition

Common data

These parameters are present in all requests/responses:

Name	Type	Comments
id	uint32	For requests - the incrementing identifier used to link requests with responses. Unique within a request direction (i.e. from agent to plugin or from plugin to agent). For responses - ID of the corresponding request.
type	uint32	The request type.

Log request

A request sent by a plugin to write a log message into the agent log file.

direction	plugin → agent
response	no

Parameters specific to log requests:

Name	Type	Comments
severity	uint32	The message severity (log level).
message	string	The message to log.

Example:

```
{"id":0,"type":1,"severity":3,"message":"message"}
```

Register request

A request sent by the agent during the agent startup phase to obtain provided metrics to register a plugin.

direction	agent → plugin
response	yes

Parameters specific to register requests:

Name	Type	Comments
version	string	The protocol version <major>.<minor>

Example:

```
{"id":1,"type":2,"version":"1.0"}
```

Register response

Plugin’s response to the register request.

direction	plugin → agent
response	n/a

Parameters specific to register responses:

Name	Type	Comments
name	string	The plugin name.
metrics	array of strings (optional)	The metrics with descriptions as used in the plugin. Returns RegisterMetrics(). Absent if error is returned.
interfaces	uint32 (optional)	The bit mask of plugin’s supported interfaces. Absent if error is returned.

Name	Type	Comments
error	string (optional)	An error message returned if a plugin cannot be started. Absent, if metrics are returned.

Examples:

```
{"id":2,"type":3,"metrics":["external.test", "External exporter Test."], "interfaces": 4}
```

or

```
{"id":2,"type":3,"error":"error message"}
```

Start request

A request to execute the Start function of the Runner interface.

direction	agent → plugin
response	no

The request doesn't have specific parameters, it only contains common data parameters.

Example:

```
{"id":3,"type":4}
```

Terminate request

A request sent by the agent to shutdown a plugin.

direction	agent → plugin
response	no

The request doesn't have specific parameters, it only contains common data parameters.

Example:

```
{"id":3,"type":5}
```

Export request

A request to execute the Export function of the Exporter interface.

direction	agent → plugin
response	no

Parameters specific to export requests:

Name	Type	Comments
key	string	The plugin key.
parameters	array of strings (optional)	The parameters for Export function.

Example:

```
{"id":4,"type":6,"key":"test.key","parameters":["foo","bar"]}
```

Export response

Response from the Export function of the Exporter interface.

direction	plugin → agent
response	n/a

Parameters specific to export responses:

Name	Type	Comments
value	string (optional)	Response value from the Export function. Absent, if error is returned.
error	string (optional)	Error message if the Export function has not been executed successfully. Absent, if value is returned.

Examples:

```
{"id":5,"type":7,"value":"response"}
```

or

```
{"id":5,"type":7,"error":"error message"}
```

Configure request

A request to execute the *Configure* function of the *Configurator* interface.

direction	agent → plugin
response	n/a

Parameters specific to *Configure* requests:

Name	Type	Comments
global_options	JSON object	JSON object containing global agent configuration options.
private_options	JSON object (optional)	JSON object containing private plugin configuration options, if provided.

Example:

```
{"id":6,"type":8,"global_options":{...},"private_options":{...}}
```

Validate request

A request to execute *Validate* function of the *Configurator* interface.

direction	agent → plugin
response	yes

Parameters specific to *Validate* requests:

Name	Type	Comments
private_options	JSON object (optional)	JSON object containing private plugin configuration options, if provided.

Example:

```
{"id":7,"type":9,"private_options":{...}}
```

Validate response

Response from *Validate* function of *Configurator* interface.

direction	plugin → agent
response	n/a

Parameters specific to *Validate* responses:

Name	Type	Comments
error	string (optional)	An error message returned if the Validate function is not executed successfully. Absent if executed successfully.

Example:

```
{"id":8,"type":10}
```

or

```
{"id":8,"type":10,"error":"error message"}
```

Java gateway

Overview Zabbix Java gateway can be installed from source code or packages.

Native support for monitoring JMX applications exists in the form of a Zabbix daemon called "Zabbix Java gateway". Zabbix Java gateway is a daemon written in Java. To find out the value of a particular JMX counter on a host, Zabbix server queries Zabbix Java gateway, which uses the JMX management API to query the application of interest remotely. The application does not need any additional software installed, it just has to be started with `-Dcom.sun.management.jmxremote` option on the command line.

Java gateway accepts incoming connection from Zabbix server or proxy and can only be used as a "passive proxy". As opposed to Zabbix proxy, it may also be used from Zabbix proxy (Zabbix proxies cannot be chained). Access to each Java gateway is configured directly in the Zabbix server or Zabbix proxy configuration file, thus only one Java gateway may be configured per Zabbix server or Zabbix proxy. If a host will have items of type *JMX agent* and items of other type, only the *JMX agent* items will be passed to Java gateway for retrieval.

When an item has to be updated over Java gateway, Zabbix server or proxy will connect to the Java gateway and request the value, which Java gateway in turn retrieves and passes back to the server or proxy. As such, Java gateway does not cache any values.

Zabbix server or proxy has a specific type of processes that connect to Java gateway, controlled by the option `StartJavaPollers`. Internally, Java gateway starts multiple threads, controlled by the `START_POLLERS` option. On the server side, if a connection takes more than "Timeout" seconds, it will be terminated, but Java gateway might still be busy retrieving value from the JMX counter. To solve this, the Java gateway provides the `TIMEOUT` option, which allows you to set a timeout for JMX network operations.

See also:

- Configuration parameters: Zabbix Java gateway
- Setup from Debian/Ubuntu packages
- Setup from RHEL packages
- Setup from sources

Zabbix server or proxy will try to pool requests to a single JMX target together as much as possible (affected by item intervals) and send them to the Java gateway in a single connection for better performance.

It is recommended to have `StartJavaPollers` less than or equal to `START_POLLERS`; otherwise, there might be situations when no threads are available in the Java gateway to service incoming requests. In such cases, the Java gateway uses `ThreadPoolExecutor.CallerRunsPolicy`, meaning that the main thread will service the incoming request and temporarily will not accept any new requests.

If you are trying to monitor Wildfly-based Java applications with Zabbix Java gateway, please install the latest `jboss-client.jar` available on the Wildfly download page.

Configuration parameters: Zabbix Java gateway

If you use `startup.sh` and `shutdown.sh` scripts for starting Zabbix Java gateway, then you can specify the necessary configuration parameters in the `settings.sh` file. The startup and shutdown scripts source the settings file and take care of converting shell variables (listed in the first column) to Java properties (listed in the second column).

If you start Zabbix Java gateway manually by running `java` directly, then you specify the corresponding Java properties on the command line.

Variable	Property	Mandatory	Range	Default	Description
LISTEN_IP	zabbix.listenIP	no		0.0.0.0	IP address to listen on.
LISTEN_PORT	zabbix.listenPort	no	1024-32767	10052	Port to listen on.
PID_FILE	zabbix.pidFile	no		/tmp/zabbix_java.pid	Name of PID file. If omitted, Zabbix Java Gateway is started as a console application.
PROPERTIES_FILE	zabbix.propertiesFile	no			Name of properties file. Can be used to set additional properties using a key-value format in such a way that they are not visible on a command line or to overwrite existing ones. For example: ”javax.net.ssl.trustStorePassword=<p
START_POLLERS	zabbix.startPollers	no	1-1000	5	Number of worker threads to start.
TIMEOUT	zabbix.timeout	no	1-30	3	How long to wait for network operations (in seconds).

Warning:
Port 10052 is not IANA registered.

Setup from Debian/Ubuntu packages

Overview

If installed from packages, the following information will help you in setting up Zabbix Java gateway.

Configuring and running Java gateway

Java gateway configuration may be tuned in the file:

```
/etc/zabbix/zabbix_java_gateway.conf
```

For more details, see Zabbix Java gateway configuration parameters.

To start Zabbix Java gateway:

```
systemctl restart zabbix-java-gateway
```

To automatically start Zabbix Java gateway on boot:

```
systemctl enable zabbix-java-gateway
```

Configuring server for use with Java gateway

With Java gateway up and running, you have to tell Zabbix server where to find Zabbix Java gateway. This is done by specifying `JavaGateway` and `JavaGatewayPort` parameters in the server configuration file. If the host on which JMX application is running is monitored by Zabbix proxy, then you specify the connection parameters in the proxy configuration file instead.

```
JavaGateway=192.168.3.14
JavaGatewayPort=10052
```

By default, server does not start any processes related to JMX monitoring. If you wish to use it, however, you have to specify the number of pre-forked instances of Java pollers. You do this in the same way you specify regular pollers and trappers.

```
StartJavaPollers=5
```

Do not forget to restart server or proxy, once you are done with configuring them.

Debugging Java gateway

Zabbix Java gateway log file is:

/var/log/zabbix/zabbix_java_gateway.log

If you like to increase the logging, edit the file:

/etc/zabbix/zabbix_java_gateway_logback.xml

and change info in level="info" to debug or even trace (for deep troubleshooting):

```
<configuration scan="true" scanPeriod="15 seconds">
[...]
      <root level="info">
              <appender-ref ref="FILE" />
      </root>

</configuration>
```

JMX monitoring

See JMX monitoring page for more details.

Setup from RHEL packages

Overview

If installed from packages, the following information will help you in setting up Zabbix Java gateway.

Configuring and running Java gateway

Configuration parameters of Zabbix Java gateway may be tuned in the file:

/etc/zabbix/zabbix_java_gateway.conf

For more details, see Zabbix Java gateway configuration parameters.

To start Zabbix Java gateway:

```
systemctl restart zabbix-java-gateway
```

To automatically start Zabbix Java gateway on boot:

RHEL 7 and later:

```
systemctl enable zabbix-java-gateway
```

RHEL prior to 7:

```
chkconfig --level 12345 zabbix-java-gateway on
```

Configuring server for use with Java gateway

With Java gateway up and running, you have to tell Zabbix server where to find Zabbix Java gateway. This is done by specifying JavaGateway and JavaGatewayPort parameters in the server configuration file. If the host on which JMX application is running is monitored by Zabbix proxy, then you specify the connection parameters in the proxy configuration file instead.

```
JavaGateway=192.168.3.14
JavaGatewayPort=10052
```

By default, server does not start any processes related to JMX monitoring. If you wish to use it, however, you have to specify the number of pre-forked instances of Java pollers. You do this in the same way you specify regular pollers and trappers.

```
StartJavaPollers=5
```

Do not forget to restart server or proxy, once you are done with configuring them.

Debugging Java gateway

Zabbix Java gateway log file is:

/var/log/zabbix/zabbix_java_gateway.log

If you like to increase the logging, edit the file:

/etc/zabbix/zabbix_java_gateway_logback.xml

and change info in level="info" to debug or even trace (for deep troubleshooting):

```
<configuration scan="true" scanPeriod="15 seconds">
[...]
      <root level="info">
              <appender-ref ref="FILE" />
      </root>

</configuration>
```

JMX monitoring

See JMX monitoring page for more details.

Setup from sources

Overview

If installed from sources, the following information will help you in setting up Zabbix Java gateway.

Overview of files

If you obtained Java gateway from sources, you should have ended up with a collection of shell scripts, JAR and configuration files under `$PREFIX/sbin/zabbix_java`. The role of these files is summarized below.

`bin/zabbix-java-gateway-$VERSION.jar`

Java gateway JAR file itself.

```
lib/logback-core-1.5.25.jar
lib/logback-classic-1.5.25.jar
lib/slf4j-api-2.0.16.jar
lib/android-json-4.3_r3.1.jar
```

Dependencies of Java gateway: Logback, SLF4J, and Android JSON library.

```
lib/logback.xml
lib/logback-console.xml
```

Configuration files for Logback.

```
shutdown.sh
startup.sh
```

Convenience scripts for starting and stopping Java gateway.

`settings.sh`

Configuration file that is sourced by startup and shutdown scripts above.

Configuring and running Java gateway

By default, Java gateway listens on port 10052. If you plan on running Java gateway on a different port, you can specify that in `settings.sh` script. See the description of Java gateway configuration file for how to specify this and other options.

> **Warning:**
> Port 10052 is not IANA registered.

Once you are comfortable with the settings, you can start Java gateway by running the `startup.sh` script:

```
./startup.sh
```

Likewise, once you no longer need Java gateway, run the `shutdown.sh` script to stop it:

```
./shutdown.sh
```

Note that unlike server or proxy, Java gateway is lightweight and does not need a database.

Configuring server for use with Java gateway

With Java gateway up and running, you have to tell Zabbix server where to find Zabbix Java gateway. This is done by specifying `JavaGateway` and `JavaGatewayPort` parameters in the server configuration file. If the host on which JMX application is running is monitored by Zabbix proxy, then you specify the connection parameters in the proxy configuration file instead.

```
JavaGateway=192.168.3.14
JavaGatewayPort=10052
```

By default, server does not start any processes related to JMX monitoring. If you wish to use it, however, you have to specify the number of pre-forked instances of Java pollers. You do this in the same way you specify regular pollers and trappers.

```
StartJavaPollers=5
```

Do not forget to restart server or proxy, once you are done with configuring them.

Debugging Java gateway

In case there are any problems with Java gateway or an error message that you see about an item in the frontend is not descriptive enough, you might wish to take a look at Java gateway log file.

By default, Java gateway logs its activities into `/tmp/zabbix_java.log` file with log level "info". Sometimes that information is not enough and there is a need for information at log level "debug". In order to increase logging level, modify file `lib/logback.xml` and change the level attribute of `<root>` tag to `debug`:

```
<root level="debug">
  <appender-ref ref="FILE" />
</root>
```

Note that unlike Zabbix server or Zabbix proxy, there is no need to restart Zabbix Java gateway after changing `logback.xml` file - changes in `logback.xml` will be picked up automatically. When you are done with debugging, you can return the logging level to "info".

If you wish to log to a different file or a completely different medium like database, adjust `logback.xml` file to meet your needs. See Logback Manual for more details.

Sometimes for debugging purposes it is useful to start Java gateway as a console application rather than a daemon. To do that, comment out `PID_FILE` variable in `settings.sh`. If `PID_FILE` is omitted, `startup.sh` script starts Java gateway as a console application and makes Logback use `lib/logback-console.xml` file instead, which not only logs to console, but has logging level "debug" enabled as well.

Finally, note that since Java gateway uses SLF4J for logging, you can replace Logback with the framework of your choice by placing an appropriate JAR file in lib directory. See SLF4J Manual for more details.

JMX monitoring

See JMX monitoring page for more details.

Sender

Overview Zabbix sender is a command line utility that may be used to send performance data to Zabbix server or proxy for processing.

The utility is usually used in long running user scripts for periodical sending of availability and performance data.

For sending results directly to Zabbix server or proxy, a *Zabbix trapper* item type must be configured.

See also:

- Zabbix sender protocol
- Zabbix sender dynamic link library for Windows
- Python library for Zabbix that has built-in functionality to act like Zabbix sender

Running Zabbix sender An example of running Zabbix UNIX sender:

```
cd bin
./zabbix_sender -z zabbix -s "Linux DB3" -k db.connections -o 43
```

where:

- `z` - Zabbix server host (IP address can be used as well);
- `s` - technical name of monitored host (as registered in Zabbix frontend);
- `k` - item key;
- `o` - value to send.

Attention:
Options that contain whitespaces, must be quoted using double quotes.

Zabbix sender can be used to send multiple values from an input file. See the Zabbix sender manpage for more information.

If a configuration file is specified, Zabbix sender uses all addresses defined in the agent `ServerActive` configuration parameter to send data. If sending to one address fails, the sender tries sending to the other addresses. If sending of batch data fails to one address, the following batches are not sent to this address.

Zabbix sender accepts strings in UTF-8 encoding (for both Unix-like systems and Windows) without byte order mark (BOM) first in the file.

Zabbix sender on Windows can be run similarly:

```
zabbix_sender.exe [options]
```

`zabbix_sender` realtime sending scenarios will gather multiple values passed to it in close succession and send them to the server in a single connection. A value that is not further apart from the previous value than 0.2 seconds can be put in the same stack, but maximum polling time still is 1 second.

Note:
Zabbix sender will terminate if invalid (not following `parameter=value` notation) parameter entry is present in the specified configuration file.

Running Zabbix sender with low-level discovery An example of running Zabbix sender for sending a JSON-formatted value for low-level discovery:

```
./zabbix_sender -z 192.168.1.113 -s "Zabbix server" -k trapper.discovery.item -o '[{"{#FSNAME}":"/","{#FST
```

For this to work, the low-level discovery rule must have the *Zabbix trapper* item type (in this example, with `trapper.discovery.item` key).

Zabbix sender protocol

Zabbix server and Zabbix proxy use a JSON-based communication protocol for receiving data from Zabbix sender. Data can be received with the help of a trapper item, or an HTTP agent item with trapping enabled.

Request and response messages must begin with header and data length.

Zabbix sender request

```
{
    "request": "sender data",
    "data": [
        {
            "host": "<hostname>",
            "key": "trap",
            "value": "test value"
        }
    ]
}
```

Zabbix server response

```
{
    "response": "success",
    "info": "processed: 1; failed: 0; total: 1; seconds spent: 0.060753"
}
```

Zabbix sender request with a timestamp

Alternatively, Zabbix sender can send a request with a timestamp and nanoseconds.

```
{
    "request": "sender data",
    "data": [
        {
            "host": "<hostname>",
            "key": "trap",
            "value": "test value",
            "clock": 1516710794,
```

```
            "ns": 592397170
        },
        {
            "host": "<hostname>",
            "key": "trap",
            "value": "test value",
            "clock": 1516710795,
            "ns": 192399456
        }
    ],
    "clock": 1516712029,
    "ns": 873386094
}
```

Zabbix server response

```
{
    "response": "success",
    "info": "processed: 2; failed: 0; total: 2; seconds spent: 0.060904"
}
```

Zabbix sender dynamic link library for Windows

In a Windows environment applications can send data to Zabbix server/proxy by using the Zabbix sender dynamic link library (`zabbix_sender.dll`) instead of having to launch an external process (`zabbix_sender.exe`).

zabbix_sender.h and zabbix_sender.lib are required for compiling user applications with zabbix_sender.dll.

Getting it

There are two ways of getting zabbix_sender.dll.

1. Download zabbix_sender.h, zabbix_sender.lib and zabbix_sender.dll files as a ZIP archive.

When choosing download options make sure to select "No encryption" under *Encryption* and "Archive" under *Packaging*. Then download Zabbix agent (not Zabbix agent 2).

The zabbix_sender.h, zabbix_sender.lib and zabbix_sender.dll files will be inside the downloaded ZIP archive in the `bin\dev` directory. Unzip the files where you need it.

2. Build zabbix_sender.dll from source (see instructions).

The dynamic link library with the development files will be located in the `bin\winXX\dev` directory. To use it, include the zabbix_sender.h header file and link with the zabbix_sender.lib library.

See also

- example of a simple Zabbix sender utility implemented with Zabbix sender dynamic link library to illustrate the library usage;
- zabbix_sender.h file for the interface functions of the Zabbix sender dynamic link library. This file contains documentation explaining the purpose of each interface function, its arguments, and return value.

Get

Overview Zabbix get is a command-line utility, which can be used to communicate with Zabbix agent and retrieve required information from the agent.

The utility is usually used for the troubleshooting of Zabbix agents.

See also the Python library for Zabbix that has built-in functionality to act like Zabbix get.

Running Zabbix get An example of running Zabbix get under UNIX to get the processor load value from the agent:

```
cd bin
./zabbix_get -s 127.0.0.1 -p 10050 -k system.cpu.load[all,avg1]
```

Another example of running Zabbix get for capturing a string from a website:

```
cd bin
./zabbix_get -s 192.168.1.1 -p 10050 -k "web.page.regexp[www.example.com,,,\"USA: ([a-zA-Z0-9.-]+)\",,\1]"
```

Note that the item key here contains a space so quotes are used to mark the item key to the shell. The quotes are not part of the item key; they will be trimmed by the shell and will not be passed to Zabbix agent.

If an item key is unsupported Zabbix get will return the exit code 1.

Zabbix get accepts the following command line parameters:

```
-s --host <host name or IP>             Specify host name or IP address of a host
-p --port <port number>                 Specify port number of agent running on the host (default: 10050)
-I --source-address <IP address>        Specify source IP address
-t --timeout <seconds>                  Specify timeout. Valid range: 1-600 seconds (default: 30 seconds)
-k --key <item key>                     Specify key of item to retrieve value for
-P --protocol <value>                   Protocol used to communicate with agent. Values:
                                            auto - connect using JSON protocol, fallback and retry with pl
                                            json - connect using JSON protocol
                                            plaintext - connect using plaintext protocol where just item k
-h --help                               Display this help message
-V --version                            Display version number

--tls-connect <value>                   How to connect to agent. Values:
                                            unencrypted - connect without encryption (default)
                                            psk - connect using TLS and a pre-shared key
                                            cert - connect using TLS and a certificate
--tls-ca-file <CA file>                 Full pathname of a file containing the top-level CA(s) certificate
--tls-crl-file <CRL file>               Full pathname of a file containing revoked certificates
--tls-agent-cert-issuer <cert issuer>   Allowed agent certificate issuer
--tls-agent-cert-subject <cert subject> Allowed agent certificate subject
--tls-cert-file <cert file>             Full pathname of a file containing the certificate or certificate
--tls-key-file <key file>               Full pathname of a file containing the private key
--tls-psk-identity <PSK-identity>       Unique, case sensitive string used to identify the pre-shared key
--tls-psk-file <PSK-file>               Full pathname of a file containing the pre-shared key
--tls-cipher13 <cipher-string>          Cipher string for OpenSSL 1.1.1 or newer for TLS 1.3. Override the
--tls-cipher <cipher-string>            GnuTLS priority string (for TLS 1.2 and up) or OpenSSL cipher stri
```

See also Zabbix get manpage for more information.

Zabbix get on Windows can be run similarly:

```
zabbix_get.exe [options]
```

JS

Overview zabbix_js is a command line utility that can be used for embedded script testing.

This utility will execute a user script with a string parameter and print the result. Scripts are executed using the embedded Zabbix scripting engine.

In case of compilation or execution errors zabbix_js will print the error in stderr and exit with code 1.

```
zabbix_js -s script-file -p input-param [-l log-level] [-t timeout]
zabbix_js -s script-file -i input-file [-l log-level] [-t timeout]
zabbix_js -h
zabbix_js -V
```

Usage zabbix_js accepts the following command line parameters:

```
-s, --script script-file          Specify the file name of the script to execute
                                  If "-" is specified as the file name, the script will be read from stdin
-i, --input input-file            Specify the file name of the input information
                                  The content will be passed to the script as a value
```

```
                                  Mutually exclusive with -p
                                  If "-" is specified as the file name, the script will be read from stdin
-p, --param input-param           Specify the input parameter: the variable to be passed to the script as
                                  Mutually exclusive with -i
-l, --loglevel log-level          Specify the log level
                                  Range: 0-5
-t, --timeout timeout             Specify the timeout in seconds
                                  Valid range: 1-600 seconds (default: 10 seconds)
-h, --help                        Display help information
-V, --version                     Display the version number
-w <webdriver url>                Enables browser monitoring
```

Examples: Example 1: Executing a script with a variable containing a value

The script (`script-file.js`):

```
return value;
```

Usage:

```
zabbix_js -s script-file.js -p example_value
```

Output: `example_value`

Example 2: Executing a script with a file containing input information

The file with the script (`script-file.js`):

```
return value;
```

The file with the input information (`example.txt`):

```
Example of input information from the file
```

Usage:

```
zabbix_js -s script-file.js -i example.txt
```

Output: `Example of input information from the file`

Example 3: Executing a script with input read from stdin

The file with the script (`script-file.js`):

```
return value;
```

Usage:

```
zabbix_js -s script-file.js -i -
```

Input (stdin):

```
Example of input from stdin
```

Output: `Example of input from stdin`

Web service

Overview Zabbix web service is a process that is used for communication with external web services. Currently, Zabbix web service is used for generating and sending scheduled reports with plans to add additional functionality in the future.

Zabbix server connects to the web service via HTTP(S). Zabbix web service requires Google Chrome to be installed on the same host; on some distributions the service may also work with Chromium (see known issues).

Installation The official zabbix-web-service package is available in the Zabbix repository.

To compile Zabbix web service from sources, specify the `--enable-webservice` configure option.

To configure Zabbix web service, update the web service configuration file parameters.

Attention:
It is strongly recommended to set up encryption between Zabbix server and Zabbix web service using certificates. By default, data transmitted between Zabbix server and Zabbix web service is not encrypted, which can lead to unauthorized access.

Configuration parameters: Zabbix web service

Overview

The Zabbix web service is a process that is used for communication with external web services.

The parameters supported by the Zabbix web service configuration file (`zabbix_web_service.conf`) are listed in this section.

The parameters are listed without additional information. Click the parameter to see the full details.

Parameter	Description
AllowedIP	A list of comma delimited IP addresses, optionally in CIDR notation, or DNS names of Zabbix servers and Zabbix proxies.
DebugLevel	The debug level.
IgnoreURLCertErrors	Specifies TLS certificate validation error handling when accessing the frontend URL.
Include	You may include individual files or all files in a directory in the configuration file.
ListenPort	The service will listen on this port for connections from the server.
LogFile	The name of the log file.
LogFileSize	The maximum size of the log file.
LogType	The type of the log output.
Timeout	The maximum time (in seconds) spent formatting the PDF report of a dashboard.
TLSAccept	What incoming connections to accept.
TLSCAFile	The full pathname of a file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.
TLSCertFile	The full pathname of a file containing the service certificate or certificate chain, used for encrypted communications between Zabbix components.
TLSKeyFile	The full pathname of a file containing the service private key, used for encrypted communications between Zabbix components.

All parameters are non-mandatory unless explicitly stated that the parameter is mandatory.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files;
- Values support environment variables;
- Zabbix supports configuration files only in UTF-8 encoding without BOM;
- Comments starting with # are only supported at the beginning of the line.

Parameter details

AllowedIP

A list of comma delimited IP addresses, optionally in CIDR notation, or DNS names of Zabbix servers and Zabbix proxies. Incoming connections will be accepted only from the hosts listed here.
If IPv6 support is enabled then `127.0.0.1`, `::127.0.0.1`, `::ffff:127.0.0.1` are treated equally and `::/0` will allow any IPv4 or IPv6 address. `0.0.0.0/0` can be used to allow any IPv4 address.

Example:

`127.0.0.1,192.168.1.0/24,::1,2001:db8::/32,zabbix.example.com`

Mandatory: yes

DebugLevel

Specify the debug level: - *0* - basic information about starting and stopping of Zabbix processes; - *1* - critical information; - *2* - error information; - *3* - warnings; - *4* - for debugging (produces lots of information); - *5* - extended debugging (produces even more information).

Default: 3
 Range: 0-5

IgnoreURLCertErrors

Specifies TLS certificate validation error handling when accessing the frontend URL:
0 - do not ignore certificate errors;
1 - ignore certificate errors.

Default: 0
 Range: 0-1

Include

Specify a directory or individual files to be included in the configuration file.
 During Zabbix installation, this parameter is set to `/usr/local/etc`, unless you change this during compile time.

General rules for the `Include` parameter:

- The path can be relative to the `zabbix_web_service.conf` file location.
- You can specify multiple `Include` parameters in one configuration file.
- Zabbix processes the `Include` parameter as soon as it reads it, so consider placing it at the end of the configuration file.

Rules for included files:

- Each file must be readable.
- Each file must have correct syntax. Otherwise, the component will not start.

Rules for including directories or file patterns:

- Do not assume any order of inclusion (e.g., files are not included in alphabetical order).
- Do not define the same parameter in more than one included file (e.g., to override a general setting with a specific one), since you cannot predict which value will be used.
- Watch out for backup copies that some text editors create automatically (e.g., editing `include/my_specific.conf` may create `include/my_specific.conf.BAK`), since both files will be included.

To include only relevant files from a directory, use the wildcard (`*`) for pattern matching.

Example:

```
Include=/absolute/path/to/config/files/*.conf
```

ListenPort

The service will listen on this port for connections from the server.

Default: 10053
 Range: 1024-32767

LogFile

The name of the log file.

Example:

`/tmp/zabbix_web_service.log`

Mandatory: Yes, if `LogType` is set to `file`; otherwise, no

LogFileSize

The maximum size of a log file in MB.
0 - disable automatic log rotation.
Note: If the log file size limit is reached and file rotation fails, for whatever reason, the existing log file is truncated and started anew.

Default: 1
 Range: 0-1024

LogType

The type of the log output: - *file* - write log to the file specified by LogFile parameter; - *system* - write log to syslog; - *console* - write log to standard output.

Default: `file`

Timeout

The maximum time (in seconds) spent formatting the PDF report of a dashboard.

Default: 10
 Range: 1-30

TLSAccept

What incoming connections to accept: - *unencrypted* - accept connections without encryption (default); - *cert* - accept connections with TLS and a certificate.

Default: `unencrypted`

TLSCAFile

The full pathname of the file containing the top-level CA(s) certificates for peer certificate verification, used for encrypted communications between Zabbix components.

TLSCertFile

The full pathname of the file containing the service certificate or certificate chain, used for encrypted communications with Zabbix components.

TLSKeyFile

The full pathname of the file containing the service private key, used for encrypted communications between Zabbix components.

Bridge adapter

Zabbix bridge adapter lets your Zabbix server communicate with Zabbix mobile app.

Zabbix mobile app doesn't connect to your Zabbix server directly. Instead, two components pass data between them:

- Zabbix Bridge - a centrally hosted service that exchanges data with Zabbix mobile app on your mobile device.
- Zabbix bridge adapter - a component that you install and run. It connects your Zabbix server to Zabbix Bridge.

The bridge adapter can be installed on the same system as your Zabbix server, or on a different system, using packages.

To configure Zabbix bridge adapter, update its configuration file (default path: `/etc/zabbix/zabbix_bridge_adapter.conf`).

To keep Zabbix bridge adapter secure, see Secure mobile access.

Data flow When you use Zabbix mobile app, data travels as follows:

1. You perform an action in Zabbix mobile app, which then sends a query to Zabbix Bridge.
2. Zabbix Bridge passes the query to Zabbix bridge adapter as a JSON-RPC request.
3. Zabbix bridge adapter checks whether the requested Zabbix API method is allowed. If the method isn't allowed, the bridge adapter rejects the request. If the method is allowed, the bridge adapter sends the request to Zabbix API.
4. Zabbix API processes the request and returns a response to Zabbix bridge adapter.
5. Zabbix bridge adapter returns the response to Zabbix Bridge.
6. Zabbix Bridge passes the response to Zabbix mobile app, which shows the data on your mobile device.

Configuration parameters: Zabbix bridge adapter

Overview

Configuration parameters allow customizing Zabbix bridge adapter behavior.

Specify parameter values in the configuration file (`zabbix_bridge_adapter.conf`). All parameters are **optional**, unless explicitly stated that the parameter is mandatory.

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
AllowedAPIMethods	Specify which Zabbix API methods can be called through Zabbix bridge adapter.
DataDir	Specify the directory where the bridge adapter state file and encryption keys are created and stored.
DebugLevel	Set the debug level.
Include	Specify a directory or individual files to be included in the configuration file.
ListenIP	Specify the bridge adapter listen IP address for incoming HTTP(S) connections from Zabbix server.
ListenHealthPort	Specify the bridge adapter listen port for incoming HTTP health-check connections.
ListenPort	Specify the bridge adapter listen port for incoming HTTP(S) connections from Zabbix server.
LogFile	Specify the full path to the bridge adapter log file.
LogFileSize	Set the maximum log file size (in MB).
LogType	Specify the log output type.
PidFile	Specify the PID file.

Parameter	Description
TLSConfig.AllowedCN	Specify a comma-separated list of permitted client Common Names (mTLS only).
TLSConfig.CAFile	Specify the full path to a file containing the top-level CA certificate(s) used to verify client certificates in mTLS mode.
TLSConfig.CertFile	Specify the full path to a file containing the certificate or certificate chain.
TLSConfig.KeyFile	Specify the full path to a file containing the private key.
TLSConfig.Mode	Set the TLS mode for incoming connections on ListenPort.
ServerId	Specify the unique identifier of Zabbix server. **Mandatory**.
ZabbixAPITLSCAFile	Specify the full pathname of a file containing the top-level CA certificate(s) used by the bridge adapter to verify the Zabbix API certificate.
ZabbixAPITLSConnect	Set the TLS mode for outbound connections to the Zabbix API.
ZabbixAPIURL	Specify the URL of the Zabbix API endpoint. **Mandatory**.
ZabbixBridgeURL	Specify Zabbix Bridge URL for Zabbix Bridge configuration retrieval (retrieved when the bridge adapter starts).

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files.
- Values support environment variables.
- Zabbix supports configuration files only in UTF-8 encoding without BOM;
- Comments starting with # are only supported at the beginning of the line.

Parameter details

AllowedAPIMethods

Specify which Zabbix API methods can be called through Zabbix bridge adapter.
 For several methods, repeat the parameter, one method per line.

Each parameter value must be either an exact method name (e.g., problem.get) or a wildcard pattern that replaces only the whole API name (e.g., *.get) or the whole method name (e.g., problem.*).

Any other value is invalid and prevents Zabbix bridge adapter from starting. This includes a lone wildcard (*), a partial wildcard pattern (host* or host*.*), or an empty value (AllowedAPIMethods=).

Example:

```
AllowedAPIMethods=apiinfo.version
AllowedAPIMethods=problem.*
AllowedAPIMethods=*.get
```

The minimal set of methods that Zabbix mobile app requires:

```
AllowedAPIMethods=apiinfo.version
AllowedAPIMethods=device.onboard
AllowedAPIMethods=device.offboard
AllowedAPIMethods=user.get
AllowedAPIMethods=role.get
AllowedAPIMethods=problem.get
AllowedAPIMethods=event.get
AllowedAPIMethods=event.acknowledge
AllowedAPIMethods=trigger.get
AllowedAPIMethods=alert.get
AllowedAPIMethods=host.get
AllowedAPIMethods=hostgroup.get
```

Default: not set (all Zabbix API methods are denied)

DataDir

Specify the directory where the bridge adapter state file and encryption keys are created and stored.
 The directory must have 0700 permissions, and all files stored in it must have 0600 permissions.

Default: /var/lib/zabbix-bridge-adapter

DebugLevel

Set the debug level:

- 0 - basic information about starting and stopping of Zabbix processes

- 1 - critical information
- 2 - error information
- 3 - warnings
- 4 - for debugging (produces lots of information)
- 5 - extended debugging (produces even more information)

Default: 3
 Range: 0-5

Include

Specify a directory or individual files to be included in the configuration file.
 During Zabbix installation, this parameter is set to `/usr/local/etc`, unless you change this during compile time.

General rules for the `Include` parameter:

- The path can be relative to the `zabbix_bridge_adapter.conf` file location.
- You can specify multiple `Include` parameters in one configuration file.
- Zabbix processes the `Include` parameter as soon as it reads it, so consider placing it at the end of the configuration file.

Rules for included files:

- Each file must be readable.
- Each file must have correct syntax. Otherwise, the component will not start.

Rules for including directories or file patterns:

- Do not assume any order of inclusion (e.g., files are not included in alphabetical order).
- Do not define the same parameter in more than one included file (e.g., to override a general setting with a specific one), since you cannot predict which value will be used.
- Watch out for backup copies that some text editors create automatically (e.g., editing `include/my_specific.conf` may create `include/my_specific.conf.BAK`), since both files will be included.

To include only relevant files from a directory, use the wildcard (`*`) for pattern matching.

Example:

```
Include=/absolute/path/to/config/files/*.conf
```

ListenIP

Specify the bridge adapter listen IP address for incoming HTTP(S) connections from Zabbix server.
 If not specified, the bridge adapter listens on all IP addresses.

ListenHealthPort

Specify the bridge adapter listen port for incoming HTTP health-check connections.
 Serves the `/health` and `/ready` endpoints.
 If not specified, the health service is not started.

Default: 8081
 Range: 1024-32767

ListenPort

Specify the bridge adapter listen port for incoming HTTP(S) connections from Zabbix server.

Default: 10054
 Range: 1024-32767

LogFile

Specify the full path to the bridge adapter log file.

Mandatory: **Yes** if `LogType` is set to `file`

Default: `/tmp/zabbix_bridge_adapter.log`

LogFileSize

Set the maximum log file size (in MB).
 0 - disable automatic log rotation.

Default: 1
 Range: 0-1024

LogType

Specify the log output type:

- `file` - write log to file specified by `LogFile` parameter.
- `system` - write log to syslog.
- `console` - write log to standard output.

Default: `file`

PidFile

Specify the PID file.

Default: `/tmp/zabbix_bridge_adapter.pid`

ServerId

Specify the unique identifier of Zabbix server.
 Available in *Reports* > *System information* in the Zabbix user interface.

Mandatory: **Yes**

TLSConfig.AllowedCN

Specify a comma-separated list of permitted client Common Names (`mTLS` only).
 If not specified, the bridge adapter does not check the client Common Name.

TLSConfig.CAFile

Specify the full path to a file containing the top-level CA certificate(s) used to verify client certificates in `mTLS` mode.

TLSConfig.CertFile

Specify the full path to a file containing the certificate or certificate chain.

TLSConfig.KeyFile

Specify the full path to a file containing the private key.

TLSConfig.Mode

Set the TLS mode for incoming connections on `ListenPort`:

- `none` - accept connections without encryption.
- `tls` - standard TLS with v1.2 minimum.
- `mtls` - extends TLS with mutual authentication.

Default: `none`

ZabbixAPITLSCAFile

Specify the full pathname of a file containing the top-level CA certificate(s) used by the bridge adapter to verify the Zabbix API certificate.

Mandatory: **Yes** if `ZabbixAPITLSConnect` is set to `verify_ca` or `verify_full`.

ZabbixAPITLSConnect

Set the TLS mode for outbound connections to the Zabbix API:

- `disabled` (or empty) - use plain HTTP; `ZabbixAPIURL` must use `http://`.
- `required` - use HTTPS without certificate verification (testing only).
- `verify_ca` - use HTTPS and verify the Zabbix API certificate.
- `verify_full` - use HTTPS and verify the Zabbix API certificate and hostname.

ZabbixAPIURL

Specify the URL of the Zabbix API endpoint.

Mandatory: **Yes**

Example:

```
ZabbixAPIURL=https://example.com/zabbix/api_jsonrpc.php
```

ZabbixBridgeURL

Specify Zabbix bridge URL for Zabbix Bridge configuration retrieval (retrieved when the bridge adapter starts).

Default: `ZabbixBridgeURL=https://test.aws.zabbix-lab.win/.well-known/bridge-configuration`

MCP server

Overview Zabbix MCP server is a process that allows AI agents to work with Zabbix using the Model Context Protocol (MCP). MCP server communicates with MCP clients and Zabbix API.

Using the MCP server, an AI agent can:

- Retrieve host groups, template groups, hosts, templates, items, LLD rules, interfaces, macros, triggers, problems, maintenance periods, and item history data (except binary items);
- Perform problem updates: add comments, acknowledge or unacknowledge, change severity, suppress or unsuppress, or close problems;
- Create, update, and delete maintenance periods.

Zabbix API tokens are used for authorization; the tokens are sent with every request. Access is limited by the permissions of the token owner: only data from host groups available to that user is visible, and actions the user is not permitted to perform are rejected.

The MCP server is stateless: no session is kept between requests.

Note:
While Zabbix MCP server may work with older/newer versions of Zabbix, its compatibility is officially supported only for the same version of Zabbix as the MCP server itself.

See also:

- Configuration parameters: MCP server
- Tools and resources

Installation The official zabbix-mcp-server package is available in the Zabbix repository.

To install Zabbix MCP server from sources, build it with `make`.

To configure Zabbix MCP server, update the MCP server configuration file parameters. At least the `FrontendURL` and `AllowedIP` parameters must be set.

The default configuration file (`zabbix_mcp_server.conf`) is located at `/etc/zabbix/`; it denies the tools that most often expose sensitive data. To enable a tool, remove the corresponding `DenyTool` line.

See also: Security best practices

Running To start the MCP server, run:

```
zabbix_mcp_server -c /etc/zabbix/zabbix_mcp_server.conf
```

To check the configuration file and exit without starting the server, use the `-T` option:

```
zabbix_mcp_server -c /etc/zabbix/zabbix_mcp_server.conf -T
```

The following command-line options are supported:

Option	Description
`-c <file>`	The path to the configuration file. The default is `/etc/zabbix/zabbix_mcp_server.conf`.
`-T`	Validate the configuration file and exit.
`-h`	Display help information.
`-V`	Display the version number.

When the MCP server is installed from sources, a systemd unit must be created manually, for example:

```
[Unit]
Description=Zabbix MCP server
After=network-online.target

[Service]
User=zabbix
ExecStart=/usr/local/sbin/zabbix_mcp_server -c /etc/zabbix/zabbix_mcp_server.conf
Restart=on-failure

[Install]
WantedBy=multi-user.target
```

Connecting an MCP client Any MCP client that supports the Streamable HTTP transport can be connected to the MCP server. The following settings are required:

Setting	Value
Transport	Streamable HTTP
URL	`https://<host>:<ListenPort>/mcp`
Header	`Authorization: Bearer <Zabbix API token>`, sent with every request

Each user must configure their own API token.

To verify the setup, send a single request to the MCP server:

```
curl -sS https://mcp.example.com:8443/mcp \
  -H "Authorization: Bearer $ZABBIX_API_TOKEN" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json, text/event-stream" \
  -d '{"jsonrpc":"2.0","id":1,"method":"tools/list"}'
```

The HTTP status code 200 and a list of the available tools in the response confirm that the MCP server is reachable and accepts the token; the list shows which tools are exposed. Other status codes indicate the following problems:

- 400 - the `Accept` header does not include both media types required by the transport;
- 401 - the `Authorization` header with the token is missing;
- 403 - the client address is not listed in the `AllowedIP` parameter.

Approval and tool access If the MCP client supports form elicitation, every change must be approved by the user before it is sent to Zabbix. The MCP client shows the user what will be changed, and the change is applied only if the user confirms it; if the user rejects or dismisses the request, nothing is changed in Zabbix. An approval is valid for five minutes and only for the exact change that was shown; within this time, the approval can be reused.

If the MCP client does not support form elicitation (including clients that support only URL elicitation), changes are applied immediately, without approval. In this case, restrict write tools and Zabbix API permissions.

Approval is particularly important because tool output may contain untrusted text from monitored systems that could influence an agent.

Use AllowTool, DenyTool, AllowToolRegexp, and DenyToolRegexp to control which tools the MCP server exposes. Denied tools cannot be called or listed, and their associated resources cannot be accessed.

> **Warning:**
> AllowTool, DenyTool, AllowToolRegexp, and DenyToolRegexp rules limit MCP access but do not change the permissions of the Zabbix token. If the frontend address is exposed to an AI agent, the agent may attempt to send direct API calls. Keep the Zabbix frontend location and MCP server configuration inaccessible to the agent so that requests must go through the MCP server and its security controls.

Security best practices API tokens

Authentication and authorization are handled entirely by Zabbix. The MCP client sends a Zabbix API token with each request; the MCP server forwards it without storing or logging it, while Zabbix validates the token, enforces permissions, and records actions in the audit log.

Use a separate, least-privileged, expiring token for each agent and user. Store tokens in a secure location and rotate API tokens regularly.

TLS encryption

To prevent tokens from being sent in clear text, set `TLSAccept=cert` and specify the certificate and the key (`TLSCertFile` and `TLSKeyFile`) or run the MCP server behind a trusted TLS-terminating reverse proxy and set `TLSAccept=unencrypted`. Unencrypted HTTP is acceptable only in a trusted network.

Access restrictions

Restrict the `AllowedIP` parameter to the addresses of trusted MCP clients, and use Zabbix permissions to limit what API tokens can do.

The MCP server does not provide request rate limiting. Each MCP request normally results in one Zabbix API call, although an agent running in a loop can still generate significant load.

Sensitive data protection

Item keys may contain credentials (for example, in SSH, JMX, ODBC and HTTP agent items), macro values may contain passwords and community strings, interfaces reveal internal addresses, and log or text item values may contain any data. For this reason, the `item_get`, `lld_get`, `interface_get`, `macro_get` and `history_get` tools are denied by default in the shipped configuration file.

To avoid sensitive macro value exposure to the LLM provider of the connected agent, use secret macros; their values are never returned by the API.

Configuration parameters: Zabbix MCP server

Overview

The Zabbix MCP server is a process that allows AI agents to work with Zabbix using the Model Context Protocol (MCP).

The parameters supported by the Zabbix MCP server configuration file (`zabbix_mcp_server.conf`) are listed in this section.

The parameters are listed without additional information. Click the parameter to see the full details.

Parameter	Description
AllowedIP	A list of comma delimited IP addresses, optionally in CIDR notation, or DNS names of MCP clients.
AllowTool	A wildcard rule that allows a tool.
AllowToolRegexp	A regular expression rule that allows tools.
DebugLevel	The debug level.
DenyTool	A wildcard rule that denies a tool.
DenyToolRegexp	A regular expression rule that denies tools.
FrontendURL	The Zabbix API endpoint.
IgnoreURLCertErrors	Specifies TLS certificate validation error handling when accessing the Zabbix API.
ListenIP	The IP address to listen on.
ListenPort	The port to listen on.
LogFile	The name of the log file.
LogFileSize	The maximum size of the log file.
LogType	The type of the log output.
MaxResponseSize	The maximum size of a response returned by a read tool.
PidFile	The name of the PID file.
Timeout	The maximum time (in seconds) to wait for a response from the Zabbix API.
TLSAccept	What incoming connections to accept.
TLSCertFile	The full pathname of a file containing the server certificate or certificate chain.
TLSKeyFile	The full pathname of a file containing the server private key.
ToolPrefix	A prefix added to the names of all tools.

All parameters are non-mandatory unless explicitly stated that the parameter is mandatory.

Note that:

- The default values reflect process defaults, not the values in the shipped configuration files;
- Values support environment variables in the `${NAME}` format;
- Zabbix supports configuration files only in UTF-8 encoding without BOM;
- Comments starting with # are only supported at the beginning of the line.

Parameter details

AllowedIP

A list of comma delimited IP addresses, optionally in CIDR notation, or DNS names of MCP clients. Incoming connections will be accepted only from the hosts listed here.

Example:

```
AllowedIP=127.0.0.1,192.168.1.0/24,::1,2001:db8::/32,mcp-client.example.com
```

Mandatory: yes

AllowTool

A wildcard rule that allows a tool. The parameter can be specified multiple times.
 Wildcard patterns use the item key syntax, where * matches any characters. The rule matches the plain tool name and ignores ToolPrefix.
 All `AllowTool`, `DenyTool`,

`AllowToolRegexp` and `DenyToolRegexp` rules form a single list in the order in which they are specified; the first matching rule decides. A tool that matches no rule is allowed.

Example:

`AllowTool=host_get`

AllowToolRegexp

A regular expression rule that allows tools. The parameter can be specified multiple times.
 Regular expressions use the RE2 syntax and match any part of the plain tool name unless anchored.

Example:

`AllowToolRegexp=^(host|hostgroup|trigger|problem)_get$`

DebugLevel

Specify the debug level: - *0* - basic information about starting and stopping of Zabbix processes; - *1* - critical information; - *2* - error information; - *3* - warnings; - *4* - for debugging (produces lots of information); - *5* - extended debugging (produces even more information).

Default: 3
 Range: 0-5

DenyTool

A wildcard rule that denies a tool. The parameter can be specified multiple times.
 A denied tool is not included in the list of tools and cannot be called. See AllowTool for details.

Example:

`DenyTool=maintenance_delete`

DenyToolRegexp

A regular expression rule that denies tools. The parameter can be specified multiple times.
 See AllowToolRegexp for details.

Example:

`DenyToolRegexp=^problem_(close|suppress)$`

FrontendURL

The Zabbix API endpoint: the URL of the Zabbix frontend ending with `api_jsonrpc.php`.

Example:

`FrontendURL=https://zabbix.example.com/api_jsonrpc.php`

Mandatory: yes

IgnoreURLCertErrors

Specifies TLS certificate validation error handling when accessing the Zabbix API:
 false - do not ignore certificate errors;
 true - ignore certificate errors (use only in a development environment).

Default: `false`

ListenIP

The IP address to listen on.

Default: `0.0.0.0`

ListenPort

The port to listen on. The MCP endpoint path is `/mcp`.

Default: 8443
 Range: 1024-32767

LogFile

The name of the log file.

Default: `/tmp/zabbix_mcp_server.log`

Mandatory: No

LogFileSize

The maximum size of a log file in MB.
0 - disable automatic log rotation.

Default: 1
 Range: 0-1024

LogType

The type of the log output: - *file* - write log to the file specified by LogFile parameter; - *system* - write log to syslog; - *console* - write log to standard output.

Default: `file`

MaxResponseSize

The maximum size of a response returned by a read tool. Whichever of this limit and the `limit` argument of the tool is reached first stops the response.
 Supports unit suffixes.

Default: 64K
 Range: 16K-256K

PidFile

The name of the PID file.

Default: `/tmp/zabbix_mcp_server.pid`

Timeout

The maximum time (in seconds) to wait for a response from the Zabbix API.

Default: 10
 Range: 3-30

TLSAccept

What incoming connections to accept: - *unencrypted* - accept connections without encryption (default); - *cert* - accept connections with TLS and a certificate.

Default: `unencrypted`

TLSCertFile

The full pathname of the file containing the server certificate or certificate chain.

Mandatory: Yes, if `TLSAccept` is set to `cert`; otherwise, no

TLSKeyFile

The full pathname of the file containing the server private key.

Mandatory: Yes, if `TLSAccept` is set to `cert`; otherwise, no

ToolPrefix

A prefix added to the names of all tools, for example, to avoid name conflicts with tools of other MCP servers. For example, with `ToolPrefix=zbx_`, the `host_get` tool is exposed as `zbx_host_get`.
 Allowed characters: letters, digits, underscore (_), hyphen (-) and dot (.).
 Tool access rules match the tool names without the prefix.

Example:

```
ToolPrefix=zbx_
```

Tools and resources

Overview

Zabbix MCP server exposes Zabbix functionality to MCP clients as *tools* and *resources*.

Tool names describe the activity. If the ToolPrefix parameter is set, the prefix is added to every tool name; the tool access rules match the plain names listed below.

Read tools

Tool	Zabbix API method	Arguments	Sorted by	Resource
`hostgroup_get`	hostgroup.get	`groupids`, `names`, `search`, `limit`	name	`zabbix://hostgroup/<i`
`templategroup_get`	templategroup.get	`groupids`, `names`, `search`, `limit`	name	`zabbix://templategrou`
`host_get`	host.get	`hostids`, `groupids`, `search`, `limit`	name	`zabbix://host/<id>`
`template_get`	template.get	`templateids`, `groupids`, `search`, `limit`	name	`zabbix://template/<id`
`item_get`	item.get	`hostids`, `itemids`, `search`, `limit`	name	`zabbix://item/<id>`
`lld_get`	discoveryrule.get	`hostids`, `ruleids`, `limit`	name	-

Tool	Zabbix API method	Arguments	Sorted by	Resource
interface_get	hostinterface.get	hostids, interfaceids, limit	interface ID	-
macro_get	usermacro.get	hostids or global, macro_names, limit	macro	-
problem_get	problem.get	eventids, groupids, hostids, severities, acknowledged, suppressed, recent, limit	event ID, descending	-
maintenance_get	maintenance.get	maintenanceids, groupids, hostids, limit	name	-
history_get	history.get	itemids, time_from, time_till, type, limit	clock, descending	-
trigger_get	trigger.get	triggerids, hostids, search, limit	description	zabbix://trigger/<id>

Note that:

- host_get and template_get also return the triggers of each object (a limited set of fields), tags and linked templates. Items, discovery rules, interfaces and macros are returned by the corresponding tools.
- history_get requires a time range that is not wider than 7 days and starts no earlier than 7 days ago. The type argument must be one of float, unsigned, character, text, log or json, and must match the value type of the items; if the type is wrong, an empty result is returned, which cannot be distinguished from an item without data. The history of binary items (value type 5) cannot be retrieved through the MCP server.
- Read tools request a fixed list of fields; all fields (output: extend) are never requested.
- The results of tools that have a resource contain the uri field, so that an agent can proceed from a list to a single entity.

Response format

All read tools return the response in the same format:

```
{
  "results": [ ... ],
  "count": 2,
  "truncated": true,
  "truncated_by": "count",
  "limit": 2,
  "note": "The result count limit was reached ..."
}
```

A response is limited by the limit argument (default: 100, maximum: 1000) and by the serialized size of the response (MaxResponseSize, default: 64K). The response is stopped when either of the limits is reached. The response is truncated only between results, never inside a result, and a truncated response indicates which limit was reached. Every tool has a fixed sort order, so a truncated response is always the beginning of the full result set. Paging is not supported; to get other results, narrow the query.

Problem tools

All problem tools call event.acknowledge with the corresponding action flags. The eventids argument is required; event IDs can be retrieved with problem_get. A message (comment) can be added with any problem tool. The message argument is optional, except for problem_annotate, where it is required; if specified, the message is recorded in the same Zabbix API call.

Tool	Effect	Additional arguments
problem_acknowledge	Acknowledge the problem.	-
problem_unacknowledge	Remove the acknowledgement.	-
problem_close	Close the problem manually (only for triggers that allow manual closing).	-
problem_annotate	Record a message.	message (required)
problem_change_severity	Change the severity of the problem.	severity (0-5)
problem_suppress	Suppress the problem until a specified time or indefinitely.	suppress_until
problem_unsuppress	Remove the suppression.	-

Problems cannot be created or deleted; all problem tools update existing problems.

Maintenance tools

Tool	Zabbix API method	Arguments
maintenance_create	maintenance.create	name, groupids and/or hostids, active_since, active_till, timeperiods, description, tags, tags_evaltype
maintenance_update	maintenance.update	maintenanceid and any of the arguments above. List arguments replace the whole set of values; an empty list clears the set.
maintenance_delete	maintenance.delete	maintenanceids

When a maintenance period is updated, list arguments (such as hostids, groupids and timeperiods) replace the existing values instead of adding to them. For example, if a maintenance period covers hosts A and B, and maintenance_update is called with hostids containing only host C, the maintenance period will cover only host C.

maintenance_create is the only operation that is not idempotent, and maintenance_delete is the only operation that cannot be reversed. To prevent them from being used, deny these tools with the DenyTool parameter.

Behavior of write tools

Problem and maintenance tools are *write tools*. They work as follows:

- The arguments are validated before any Zabbix API call; if the validation fails, the error names the constraint that was not met.
- If the client supports it, the change is applied only after it is approved by the user.
- The change is performed by Zabbix using the token of the caller. Zabbix checks the permissions and records the change in the audit log against the token owner.
- Zabbix errors, including permission errors, are returned unchanged as an error object with the code, message and data fields.
- Write tools return the Zabbix result directly; the response limits of read tools do not apply.
- Write calls are never retried by the MCP server.

Resources

Resources use the zabbix://<kind>/<id> URI scheme. Only zabbix://about is included in the resources/list response; entity IDs cannot be enumerated and are obtained from the results of tools. The resources/templates/list response lists the kinds of entities.

URI	Content
zabbix://about	The API URL, the frontend URL, the Zabbix version, the identity to which the token of the caller resolves, and the list of exposed tools.
zabbix://host/<hostid>	The host record without items, triggers, interfaces and macros.
zabbix://hostgroup/<groupid>	The host group record without the list of members.
zabbix://template/<templateid>	The template record without items, triggers and discovery rules.
zabbix://templategroup/<groupid>	The template group record without the list of members.
zabbix://item/<itemid>	The item definition without the current value.
zabbix://trigger/<triggerid>	The trigger definition without hosts, dependencies, functions and tags.

A resource whose corresponding tool is denied does not exist. A nonexistent entity and an entity that the token owner has no access to are both reported as *resource not found*; these cases are deliberately indistinguishable. The zabbix://about resource is not affected by the tool access rules.

Environment variables

Overview Environment variables allow configuring Zabbix components without hardcoding values in configuration files. This makes it easy to manage configurations in dynamic environments, such as Docker, where variables can be passed at runtime to adapt to different setups.

In the simplest case, by setting the Zabbix server DebugLevel configuration parameter value to an environment variable, you can then use it to configure the server on startup:

```
### Zabbix server configuration file:
DebugLevel=${NEW_DEBUG_LEVEL}

### Starting Zabbix server:
NEW_DEBUG_LEVEL=5 /usr/sbin/zabbix_server
```

Environment variables are supported by the following Zabbix components:

- Server
- Proxy
- Agent (UNIX or Windows)
- Agent 2 (UNIX or Windows), including plugins
- Web service
- Zabbix sender (when using the -c, --config option)

Important notes

- When a configuration parameter is set to an environment variable, which is not specified when running the component, the default value of the parameter is used.
- When using runtime commands (e.g., to increase the agent log level), any previously used environment variables must be specified. This is because Zabbix components use their configuration file to execute runtime commands; if environment variables are omitted, configuration parameter default values will be used. See Examples.
- The `userparameter_reload` runtime command does not support reloading environment variables. During reload, variables are ignored, and only parameters with regular values are reloaded.
- Process current environment variables, which were used in configuration files, are cleared after the Zabbix component is started. This ensures that child processes (e.g., remote scripts executed by Zabbix) cannot access these variables. However, note that process initial variables can still be retrieved (e.g., via the `/proc/<PID>/environ` file).

Syntax Environment variables must use the following syntax: `${alphanumerics/underscores}`.

The variable name may only include letters (a-z, A-Z), underscores (_), and digits (0-9), and must not begin with a digit.

Variables that do not match the required syntax or are combined with a regular value will be treated as regular values, which may produce errors.

Correct variable syntax:

```
DebugLevel=${NEW_DEBUG_LEVEL}
Hostname=${ZBX_HOSTNAME}
LogFile=${LogFile_001}
```

Incorrect variable syntax:

```
DebugLevel=${5_DebugLevel}
Hostname=${ZBX.HOSTNAME 1}
LogFile=/${HOME}/zabbix/zabbix_server.log
```

> **Note:**
> In Windows, environment variable names are case-insensitive.

Examples The following examples show how to configure and use environment variables with Zabbix components.

Example 1: Configuring and testing Zabbix agent

1. Set environment variables in the agent configuration file:

```
Hostname=${ZBX_HOSTNAME}
ServerActive=${ServerActive}
```

2. Test the configuration file:

```
ZBX_HOSTNAME="New Zabbix agent" ServerActive=127.0.0.1 /usr/sbin/zabbix_agentd -c /etc/zabbix/zabbix_agent
```

3. Start agent with environment variables:

```
ZBX_HOSTNAME="New Zabbix agent" ServerActive=127.0.0.1 /usr/sbin/zabbix_agentd -c /etc/zabbix/zabbix_agent
```

When using runtime commands (e.g., to increase the agent log level), any previously used environment variables must be specified:

```
ZBX_HOSTNAME="New Zabbix agent" ServerActive=127.0.0.1 /usr/sbin/zabbix_agentd -c /etc/zabbix/zabbix_agent
```

This is because agent uses its configuration file to execute runtime commands; if environment variables are omitted, configuration parameter default values will be used.

Alternatively, after setting environment variables in the agent configuration file, you may make them available to processes (e.g., by using the `export` command). This reduces the risk of unexpected behavior due to missing or incorrectly set variables.

```
export ZBX_HOSTNAME="New Zabbix agent"
export ServerActive=127.0.0.1
/usr/sbin/zabbix_agentd -c /etc/zabbix/zabbix_agentd.conf --test-config
/usr/sbin/zabbix_agentd -c /etc/zabbix/zabbix_agentd.conf
/usr/sbin/zabbix_agentd -c /etc/zabbix/zabbix_agentd.conf -R log_level_increase
```

Example 2: Configuring Zabbix agent for a container

If you're creating and configuring your own custom image for Zabbix components (e.g., Zabbix agent), you can define configuration parameters using environment variables and then start the container with those variables.

1. When preparing the image, set environment variables in the agent configuration file:

```
Hostname=${ZBX_HOSTNAME}
BufferSize=${BUFSZ}
ListenPort=${LISTENPORT}
UserParameter=${_UsrPar01}
UserParameter=${_UsrPar02}
```

2. After building the container image, start the agent container (e.g., Docker) with environment variables:

```
docker run --name my-zabbix-agent -e ZBX_HOSTNAME="new-hostname" -e BUFSZ=1000 -e LISTENPORT=20050 -e _Usr
```

3. When using runtime commands (e.g., to increase the agent log level), access the container shell and execute the runtime command:

```
docker exec -it <containerid> sh
/usr/sbin/zabbix_agentd -R log_level_increase
```

> **Attention:**
> The `userparameter_reload` runtime command does not support reloading environment variables. During reload, variables are ignored, and only parameters with regular values are reloaded.

Configuration

Overview Configuration in Zabbix involves specifying which hosts and systems to monitor, defining what data to collect, and establishing how notifications are delivered when problems occur.

Use the sidebar to navigate the relevant configuration aspects.

Key configuration tasks

- Set up a new host or host group → Hosts and host groups
- Define what data to collect from hosts → Items
- Create rules to detect problems → Problem detection with triggers
- Configure how notifications are delivered → Notifications upon events
- Display monitoring data through dashboards, graphs, and maps → Visualization
- Reuse configurations across multiple hosts → Templates and template groups, Templates out of the box
- Control user access and permissions → Users and user groups

Configuration best practices Configuration choices affect performance and maintainability. While every Zabbix setup is different, these best practices can be used as a foundation to adjust as needed.

Separate host for each monitoring instance

A host in Zabbix is a logical entity rather than a physical machine or device. Create a separate host for each independent instance such as a virtual machine, database, container, or network switch.

This approach provides:

1. Cleaner monitoring data with separate items, triggers, and alert notifications for each host.
2. Fine-grained user-access control via user-roles aligned with the principle of least privilege.

Use templates and mass update for similar hosts

For similar hosts like *Network switch 1* and *Network switch 2*, you can create a template with the required metrics (*Network switch template*) and group the hosts into a host group (*Network switches*).

When adding new switches, you can use mass update in *Data Collection* > *Hosts* to link the template instead of cloning hosts manually.

Configure master and dependent items to minimize requests

To minimize requests to the target entity, configure a master item to collect data in a single request, then use dependent items with preprocessing to extract specific values.

For example, the master item might collect a JSON or XML response with multiple metrics or execute a database query returning multiple columns (e.g., number of open connections, aborted connections, maximum allowed simultaneous connections, and total cumulative connections since startup), and the dependent items will parse and store each value separately.

Discard the master item's history after collection and only keep only the dependent items' data.

Deploy a Zabbix proxy when needed

If all your hosts are on the same local network as the Zabbix server with no scalability or performance concerns, you may not need a proxy. In larger or more complex environments, a proxy allows for more even load distribution.

Add a Zabbix proxy when:

1. Monitoring multiple hosts using various metric collection methods behind a firewall. The proxy gathers data from hosts and forwards it to Zabbix server, reducing the number of open firewall ports.
2. Monitoring remote locations, branches, or networks. Zabbix proxies in remote locations continue data collection during network interruptions and send the data back when the connection is restored.
3. Managing large-scale deployments. If you have many hosts and/or collect a high number of values per second, deploy a proxy to reduce server load and improve performance.

Configure independent secret macro retrieval

You may want to use secret user macros as either secret text or secret vault macros.

By default, secret macro values are retrieved by Zabbix server and propagated to proxies. For enhanced security with secret vault macros, configure the macro values to be retrieved independently by Zabbix server and proxies.

Configuring a template

Overview

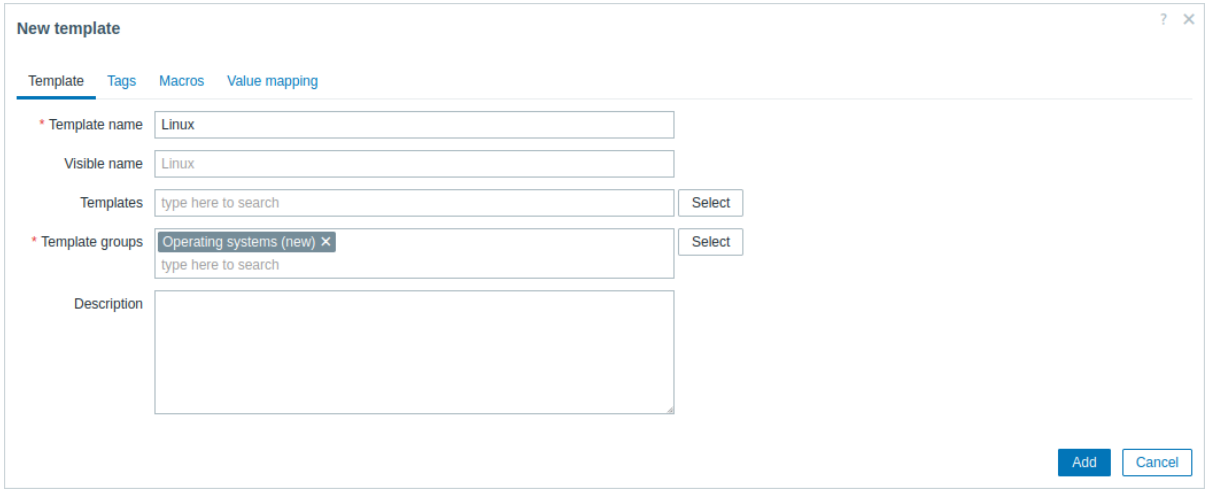
Configuring a template requires that you first create a template by defining its general parameters and then you add entities (items, triggers, graphs, etc.) to it.

Creating a template

To create a template, do the following:

1. Go to *Data collection* > *Templates*.
2. Click *Create template*.
3. Edit template attributes.

The **Template** tab contains general template attributes.

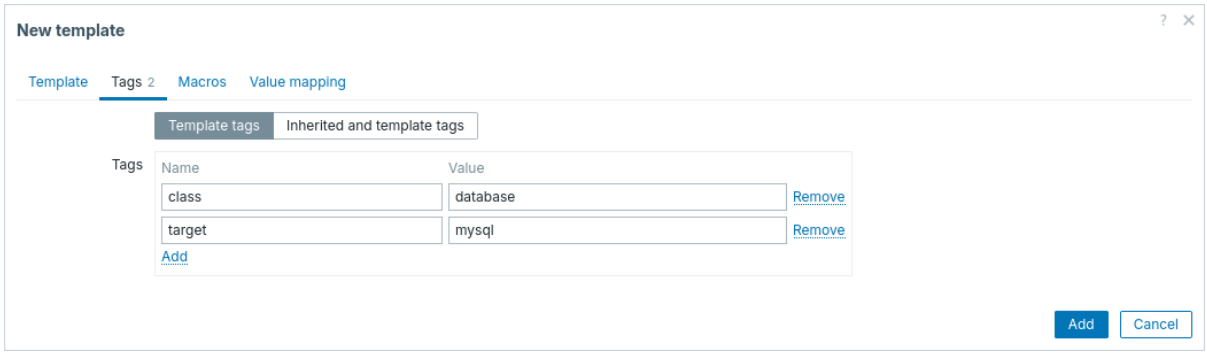


All mandatory input fields are marked with a red asterisk.

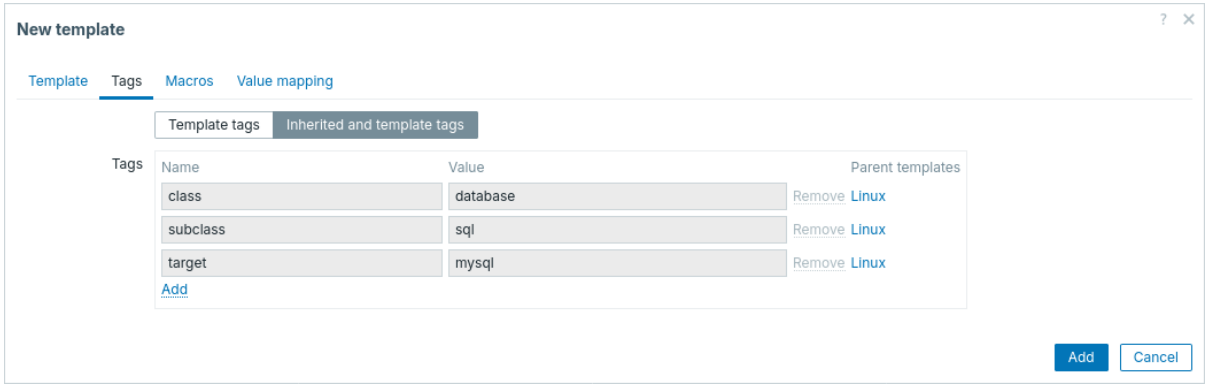
Template attributes:

Parameter	Description
Template name	Unique template name. Alphanumerics, spaces, dots, dashes, and underscores are allowed. Leading and trailing spaces are not allowed.
Visible name	If you set this name, it will be the one visible in lists, maps, etc.
Templates	Link one or more templates to this template. All entities (items, triggers, etc.) will be inherited from the linked templates. To link a new template, type its name in the *Templates* field—this will display a drop-down list of matching templates. Alternatively, click *Select* to open the *Templates* pop-up, which initially shows no templates. To populate the list, either type a template group name in the *Template group* field (and select one from the suggested list) or click the *Select* button next to the *Template group* field to open a *Template groups* pop-up. Once you choose a template group, the *Templates* pop-up is updated to show the templates belonging to that group. The selected templates will be linked when the configuration is saved or updated. To unlink a template, use one of the two options in the *Templates* block: *Unlink* - unlink the template, but preserve its entities (items, triggers, etc.); *Unlink and clear* - unlink the template and remove all of its entities (items, triggers, etc.).
Template groups	Template groups the template belongs to.
Description	Template description.
Vendor and version	Template vendor and version; displayed only when updating existing templates (out-of-the-box templates provided by Zabbix, imported templates, or templates modified through the Template API) if the template configuration contains such information. Cannot be modified in Zabbix frontend. For out-of-the-box templates, version is displayed as follows: major version of Zabbix, delimiter (”-”), revision number (increased with each new version of the template, and reset with each major version of Zabbix). For example, 8.0-0, 8.0-5, 8.2-0, 8.2-3.

The **Tags** tab allows you to define template-level tags. All problems of hosts linked to this template will be tagged with the values entered here.



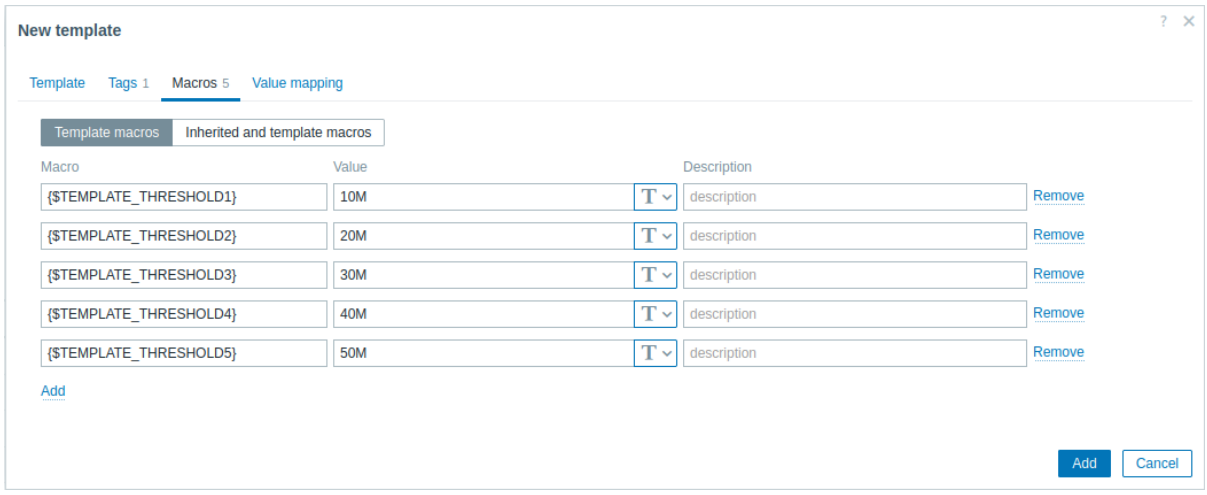
If you select the *Inherited and template tags* option, you may also view tags from linked templates that the template will inherit.



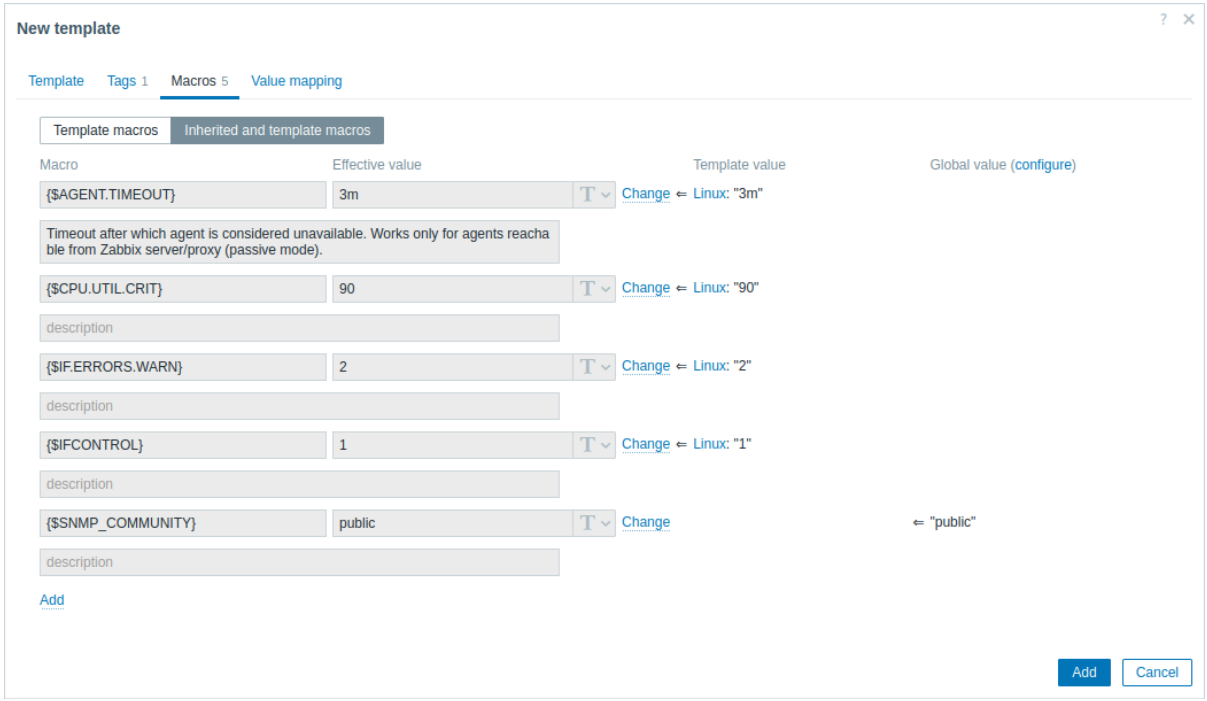
For convenience, links to the respective templates is provided.

User macros, {INVENTORY.*} macros, {HOST.HOST}, {HOST.NAME}, {HOST.CONN}, {HOST.DNS}, {HOST.IP}, {HOST.PORT} and {HOST.ID} macros are supported in tags.

The **Macros** tab allows you to define template-level user macros as a name-value pairs. Note that macro values can be kept as plain text, secret text, or Vault secret. Adding a description is also supported.



If you select the *Inherited and template macros* option, you may also view macros from linked templates and global macros that the template will inherit, as well as the values that the macros will resolve to.



For convenience, links to the respective templates, as well as a link to global macro configuration is provided. It is also possible to edit a linked template macro or global macro on the template level, effectively creating a copy of the macro on the template.

The **Value mapping** tab allows you to configure human-friendly representation of item data in value mappings.

Buttons:

Add	Add the template. The added template should appear in the list.
Update	Update the properties of an existing template.
Clone	Create another template based on the properties of the current template. This **includes** the entities (items, triggers, etc.) both inherited from linked templates and directly attached to the current template, but **excludes** the current template's vendor and version for the cloned template to be distinguishable from the original.
Delete	Delete the template; entities of the template (items, triggers, etc.) remain with the linked hosts.
Delete and clear	Delete the template and all its entities from linked hosts.
Cancel	Cancel the editing of template properties.

Adding items, triggers, graphs

> **Attention:**
> Items have to be added to a template first. Triggers and graphs cannot be added without the corresponding item.

You can add existing items to a template or create new items.

To add items to the template:

1. Go to *Data collection* > *Hosts* (or *Templates*).
2. Click *Items* in the row of the required host/template.
3. Mark the checkboxes of items you want to add to the template.
4. Click *Copy* below the item list.
5. Select the template (or group of templates) the items should be copied to and click *Copy*.
 All the selected items should be copied to the template.

Adding triggers and graphs is done in a similar fashion (from the list of triggers and graphs respectively), again, keeping in mind that they can only be added if the required items are added first.

Adding dashboards

To add dashboards to a template in *Data collection* > *Templates*, do the following:

1. Click *Dashboards* in the row of the template.
2. Configure a dashboard following the guidelines of configuring dashboards.

> **Attention:**
> When configuring widgets on a template dashboard (instead of a global dashboard), the host-related parameters are not available, and some parameters have a different label. This is because template dashboards display data only from the host that the template is linked to. For example, the parameters *Host groups*, *Exclude host groups* and *Hosts* in the *Problems* widget are not available, the parameter *Host groups* in the *Host availability* widget is not available, and the parameter *Show hosts in maintenance* is renamed to *Show data in maintenance*, etc. For more information on the availability of parameters in template dashboard widgets, see specific parameters for each dashboard widget.

> **Note:**
> For details on accessing host dashboards that are created from template dashboards, see the host dashboards section.

Configuring low-level discovery rules

See the low-level discovery section of the manual.

Adding web scenarios

To add web scenarios to a template in *Data collection* > *Templates*, do the following:

1. Click *Web* in the row of the template.
2. Configure a web scenario following the usual method of configuring web scenarios.

Configuring a template group

Overview

Template groups are used for the logical grouping of templates and assigning user permissions to them.

Each template must have at least one template group assigned. A template may belong to multiple template groups, and each template group may contain multiple templates.

Note that in Zabbix, all permissions are based on groups: user groups, host groups, and template groups. So, even if a single user needs access to a single template, it is granted by adding the user to a user group and assigning that user group permissions to the template group containing that template.

Configuration

> **Attention:**
> Only Super admin users can create template groups.

There are two options of creating a template group in Zabbix frontend.

Option one:

1. Go to *Data collection* > *Template groups*.
2. Click *Create template group* in the upper-right corner of the screen.
3. Enter the group name in the form.



Option two: when configuring a template, enter a non-existing group name in the *Template groups* input field.

Once the template group is created, you can click the template name in the list under *Data collection* > *Template groups* to edit the group name, clone the group, or delete the group.

Deleting a template group only deletes the logical group, not the templates in the group. It is not possible to delete a template group that is the only group for any existing template.

Creating template subgroups

A template subgroup (or nested template group) is a child of the parent template group that contains it.

A subgroup is created by using the forward slash '/' in the group name input field to denote its relation to the parent group(s). For example:

- inputting `Linux servers/Databases` creates the `Linux servers/Databases` subgroup of the parent group `Linux servers`.
- inputting `Linux servers/Databases/MySQL/Tokyo` creates the respective subgroup within the nested parent groups `Linux servers`, `Linux servers/Databases`, `Linux servers/Databases/MySQL`.

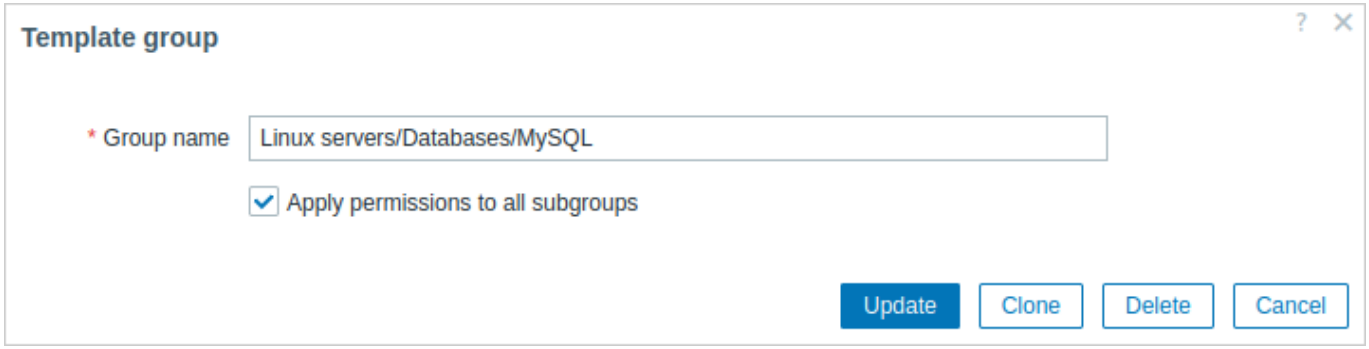
When creating a subgroup, using leading or trailing slashes, or several slashes in a row is not allowed. Escaping of '/' is not supported.

It is not required to create any parent template group(s) before creating a subgroup. You can choose whether to start by creating a subgroup (for example, `Linux servers/Databases`) or any parent template group(s) (in our example, `Linux servers`). If you start by creating a subgroup, parent template group(s) will **not** be created automatically.

Permissions to template groups

- When creating a subgroup to an existing parent template group (for example, creating `Linux servers/Databases` when `Linux servers` already exists), user group permissions to the subgroup are inherited from the parent.
- When creating a parent template group to an existing subgroup (for example, creating `Linux servers` when `Linux servers/Databases` already exists), no permissions to the parent are set.

When editing any template group, you can also set an additional option, *Apply permissions to all subgroups*.



Marking this checkbox and clicking *Update* will apply the same level of permissions to all current and future subgroups of the template group being edited.

So, if any user groups have been given varying permissions to the subgroups of the template group being edited, marking the checkbox will grant all current and future subgroups the same user permissions as the group being edited.

Note that this option is not saved in the database and will override existing permissions. Any changes made through this option can be reverted only manually.

Linking/unlinking

Overview

Linking is a process whereby templates are applied to hosts, whereas unlinking removes the association with the template from a host.

Linking a template

To link a template to the host, do the following:

1. Go to *Data collection* > *Hosts*.
2. Click the required host.
3. Start typing the template name in the *Templates* field. A list of matching templates will appear; scroll down to select. Alternatively, you may click *Select* next to the field and select one or several templates from the list in a popup window.
4. Click *Add/Update* in the host attributes form.

The host will now have all the entities of the template. This includes items, triggers, graphs, low-level discovery rules, web scenarios, as well as dashboards.

Attention:
Linking multiple templates to the same host will fail if those templates contain items with the same item key. And, as triggers and graphs use items, they cannot be linked to a single host from multiple templates either, if using identical item keys.

When entities (items, triggers, etc.) are added from the template:

- previously existing identical entities on the host are updated as entities of the template, and **any existing host-level customizations to the entity are lost**;
- entities from the template are added;
- any directly linked entities that, prior to template linkage, existed only on the host remain untouched.

In the lists, all entities from the template now are prefixed by the template name, indicating that these belong to the particular template. The template name itself (displayed in gray) is a link that allows you to access the list of those entities at the template level.

Note:
For some items, such as external checks, HTTP agent checks, simple checks, SSH checks and Telnet checks, a host interface is optional. If, at the time of linking a template, the host does not have an interface defined these items will be added without an interface. If you add a host interface later it will not be assigned automatically to already existing items. To assign the newly added host interface to all template items at once, unlink the template from the host and then link it back again. To preserve item history use the option *Unlink*, do not use *Unlink and clear*.

If some entity is not prefixed by the template name, it means that it existed on the host before and was not added by the template.

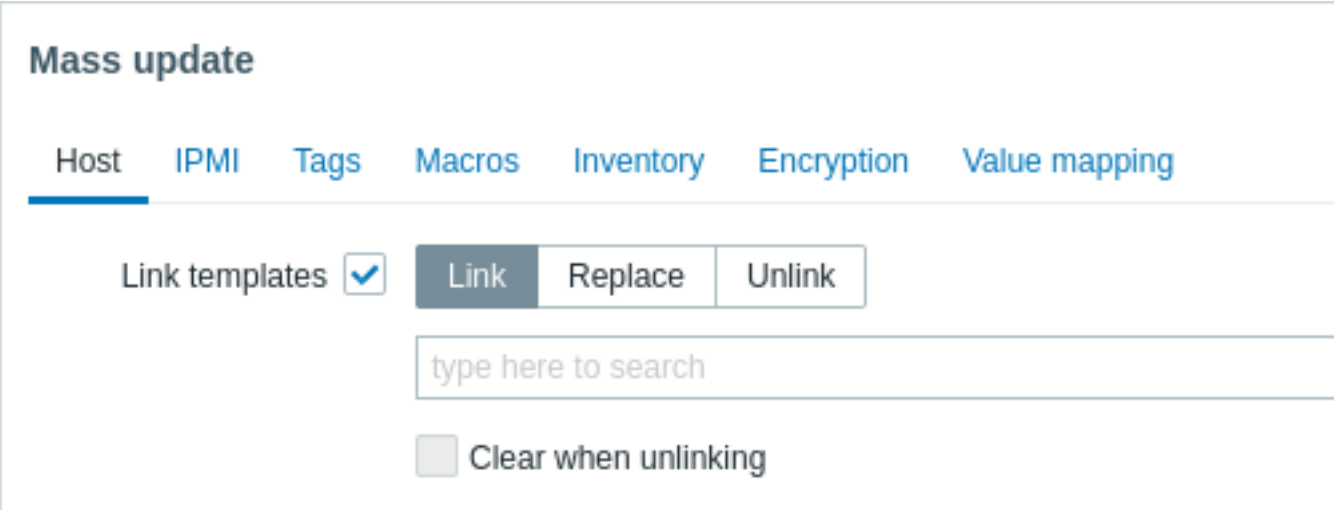
Entity uniqueness criteria

When adding entities (items, triggers, etc.) from a template it is important to know what of those entities already exist on the host and need to be updated and what entities differ. The uniqueness criteria for deciding upon the sameness/difference are:

- for items - the item key;
- for triggers - trigger name and expression;
- for custom graphs - graph name and its items.

Linking templates to several hosts

To update template linkage of many hosts, in *Data collection* > *Hosts* select some hosts by marking their checkboxes, then click **Mass update** below the list and then select *Link templates*:



To link additional templates, start typing the template name in the auto-complete field until a dropdown appears offering the matching templates. Just scroll down to select the template to link.

The *Replace* option allows you to link a new template while unlinking any template that was linked to the hosts before. The *Unlink* option allows you to specify which templates to unlink. The *Clear when unlinking* option allows you not only to unlink previously linked templates but also to remove all entities inherited from them (items, triggers, etc.).

Note:
Zabbix offers a sizable set of predefined templates. You can use these for reference, but beware of using them unchanged in production as they may contain too many items and poll for data too often. If you feel like using them, fine-tune them to fit your real needs.

Editing linked entities

If you try to edit an item or a trigger that was linked from the template, you may realize that many key options are disabled for editing. This makes sense as the idea of templates is that things are edited in a one-touch manner on the template level. However, you still can, for example, enable/disable an item on individual hosts and set the update interval, history length and some other parameters.

If you want to edit the entity fully, you have to edit it on the template level (template level shortcut is displayed in the form name), keeping in mind that these changes will affect all hosts that have this template linked to them.

Attention:
Any customizations to the entities implemented on a template-level will override the previous customizations of the entities on a host-level.

Unlinking a template

To unlink a template from a host, do the following:

1. Go to *Data collection* > *Hosts*.
2. Click the required host and find the *Templates* field.
3. Click *Unlink* or *Unlink and clear* next to the template to unlink the template.
4. Click *Update* in the host attributes form.

Choosing the *Unlink* option will simply remove association with the template, while leaving all its entities with the host. This includes items, triggers, graphs, low-level discovery rules, and web scenarios, but excludes dashboards. Note that value maps and

tags inherited from linked templates will also be removed.

Choosing the *Unlink and clear* option will remove both the association with the template and all its entities (items, triggers, etc.).

Nesting

Overview

Nesting is a way of one template encompassing one or more other templates.

As it makes sense to separate out entities on individual templates for various services, applications, etc., you may end up with quite a few templates all of which may need to be linked to quite a few hosts. To simplify the picture, it is possible to link some templates together in a single template.

The benefit of nesting is that you have to link only one template to the host, and the host will automatically inherit all entities from the templates that are linked to the one template. For example, if we link *T1* and *T2* to *T3*, we supplement *T3* with all entities from *T1* and *T2*, but not vice versa. If we link *T1* to *T2* and *T3*, we supplement *T2* and *T3* with entities from *T1*.

Configuring nested templates

To link templates, you need to take an existing template (or create a new one), and then:

1. Open the template configuration form.
2. Find the *Templates* field.
3. Click *Select* to open the *Templates* pop-up window.
4. In the pop-up window, choose the required templates, and then click *Select* to add the templates to the list.
5. Click *Add* or *Update* in the template configuration form.

Thus, all entities of the configured template, as well as all entities of linked templates will now appear in the template configuration. This includes items, triggers, graphs, low-level discovery rules, and web scenarios, but excludes dashboards. However, linked template dashboards will, nevertheless, be inherited by hosts.

To unlink any of the linked templates, click *Unlink* or *Unlink and clear* in the template configuration form, and then click *Update*.

The *Unlink* option will simply remove the association with the linked template, while not removing any of its entities (items, triggers, etc.).

The *Unlink and clear* option will remove both the association with the linked template, as well as all its entities (items, triggers, etc.).

Mass update

Overview

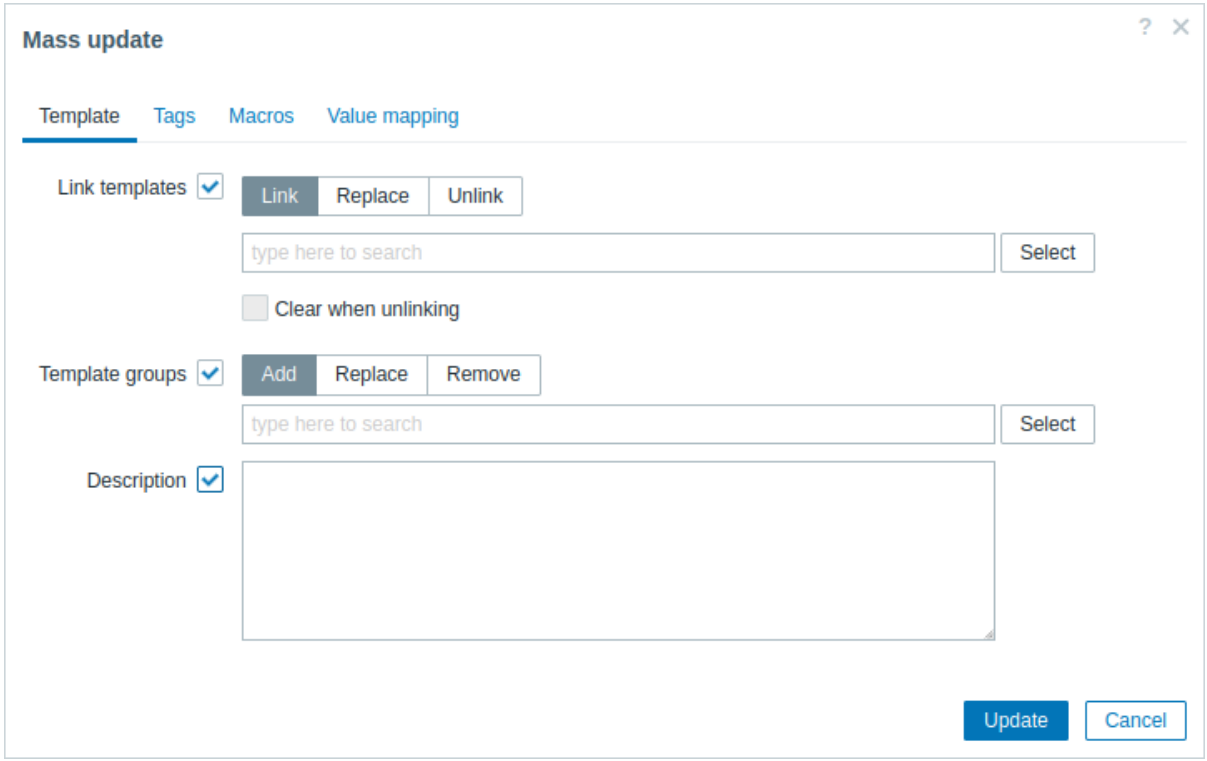
Sometimes you may want to change some attribute for a number of templates at once. Instead of opening each individual template for editing, you may use the mass update function for that.

Using mass update

To mass-update some templates, do the following:

1. Mark the checkboxes before the templates you want to update in the template list.
2. Click *Mass update* below the list.
3. Navigate to the tab with required attributes (*Template*, *Tags*, *Macros* or *Value mapping*).
4. Mark the checkboxes of any attribute to update and enter a new value for them.

The **Template** tab contains general template mass update options.



The following options are available when selecting the respective button for the *Link templates* update:

- *Link* - specify which additional templates to link;
- *Replace* - specify which templates to link while at the same time unlinking any previously linked templates;
- *Unlink* - specify which templates to unlink.

To specify the templates to link/unlink, start typing the template name in the auto-complete field until a dropdown appears offering the matching templates. Just scroll down to select the required templates.

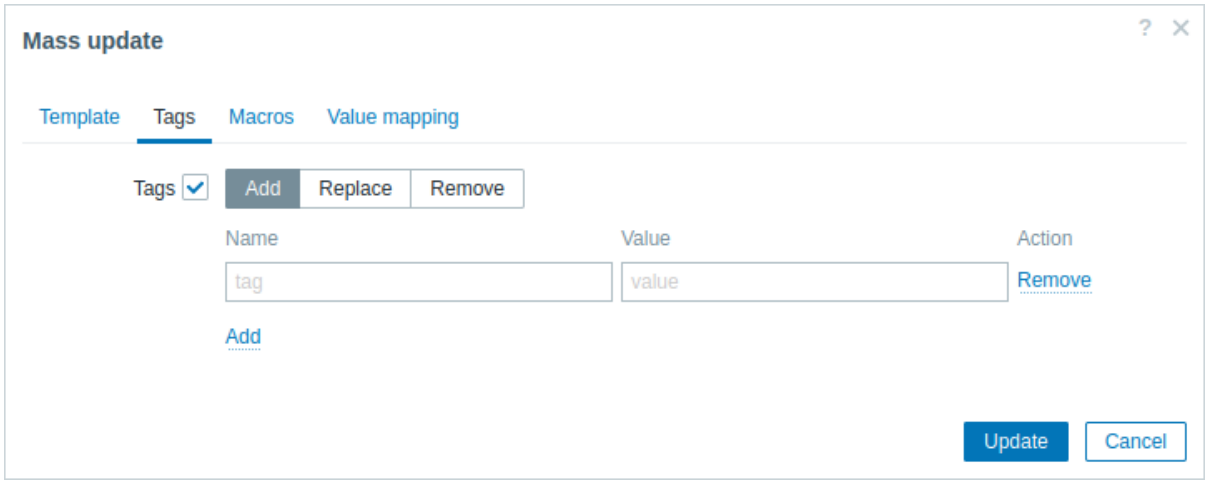
The *Clear when unlinking* option allows you not only to unlink previously linked templates but also to remove all entities inherited from them (items, triggers, graphs, etc.).

The following options are available when selecting the respective button for the *Template groups* update:

- *Add* - specify additional template groups from the existing ones or enter completely new template groups for the templates;
- *Replace* - remove the template from any existing template groups and replaces them with the one(s) specified in this field (existing or new template groups);
- *Remove* - remove specific template groups from templates.

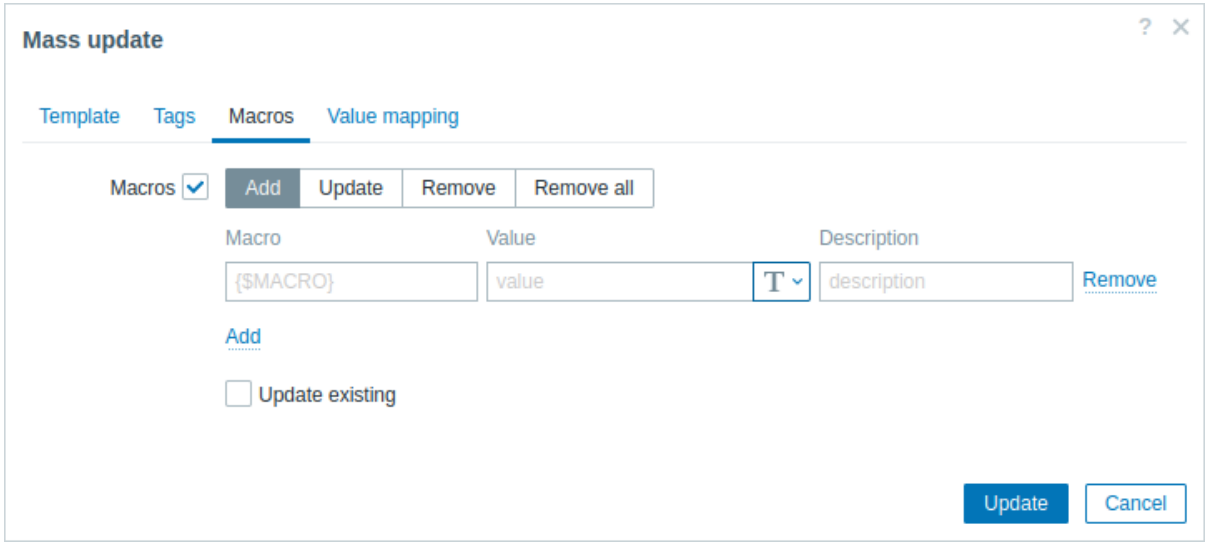
These fields are auto-complete - starting to type in them offers a dropdown of matching template groups. If the template group is new, it also appears in the dropdown and it is indicated by *(new)* after the string. Just scroll down to select.

The **Tags** tab allows you to mass update template-level tags.



User macros, {INVENTORY.*} macros, {HOST.HOST}, {HOST.NAME}, {HOST.CONN}, {HOST.DNS}, {HOST.IP}, {HOST.PORT} and {HOST.ID} macros are supported in tags. Note that tags with the same name, but different values are not considered 'duplicates' and can be added to the same template.

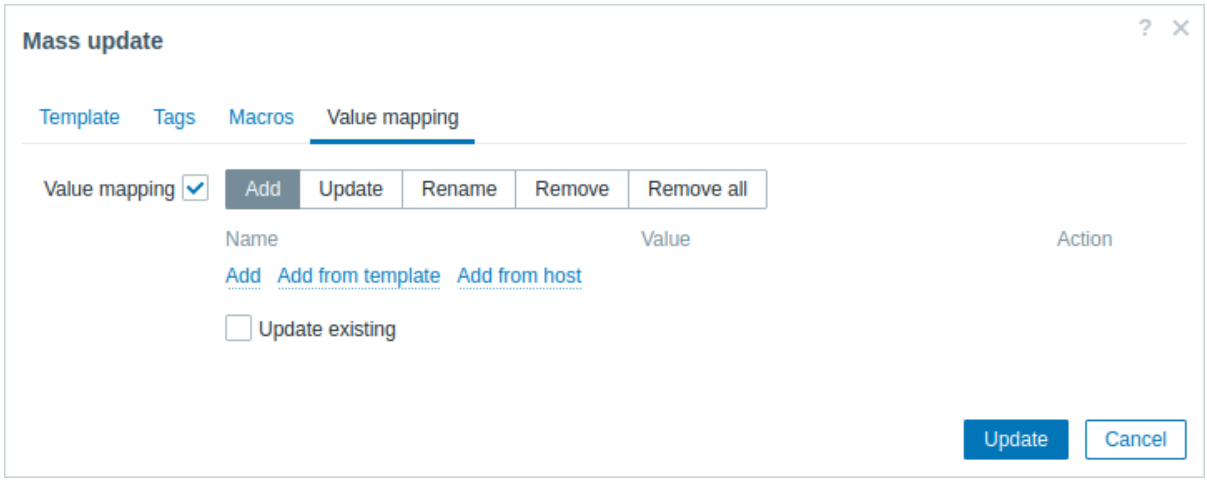
The **Macros** tab allows you to mass update template-level macros.



The following options are available when selecting the respective button for macros update:

- *Add* - specify additional user macros for the templates. If the *Update existing* checkbox is checked, value, type and description for the specified macro name will be updated. If unchecked, if a macro with that name already exist on the template(s), it will not be updated.
- *Update* - replace values, types and descriptions of macros specified in this list. If the *Add missing* checkbox is checked, macro that didn't previously exist on a template will be added as new macro. If unchecked, only macros that already exist on a template will be updated.
- *Remove* - remove specified macros from templates. If the *Except selected* checkbox is checked, all macros except specified in the list will be removed. If unchecked, only macros specified in the list will be removed.
- *Remove all* - remove all user macros from templates. If the *I confirm to remove all macros* checkbox is not checked, a new popup window will open asking to confirm removal of all macros.

The **Value mapping** tab allows you to mass update value mappings.



Buttons with the following options are available for value map update:

- *Add* - add value maps to the templates. If you mark *Update existing*, all properties of the value map with this name will be updated. Otherwise, if a value map with that name already exists, it will not be updated.
- *Update* - update existing value maps. If you mark *Add missing*, a value map that didn't previously exist on a template will be added as a new value map. Otherwise, only the value maps that already exist on a template will be updated.
- *Rename* - give new name to an existing value map.
- *Remove* - remove the specified value maps from the templates. If you mark *Except selected*, all value maps will be removed **except** the ones that are specified.
- *Remove all* - remove all value maps from the templates. If the *I confirm to remove all value maps* checkbox is not marked, a new popup window will open asking to confirm the removal.

The *Add from template* and *Add from host* options are available for value mapping add/update operations. They allow you to select value mappings from a template or a host respectively.

When done with all required changes, click *Update*. The attributes will be updated accordingly for all the selected templates.

Hosts and host groups

What is a ”host”?

In Zabbix, a ”host” refers to any physical or virtual device, application, service, or any other logically-related collection of monitored parameters.

Creating hosts is one of the first monitoring tasks in Zabbix. For example, if you want to monitor some parameters on a server ”x”, you must first create a host called, say, ”Server X” and then you can look to add monitoring items to it.

Hosts are organized into host groups.

Proceed to:

- Host Wizard for a guided, step-by-step interface for creating and configuring a new or existing host.
- Configuring a host for a classic approach to creating and configuring a host.

Host Wizard

Overview

The Host Wizard is a guided, step-by-step interface for setting up your monitoring target (device, application, service, etc.) in Zabbix. It walks you through the following steps:

- Selecting a template
- Creating or selecting a host
- Installing Zabbix agent or agent 2
- Adding a host interface
- Applying additional configuration to your monitoring target or the Zabbix host (if required by the template)

To access the Host Wizard in Zabbix frontend, do the following:

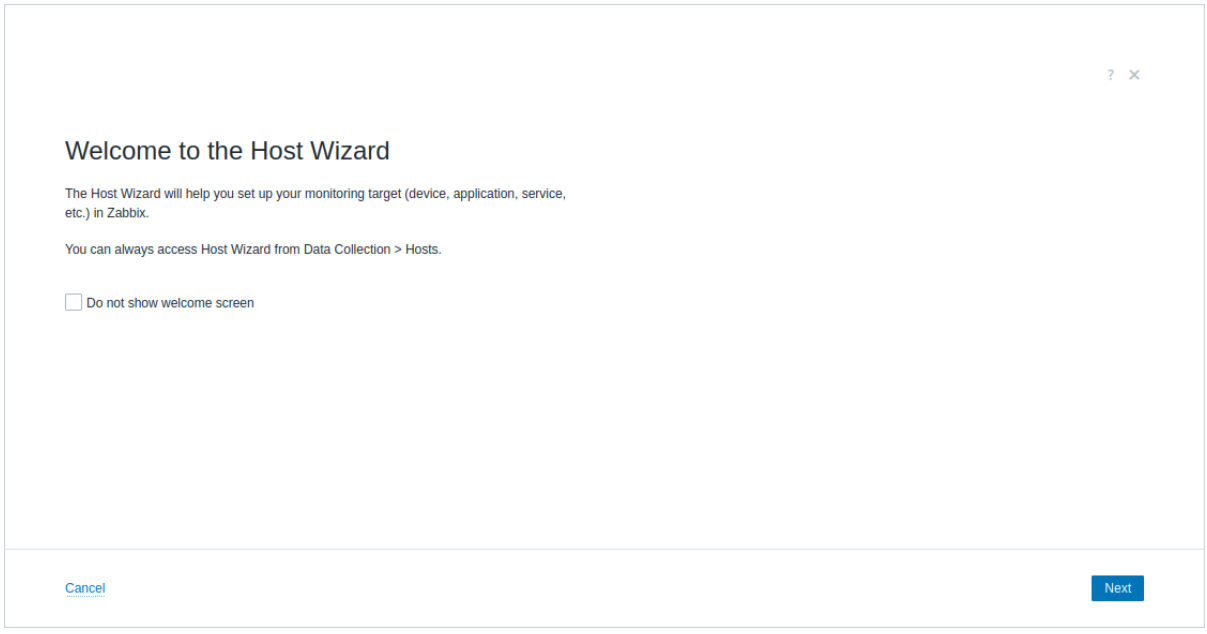
1. Go to: *Data collection* > *Hosts*.
2. Click *Host Wizard* in the upper-right corner of the screen.
3. Follow the instructions in the wizard.

You can also edit an existing host by clicking its name (or the three dots icon next to it) in various frontend sections to access the Host Wizard from the host menu. Note that some steps are skipped when editing an existing host if it already includes the configuration required by the selected template.

Alternatively, you can use the classic approach to creating and configuring a host.

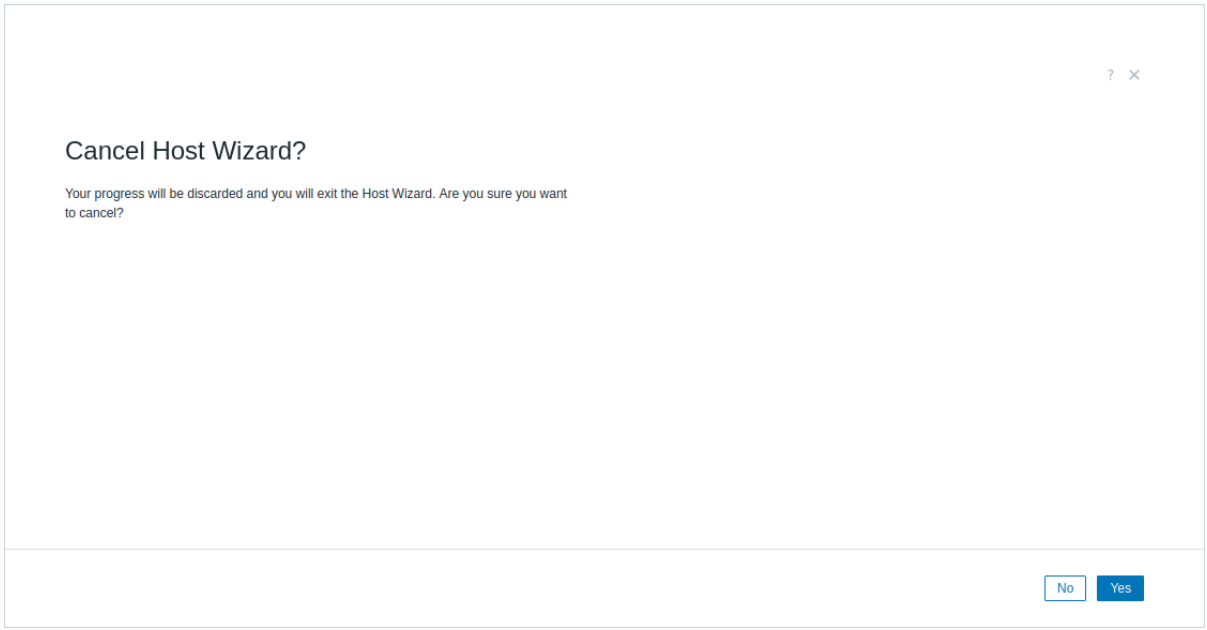
Welcome to the Host Wizard

When you start the Host Wizard, a welcome screen is displayed:



To skip this screen in future sessions, you can mark the *Do not show welcome screen* checkbox and click *Next*.

If you start configuring settings in the next Host Wizard steps and try to close the wizard by clicking the close icon or pressing Esc, a confirmation screen will appear:



Clicking *Yes* will exit the Host Wizard without saving your progress. Clicking *No*, Esc, or the close icon will return to the last step.

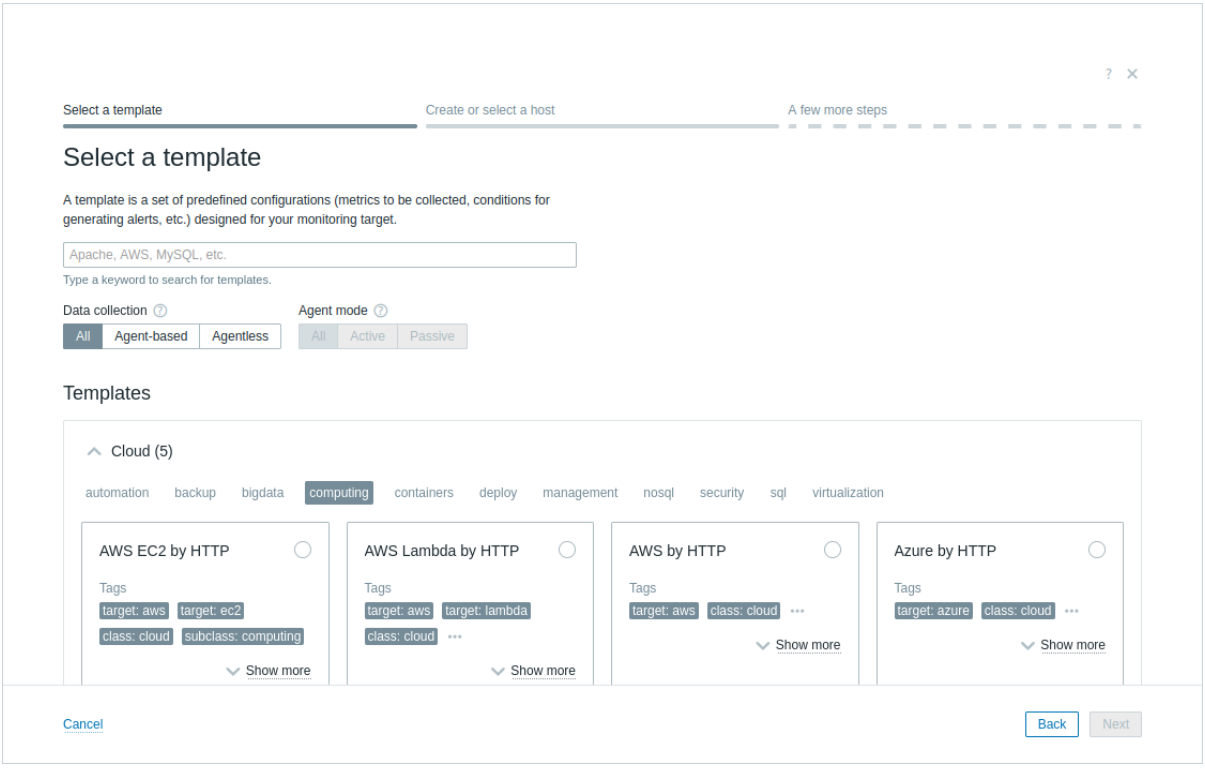
Select a template

The first step of setting up your monitoring target is to select a template, which is a set of predefined configurations (metrics to be collected, conditions for generating alerts, etc.) designed for your monitoring target.

Zabbix out-of-the-box templates offer a variety of predefined monitoring configurations for operating systems, applications, databases, network devices, services, and more.

In this step, you can:

- Browse templates by class (Cloud, Database, Network, etc.) and subclass (e.g., automation, backup), which are based on template tags.
- Search templates by keywords in template name or template tag names and values.
- Filter templates by data collection method (agent-based or agentless; agent-based templates include at least one Zabbix agent item).
- Filter templates by agent mode (active or passive; active templates include at least one Zabbix agent (active) item, while passive templates include at least one passive item; for details, see Passive and active agent checks).
- When using the Host Wizard to configure an existing host, you can also filter templates already linked to the host.



Attention:
When selecting a template, you may see the following message: *Some templates (n) are incompatible with the Host Wizard. See how to update them. Custom templates are not supported.* This indicates that some out-of-the-box templates are not yet compatible with the Host Wizard and need to be upgraded; see template upgrade instructions.

The template you select will determine the next steps in the Host Wizard. For example, if you select the *MySQL by Zabbix agent* template, which uses Zabbix agent for data collection, the wizard will guide you through the agent installation process.

Create or select a host

A template must be linked to a host, which is an entity in Zabbix that represents your monitoring target. When linked, the host receives all template entities, such as items (metrics to be collected) and triggers (conditions for generating alerts), that are required for monitoring.

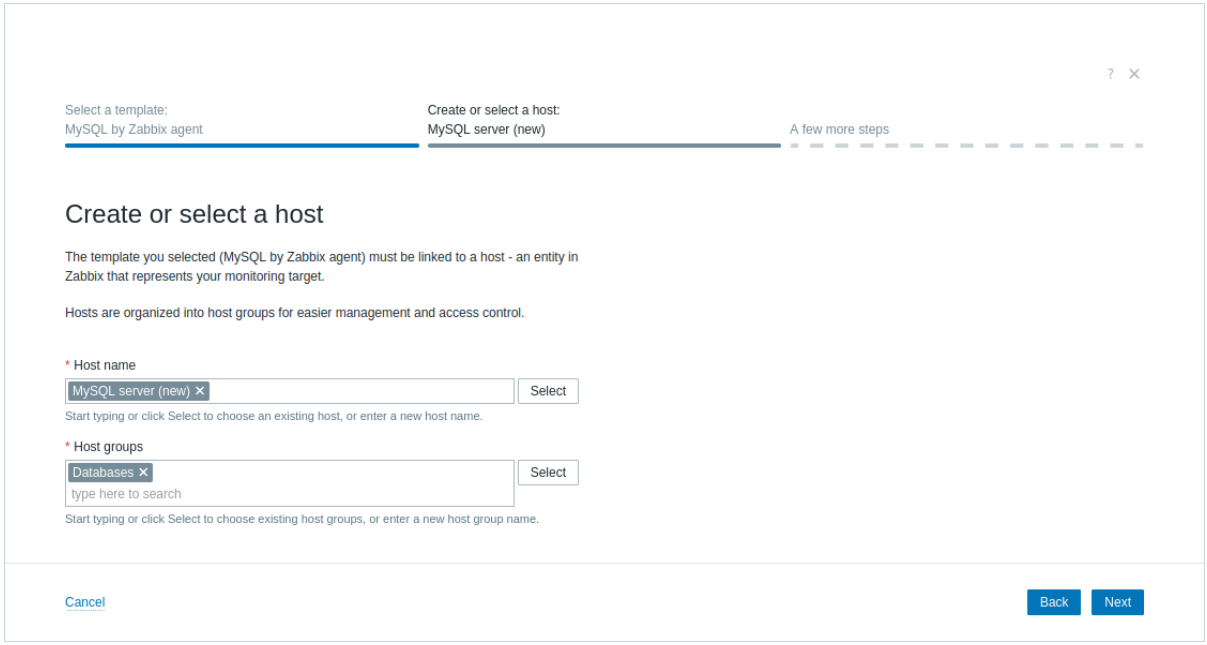
Each host must also belong to at least one host group, which is used for organizing hosts and assigning user permissions to them.

In this step, you can:

- Create a new host and host group by entering names for both.
- Create a new host by entering its name, and select an existing host group, to which the new host will be assigned.
- Select an existing host; you can also assign this host to additional host groups without removing it from any current groups.

Note:
Selecting an existing host does not remove or overwrite its current configuration unless stated explicitly in later steps. Discovered hosts cannot be selected.

For example, you could create a new host named *MySQL server* to represent the locally installed MySQL server that you want to monitor. The previously selected *MySQL by Zabbix agent* template will be linked to this host. Additionally, you could select the existing *Databases* host group (or create a new one, e.g., *MySQL servers*) to organize the host with other database-related hosts and later simplify permission management.



Install Zabbix agent

Agent-based templates (such as *MySQL by Zabbix agent*) require installing Zabbix agent or agent 2 - a process deployed on your monitoring target to actively monitor local resources and applications (if required by the template).

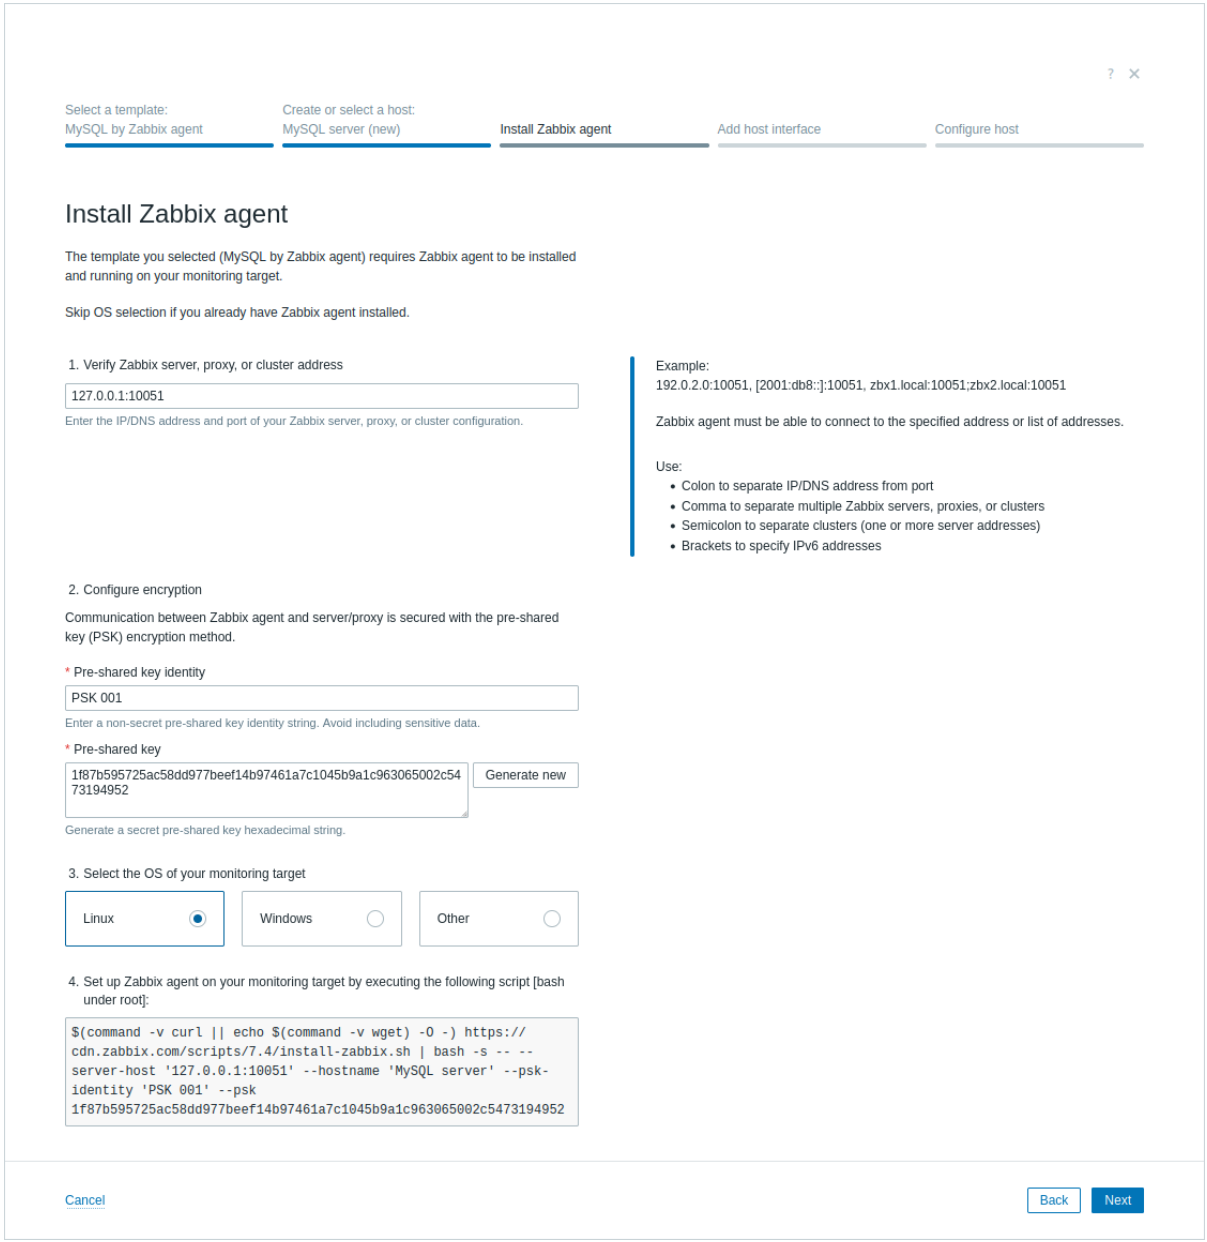
In this step:

- Verify Zabbix server, proxy, or cluster address (e.g., 192.0.2.1:10051). This value will be used to configure the Server and ServerActive parameter for Zabbix agent.
- Enter a new, unique, and non-secret pre-shared key identity (e.g., *PSK 001* or *mysql-agent-psk1*), and generate the pre-shared key. These values will be used to configure pre-shared key (PSK) encryption for the Zabbix host (e.g., *MySQL server*) and Zabbix agent.

Attention:
When editing an existing host, the PSK configuration will overwrite all existing encryption settings on the host.

- Select the operating system of the machine that hosts your monitoring target.
- Install Zabbix agent on that system by running the provided script on it or by following the installation instructions.

For example, you could set *Pre-shared key identity* to *PSK 001* and generate a new *Pre-shared key*. Then, you could select *Linux* as the operating system and run the provided script on that system. After installing Zabbix agent, return to the Host Wizard.

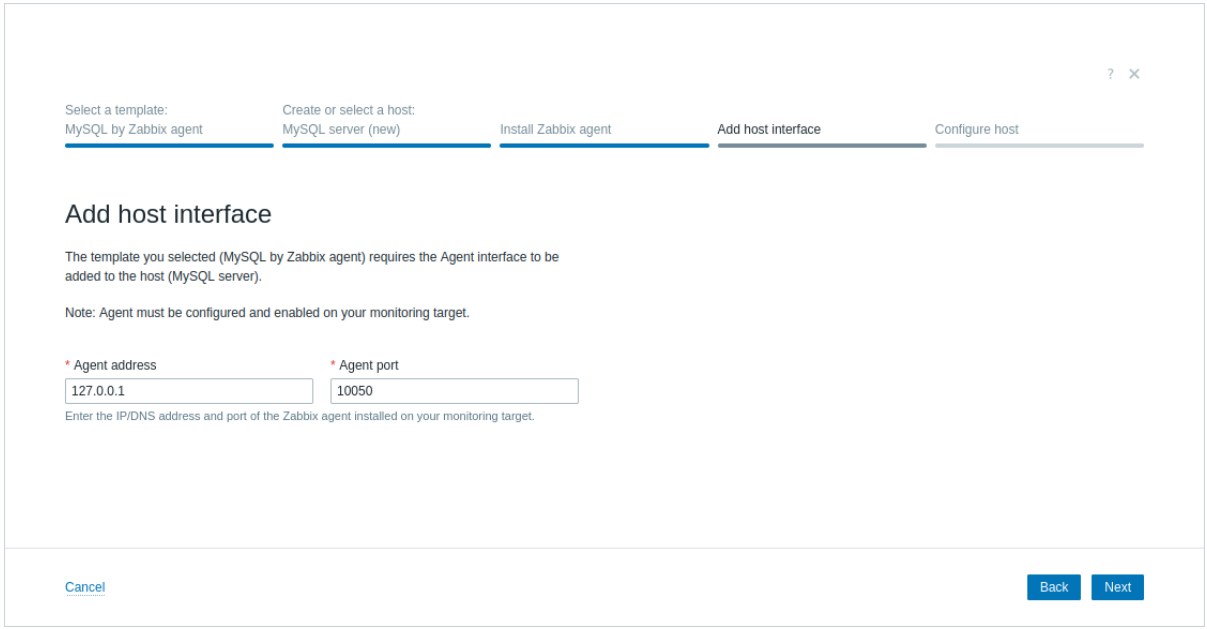


Add host interface

The host interface defines how Zabbix server connects to your monitoring target over the network. It defines connection parameters such as IP address, DNS name, port, and interface type (Agent, SNMP, JMX, or IPMI), depending on the data collection method required by the selected template.

In this step, enter the host interface details required by the selected template.

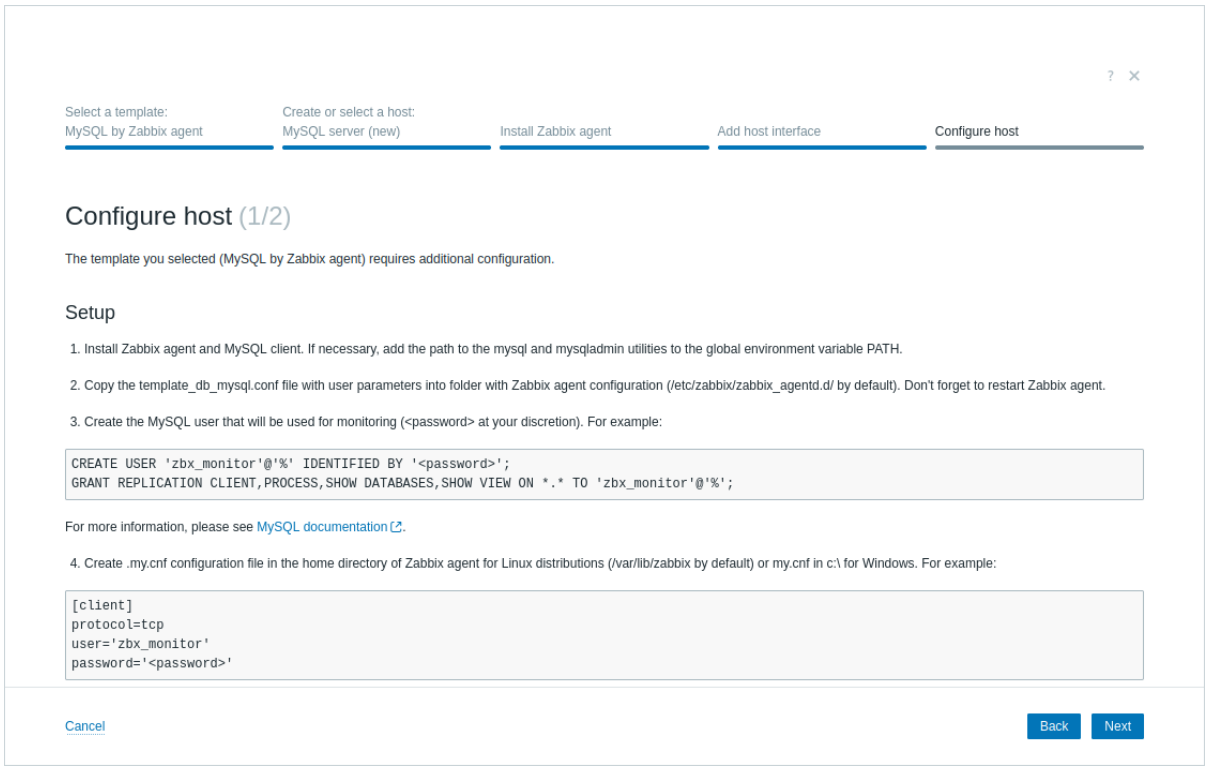
For example, you could use the default agent address (`127.0.0.1`) and port (10050) if Zabbix server, Zabbix agent, and the MySQL server are all running locally on the same machine.



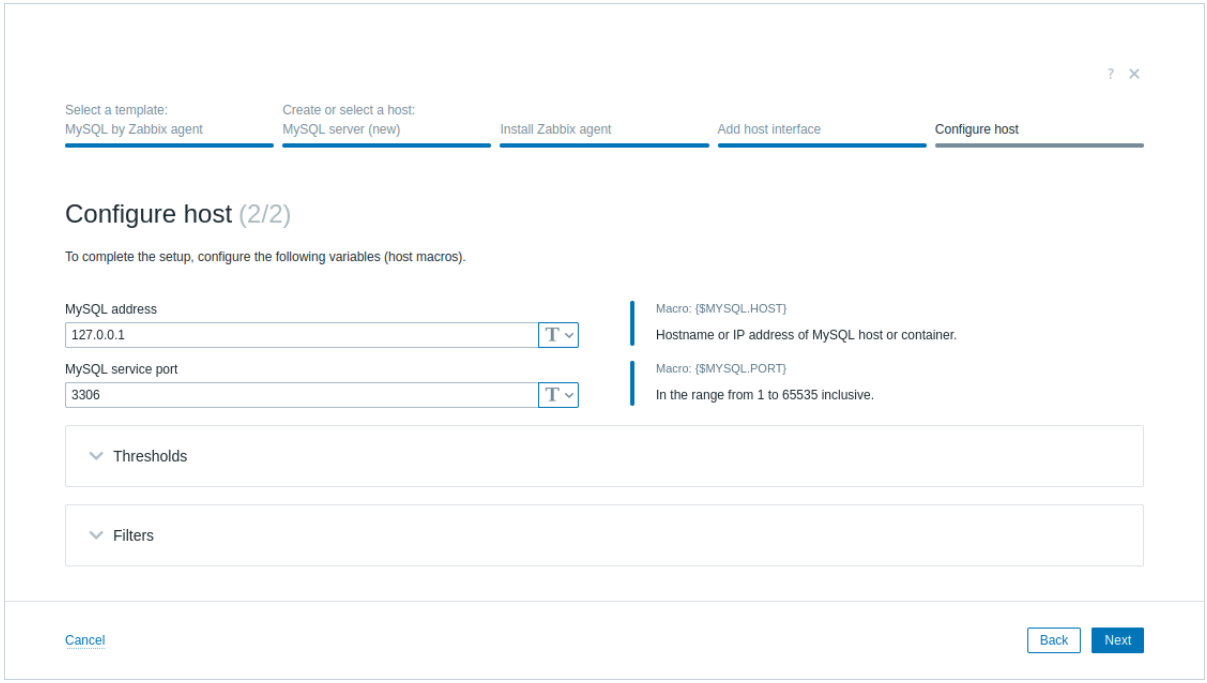
Configure host

Some templates require additional configuration before the host can be created. These may include:

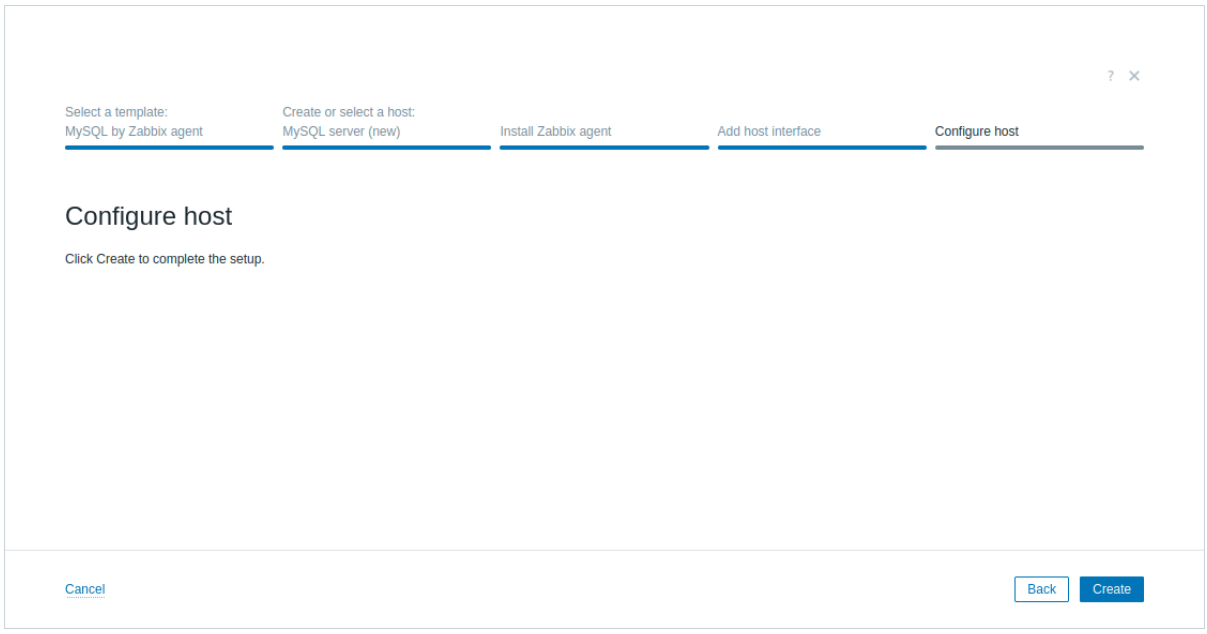
- Manual configuration of your monitoring target (e.g., enabling specific services, creating service users, or granting permissions):



- Defining host macros - variables that control item behavior, connection settings, authentication credentials, and other parameters:



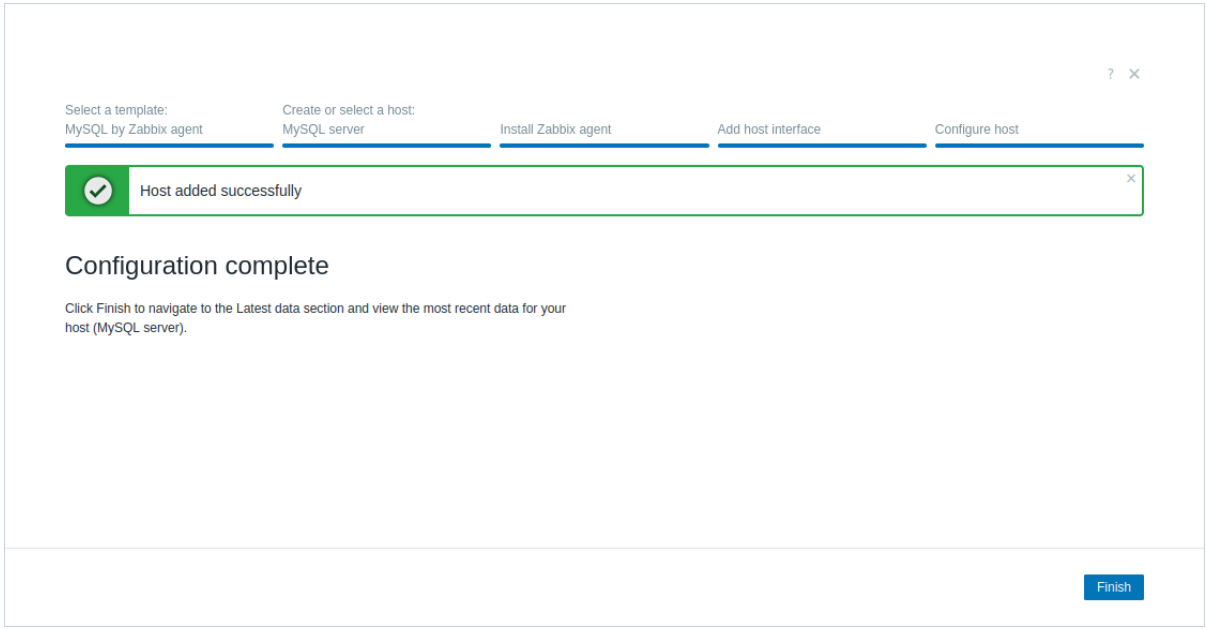
After completing all required configuration steps, click *Create* to add the host to Zabbix (or *Update* to update the configuration of an existing host):



Configuration complete

At this point, Zabbix is already monitoring your target.

For new hosts, click *Finish* to navigate to the Latest data section and view the most recent data for your host. For existing hosts, click *Finish* to close the Host Wizard and return to the screen where it was opened.



Configuring a host

Overview

To configure a host in Zabbix frontend, do the following:

1. Go to *Data collection* > *Hosts* or *Monitoring* > *Hosts*.
2. Click *Create host* in the upper-right corner of the screen (or the host name to edit an existing host).
3. Enter parameters of the host in the form.

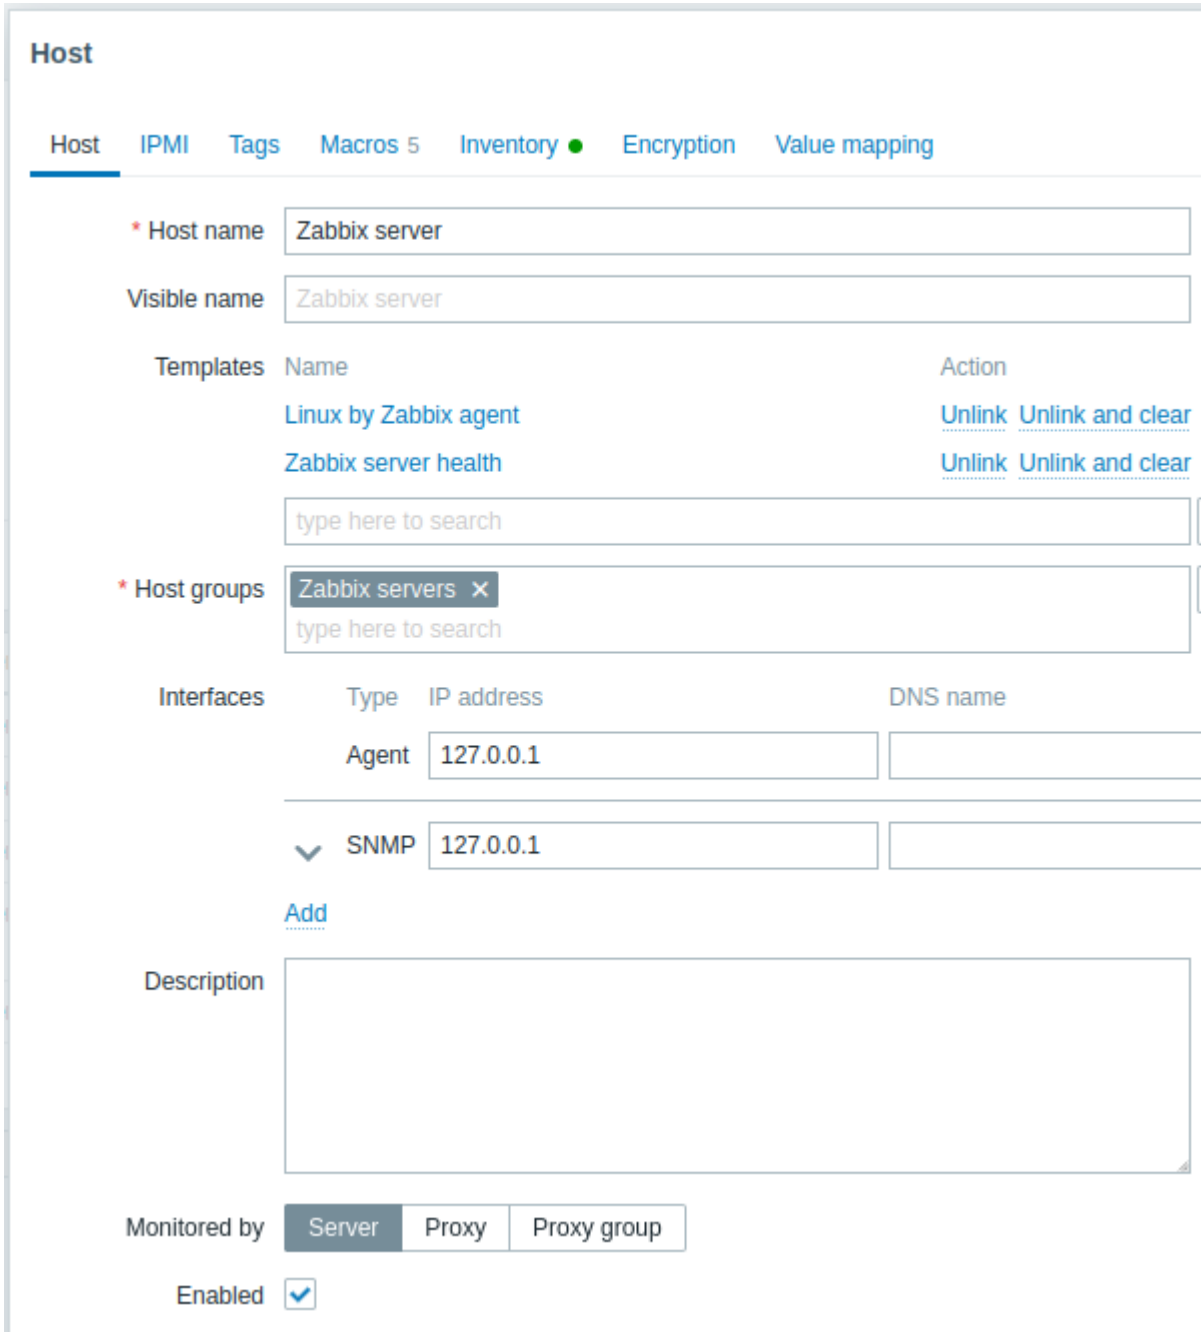
You can also use the *Clone* button in the configuration form of an existing host to create a new host. This host will have all of the properties of the existing host, including linked templates, entities (items, triggers, etc.) from those templates, as well as the entities directly attached to the existing host.

Note that when a host is cloned, it will retain all template entities as they are originally on the template. Any changes to those entities made on the existing host level (such as changed item interval, modified regular expression or added prototypes to the low-level discovery rule) will not be cloned to the new host; instead they will be as on the template.

Alternatively, you can use the Host Wizard to configure a host through a guided, step-by-step interface.

Configuration

The **Host** tab contains general host attributes:



All mandatory input fields are marked with a red asterisk.

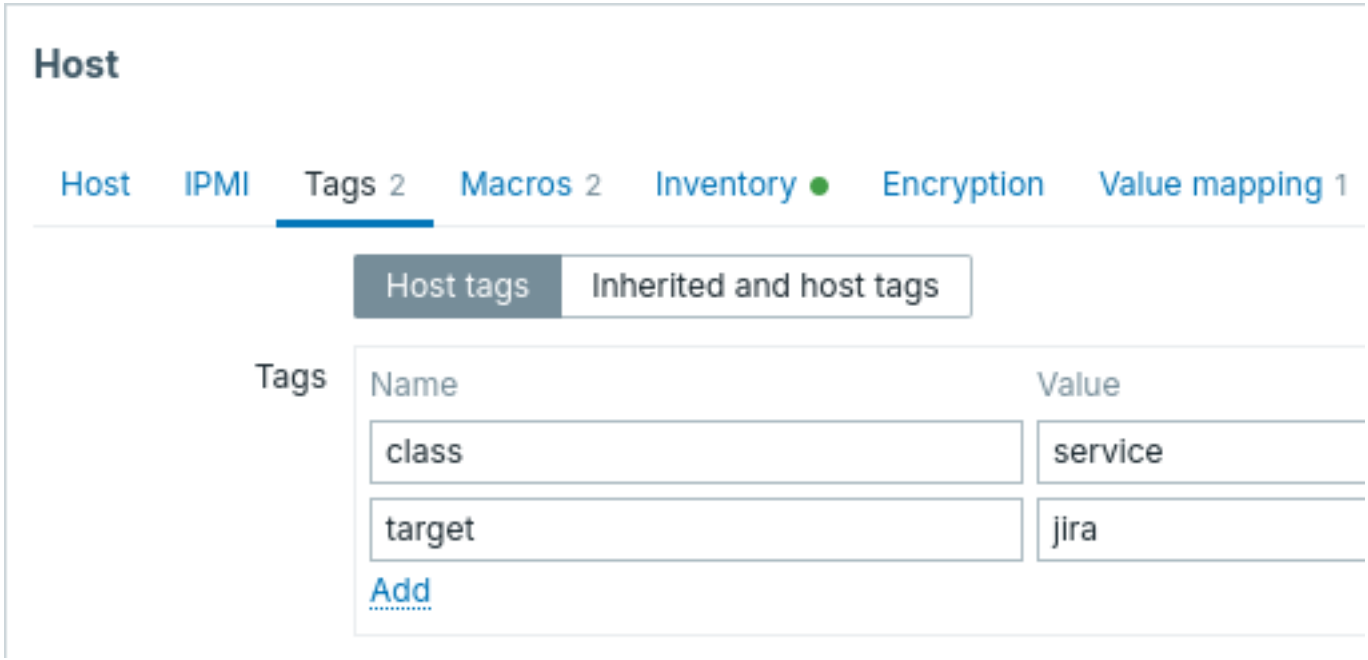
Parameter	Description
Host name	Enter a unique host name. Alphanumerics, spaces, dots, dashes and underscores are allowed. However, leading and trailing spaces are disallowed. *Note:* With Zabbix agent running on the host you are configuring, the agent configuration file parameter *Hostname* must have the same value as the host name entered here. The name in the parameter is needed in the processing of active checks.
Visible name	Enter a unique visible name for the host. If you set this name, it will be the one visible in lists, maps, etc instead of the technical host name. This attribute has UTF-8 support.
Templates	Link templates to the host. All entities (items, triggers, etc.) will be inherited from the template. To link a new template, start typing the template name in the text input field. A list of matching templates will appear; scroll down to select. Alternatively, you may click *Select* next to the field and select templates from the list in a popup window. All selected templates will be linked to the host when the host configuration form is saved or updated. To unlink a template, use one of the two options in the *Linked templates* block: *Unlink* - unlink the template, but preserve its entities (items, triggers, etc.); *Unlink and clear* - unlink the template and remove all its entities (items, triggers, etc.). Listed template names are clickable links leading to the template configuration form.

Parameter	Description
Host groups	Select host groups the host belongs to. A host must belong to at least one host group. A new group can be created and linked to the host by adding a non-existing group name.
Interfaces	Several host interface types are supported for a host: *Agent*, *SNMP*, *JMX* and *IPMI*. No interfaces are defined by default. To add a new interface, click *Add* in the *Interfaces* block, select the interface type and enter *IP/DNS*, *Connect to* and *Port* info. *Note:* Interfaces that are used in any items cannot be removed and link *Remove* is grayed out for them. The "IP" or "DNS" from an SNMP interface is also used for SNMP traps. During matching, only the selected "IP" or "DNS" in the host interface is used. See Configuring SNMP monitoring for additional details on configuring an SNMP interface (v1, v2 and v3).
IP address	Host IP address (optional).
DNS name	Host DNS name (optional).
Connect to	Clicking the respective button will tell Zabbix server what to use to retrieve data from agents: **IP** - Connect to the host IP address (recommended) **DNS** - Connect to the host DNS name
Port	TCP/UDP port number. Default values are: 10050 for Zabbix agent, 161 for SNMP agent, 12345 for JMX and 623 for IPMI.
Default	Check the radio button to set the default interface.
Description	Enter the host description.
Monitored by	Select if the host is monitored by: **Server** - host is monitored by Zabbix server; **Proxy** - host is monitored by standalone proxy; **Proxy group** - host is monitored by proxy group.
Proxy	The assigned proxy name is displayed (only if Zabbix server has assigned one from the selected proxy group). This field is displayed only if the host is monitored by a proxy group.
Enabled	When the checkbox is checked, the host is enabled - ready for monitoring. When the checkbox is unchecked, the host is disabled - not monitored: For passive data requests initiated by Zabbix server/proxy (for example, Zabbix agent, SNMP agent, simple checks), monitoring is disabled after configuration synchronization. Triggers and actions linked to the host are also disabled only after the configuration is reloaded. For Zabbix agent active checks, monitoring stops within the time frame (approx. 5 seconds) that Zabbix agent receives information about the host having been disabled. During this brief interval, the host will continue to locally collect data for the active checks and try sending it to the server/proxy; however, since the host is marked as *Disabled*, the server/proxy will reject the data. When you disable the host, its items are immediately removed from the history cache (except for their last values, which are kept for logs).

The **IPMI** tab contains IPMI management attributes.

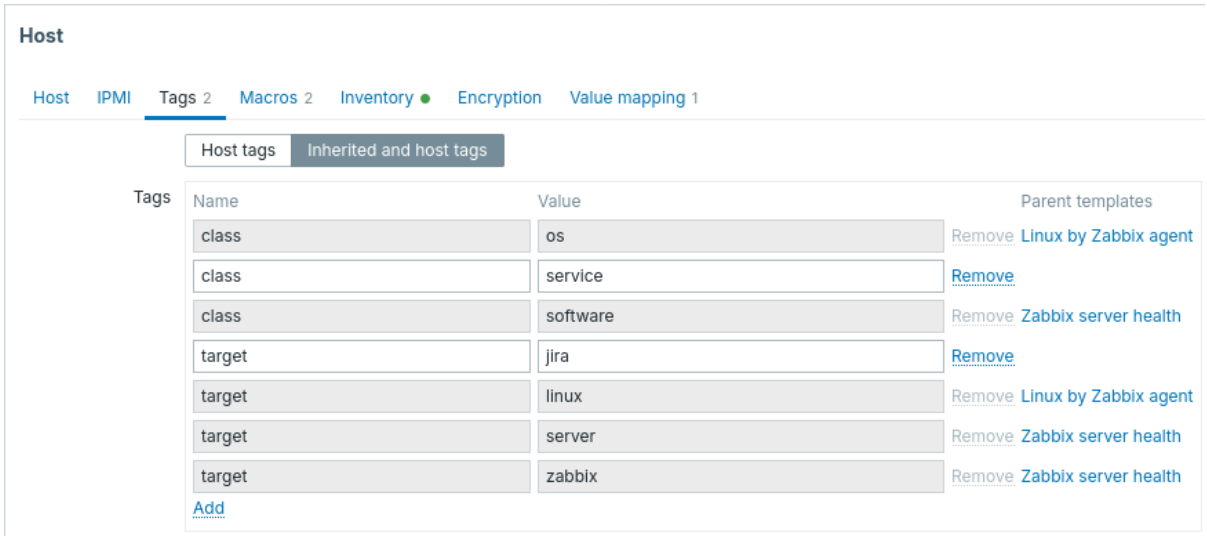
Parameter	Description
Authentication algorithm	Select the authentication algorithm.
Privilege level	Select the privilege level.
Username	User name for authentication. User macros may be used.
Password	Password for authentication. User macros may be used.

The **Tags** tab allows you to define host-level tags. All problems of this host will be tagged with the values entered here.



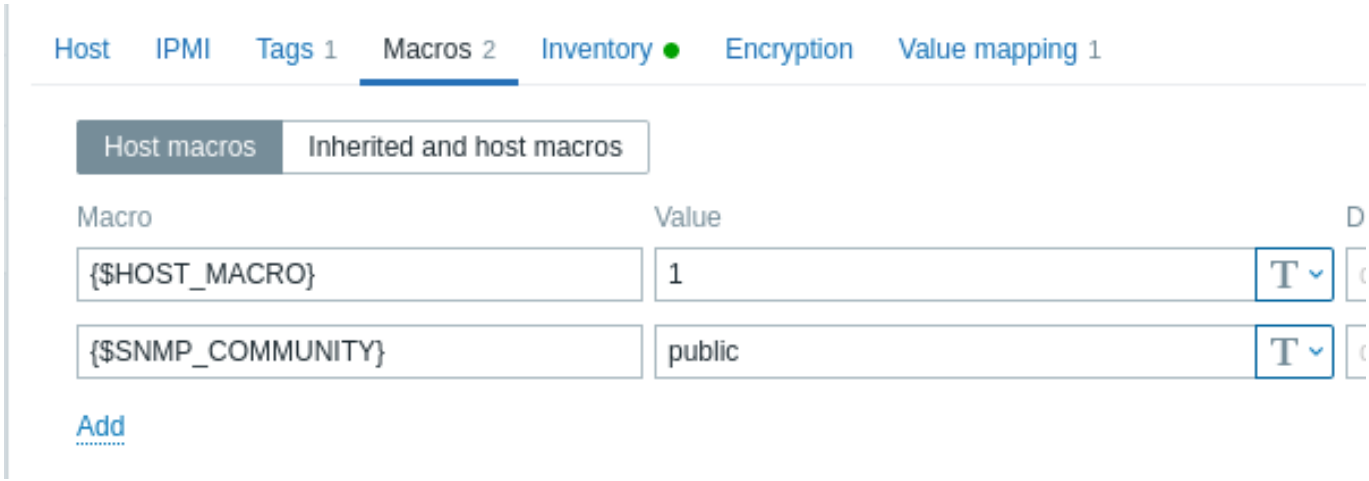
User macros, `{INVENTORY.*}` macros, `{HOST.HOST}`, `{HOST.NAME}`, `{HOST.CONN}`, `{HOST.DNS}`, `{HOST.IP}`, `{HOST.PORT}` and `{HOST.ID}` macros are supported in tags.

You may also view here template-level tags if you select the *Inherited and host tags* option. That is where all defined user tags for the host are displayed as well as their origin.

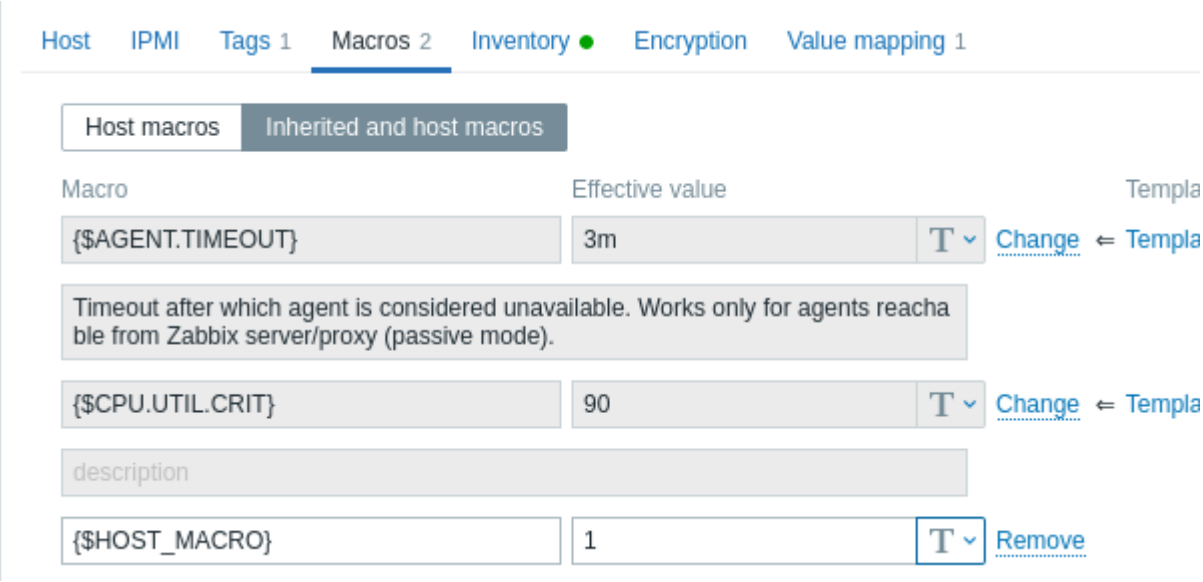


For convenience, links to respective templates are provided.

The **Macros** tab allows you to define host-level user macros as a name-value pairs. Note that macro values can be kept as plain text, secret text or Vault secret. Adding a description is also supported.

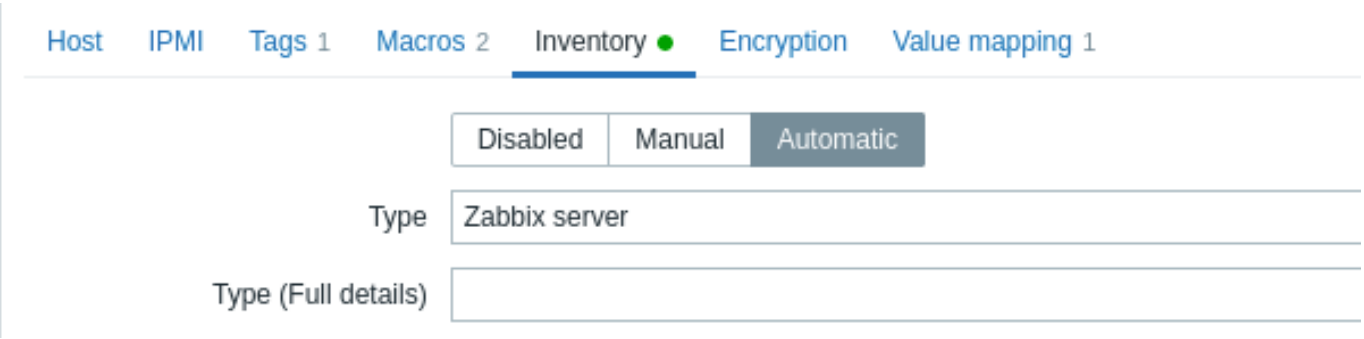


You may also view here template-level and global user macros if you select the *Inherited and host macros* option. That is where all defined user macros for the host are displayed with the value they resolve to as well as their origin.



For convenience, links to respective templates and global macro configuration are provided. It is also possible to edit a template/global macro on the host level, effectively creating a copy of the macro on the host.

The **Inventory** tab allows you to manually enter inventory information for the host. You can also select to enable *Automatic* inventory population, or disable inventory population for this host.



If inventory is enabled (manual or automatic), a green dot is displayed with the tab name.

Encryption

The **Encryption** tab allows you to require encrypted connections with the host.

Parameter	Description
Connections to host	How Zabbix server or proxy connects to Zabbix agent on a host: no encryption (default), using PSK (pre-shared key) or certificate.
Connections from host	Select what type of connections are allowed from the host (i.e. from Zabbix agent and Zabbix sender). Several connection types can be selected at the same time (useful for testing and switching to other connection type). Default is "No encryption".
Issuer	Allowed issuer of certificate. Certificate is first validated with CA (certificate authority). If it is valid, signed by the CA, then the *Issuer* field can be used to further restrict allowed CA. This field is intended to be used if your Zabbix installation uses certificates from multiple CAs. If this field is empty then any CA is accepted.
Subject	Allowed subject of certificate. Certificate is first validated with CA. If it is valid, signed by the CA, then the *Subject* field can be used to allow only one value of *Subject* string. If this field is empty then any valid certificate signed by the configured CA is accepted.
PSK identity	Pre-shared key identity string. Do not put sensitive information in the PSK identity, it is transmitted unencrypted over the network to inform a receiver which PSK to use.
PSK	Pre-shared key (hex-string). Maximum length: 512 hex-digits (256-byte PSK) if Zabbix uses GnuTLS or OpenSSL library, 64 hex-digits (32-byte PSK) if Zabbix uses mbed TLS (PolarSSL) library. Example: 1f87b595725ac58dd977beef14b97461a7c1045b9a1c963065002c5473194952

Value mapping

The **Value mapping** tab allows you to configure human-friendly representation of item data in value mappings.

Configuring a host group

Overview

Host groups are used for the logical grouping of hosts and assigning user permissions to them.

Each host must have at least one host group assigned. A host may belong to multiple host groups, and each host group may contain multiple hosts.

Note that in Zabbix, all permissions are based on groups: user groups, host groups, and template groups. So, even if a single user needs access to a single host, it is granted by adding the user to a user group and assigning that user group permissions to the host group containing that host.

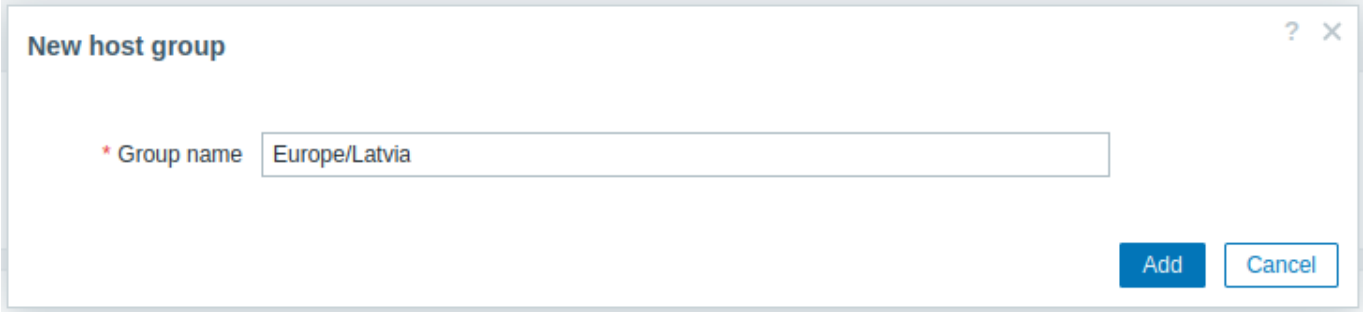
Configuration

> **Attention:**
> Only Super admin users can create host groups.

There are two options of creating a host group in Zabbix frontend.

Option one:

1. Go to *Data collection* > *Host groups*.
2. Click *Create host group* in the upper-right corner of the screen.
3. Enter the group name in the form.



Option two: when configuring a host, enter a non-existing group name in the *Host groups* input field.

Once the host group is created, you can click the group name in the list under *Data collection* > *Host groups* to edit the group name, clone the group, or delete the group.

Deleting a host group only deletes the logical group, not the hosts in the group. It is not possible to delete a host group that is the only group for any existing host.

Creating host subgroups

A host subgroup (or nested host group) is a child of the parent host group that contains it.

A subgroup is created by using the forward slash '/' in the group name input field to denote its relation to the parent group(s). For example:

- inputting `Europe/Latvia` creates the `Europe/Latvia` subgroup of the parent group `Europe`.
- inputting `Europe/Latvia/Riga/Zabbix servers` creates the respective subgroup within the nested parent groups `Europe`, `Europe/Latvia`, `Europe/Latvia/Riga`.

When creating a subgroup, using leading or trailing slashes, or several slashes in a row is not allowed. Escaping of '/' is not supported.

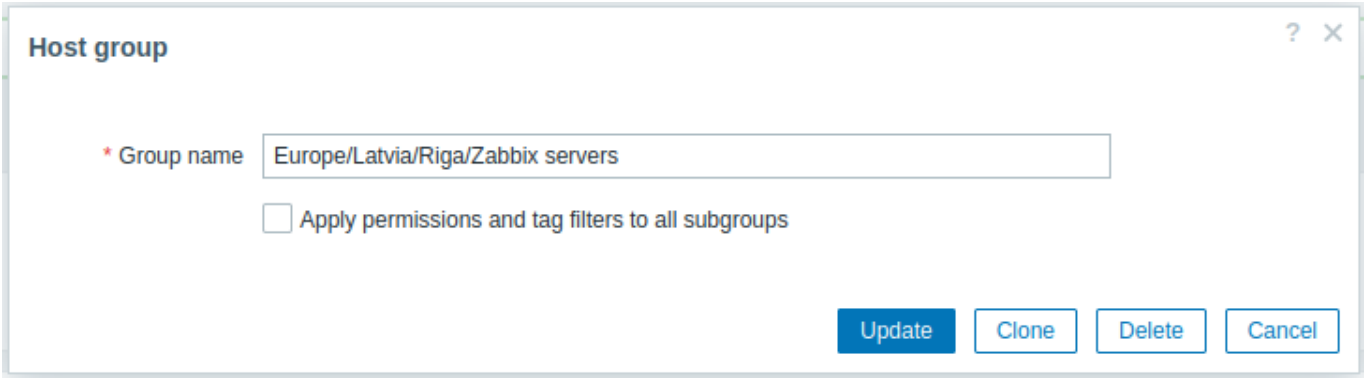
It is not required to create any parent host group(s) before creating a subgroup. You can choose whether to start by creating a subgroup (for example, `Europe/Latvia`) or any parent host group(s) (in our example, `Europe`). If you start by creating a subgroup, parent host group(s) will **not** be created automatically.

Permissions to host groups

- When creating a subgroup to an existing parent host group (for example, creating `Europe/Latvia` when `Europe` already exists), user group permissions to the subgroup are inherited from the parent.

- When creating a parent host group to an existing subgroup (for example, creating `Europe` when `Europe/Latvia` already exists), no permissions to the parent are set.

When editing any host group, you can also set an additional option, *Apply permissions and tag filters to all subgroups*.



Marking this checkbox and clicking *Update* will apply the same level of permissions and tag filters to all current and future subgroups of the host group being edited.

So, if any user groups have been given varying permissions to the subgroups of the host group being edited, marking the checkbox will grant all current and future subgroups the same user permissions and tag-based permissions as the group being edited.

Note that this option is not saved in the database and will override existing permissions. Any changes made through this option can be reverted only manually.

Inventory

Overview

You can keep the inventory of networked devices in Zabbix.

There is a special *Inventory* menu in the Zabbix frontend. However, you will not see any data there initially and it is not where you enter data. Building inventory data is done manually when configuring a host or automatically by using some automatic population options.

Building inventory

Manual mode

When configuring a host, in the *Inventory* tab you can enter such details as the type of device, serial number, location, responsible person, URLs, etc. - data that will populate inventory information.

If a URL is included in host inventory information and it starts with 'http' or 'https', it will result in a clickable link in the *Inventory* section.

Automatic mode

Host inventory can also be populated automatically. For that to work, when configuring a host the inventory mode in the *Inventory* tab must be set to *Automatic*.

Then you can configure host items to populate any host inventory field with their value, indicating the destination field with the respective attribute (called *Item will populate host inventory field*) in item configuration.

Items that are especially useful for automated inventory data collection:

- system.hw.chassis[full|type|vendor|model|serial] - default is [full], root permissions needed
- system.hw.cpu[all|cpunum,full|maxfreq|vendor|model|curfreq] - default is [all,full]
- system.hw.devices[pci|usb] - default is [pci]
- system.hw.macaddr[interface,short|full] - default is [all,full], interface is regexp
- system.sw.arch
- system.sw.os[name|short|full] - default is [name]
- system.sw.packages[regexp,manager,short|full] - default is [all,all,full]

Inventory mode selection

Inventory mode can be selected in the host configuration form.

Inventory mode by default for new hosts is selected based on the *Default host inventory mode* setting in *Administration* > *General* > *Other*.

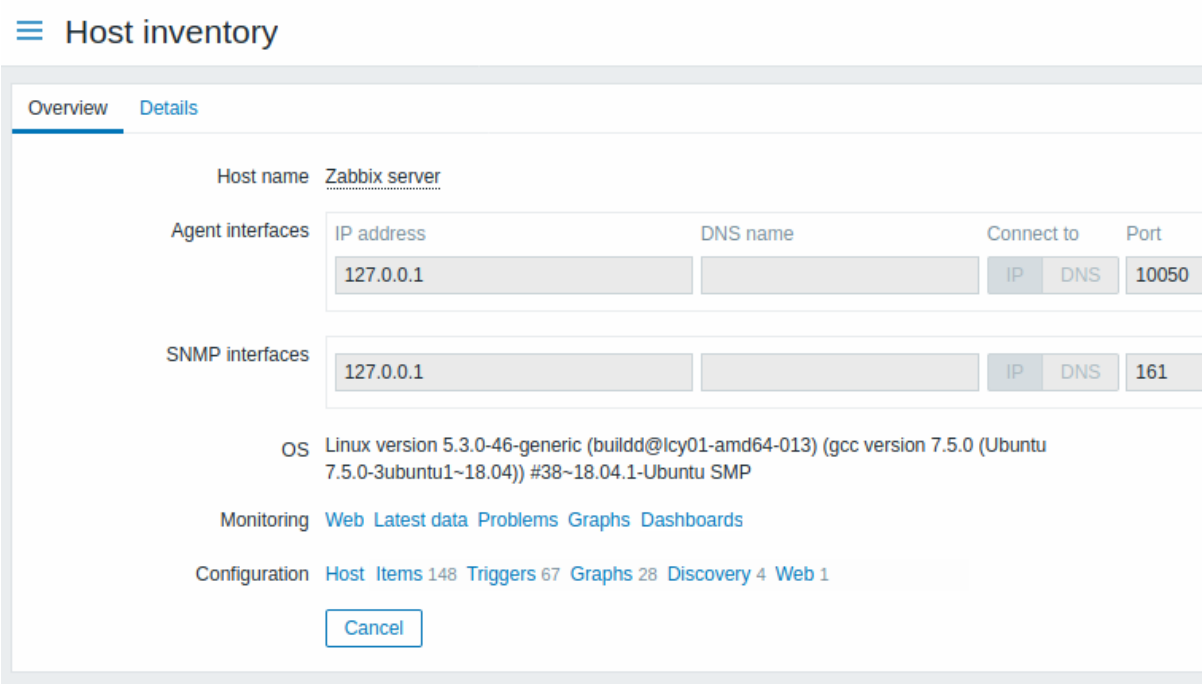
For hosts added by network discovery or autoregistration actions, it is possible to define a *Set host inventory mode* operation selecting manual or automatic mode. This operation overrides the *Default host inventory mode* setting.

Inventory overview

The details of all existing inventory data are available in the *Inventory* menu.

In *Inventory* > *Overview* you can get a host count by various fields of the inventory.

In *Inventory* > *Hosts* you can see all hosts that have inventory information. Clicking the host name will reveal the inventory details in a form.



The **Overview** tab shows:

Parameter	Description
Host name	Name of the host. Clicking the name opens a menu with the scripts defined for the host. Host name is displayed with an orange icon, if the host is in maintenance.
Visible name	Visible name of the host (if defined).
Host (Agent, SNMP, JMX, IPMI) interfaces	This block provides details of the interfaces configured for the host.
OS	Operating system inventory field of the host (if defined).
Hardware	Host hardware inventory field (if defined).
Software	Host software inventory field (if defined).
Description	Host description.
Monitoring	Links to monitoring sections with data for this host: *Web*, *Latest data*, *Problems*, *Graphs*, *Dashboards*.
Configuration	Links to configuration sections for this host: *Host*, *Items*, *Triggers*, *Graphs*, *Discovery*, *Web*. The amount of configured entities is listed after each link.

The **Details** tab shows all inventory fields that are populated (are not empty).

Inventory macros

There are host inventory macros `{INVENTORY.*}` available for use in notifications, for example:

"Server in {INVENTORY.LOCATION1} has a problem, responsible person is {INVENTORY.CONTACT1}, phone number {INVENTORY.POC.PRIMARY.PHONE.A1}."

For more details, see the Supported macros page.

Mass update

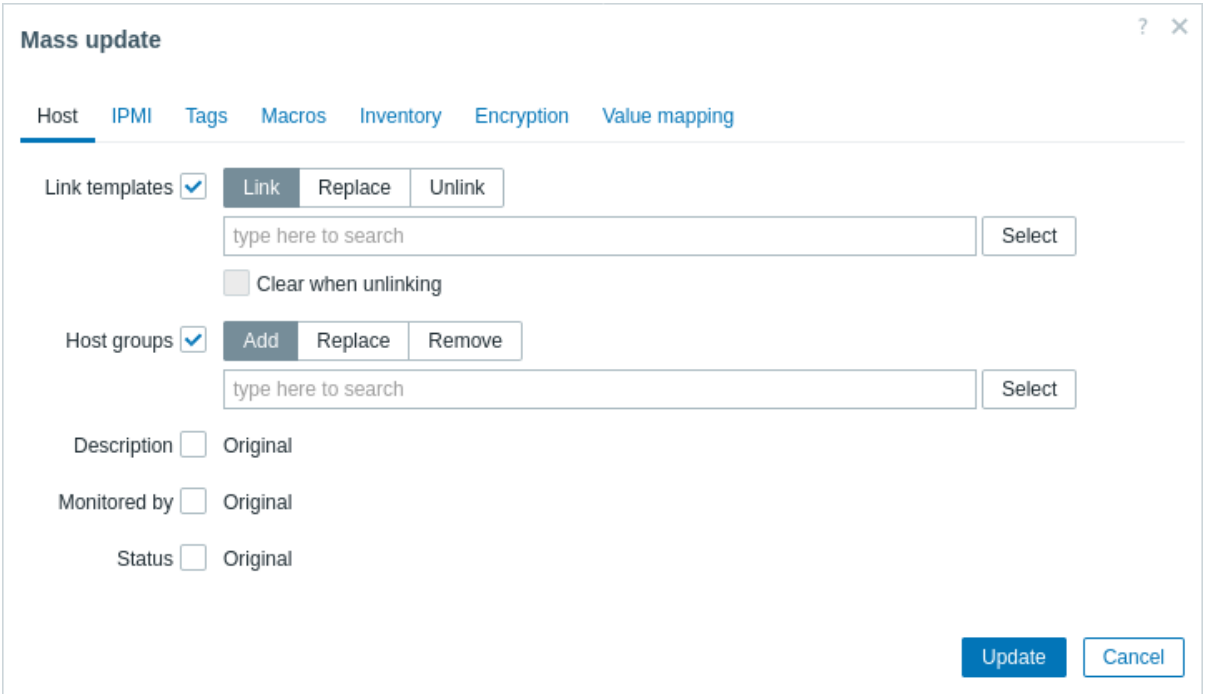
Overview

Sometimes you may want to change some attribute for a number of hosts at once. Instead of opening each individual host for editing, you may use the mass update function for that.

Using mass update

To mass-update some hosts, do the following:

- Mark the checkboxes before the hosts you want to update in the host list
- Click *Mass update* below the list
- Navigate to the tab with required attributes (*Host*, *IPMI*, *Tags*, *Macros*, *Inventory*, *Encryption* or *Value mapping*)
- Mark the checkboxes of any attribute to update and enter a new value for them



The following options are available when selecting the respective button for **template** linkage update:

- *Link* - specify which additional templates to link
- *Replace* - specify which templates to link while unlinking any template that was linked to the hosts before
- *Unlink* - specify which templates to unlink

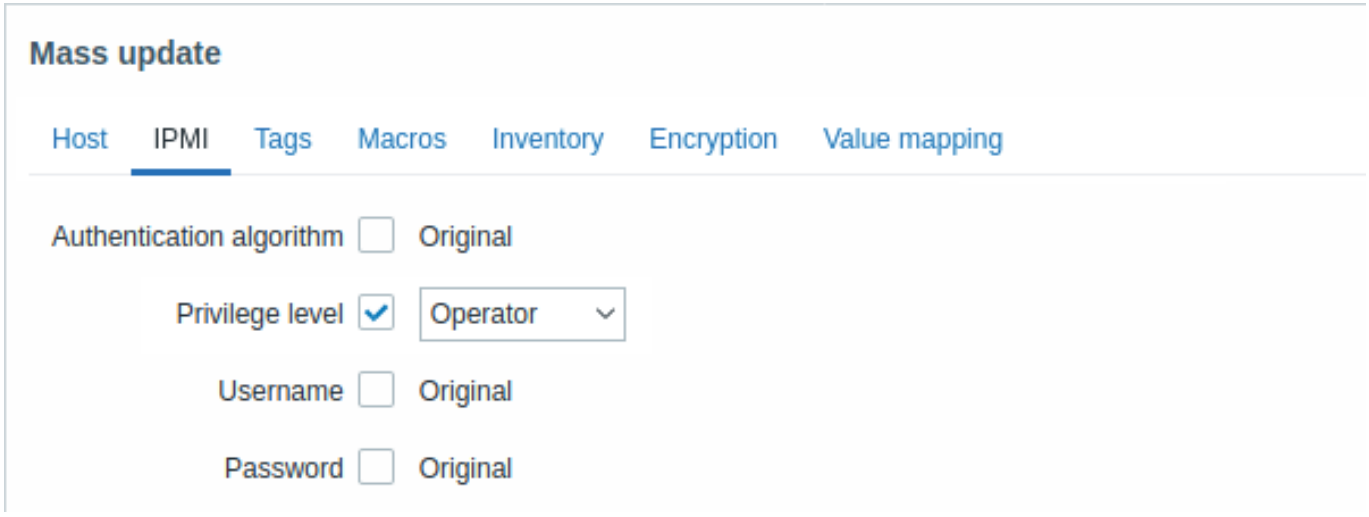
To specify the templates to link/unlink start typing the template name in the auto-complete field until a dropdown appears offering the matching templates. Just scroll down to select the required template.

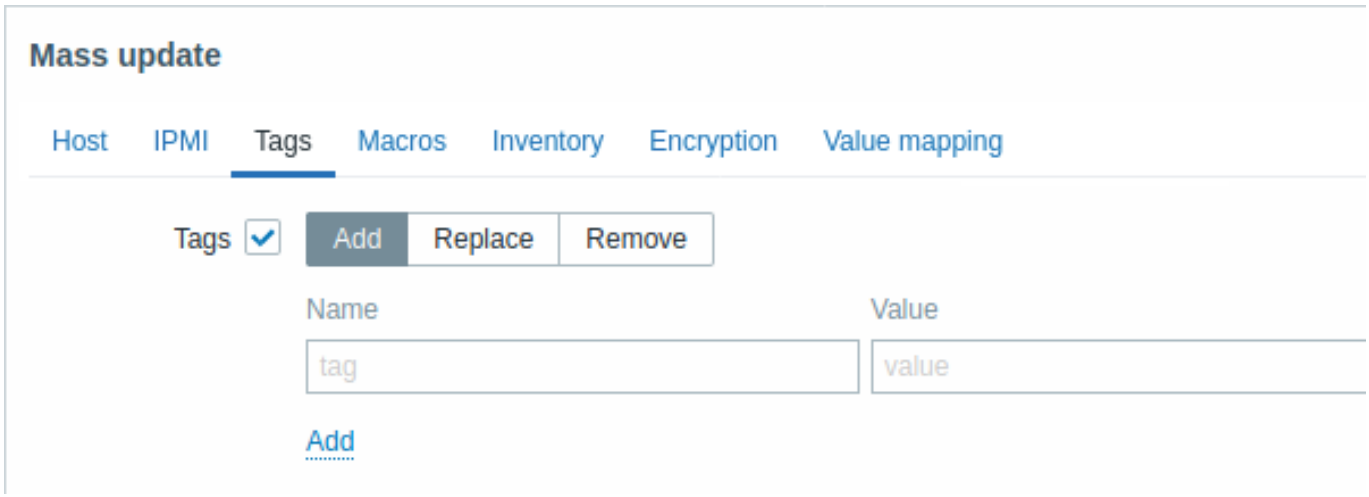
The *Clear when unlinking* option allows you not only to unlink previously linked templates but also to remove all entities inherited from them (items, triggers, graphs, etc.).

The following options are available when selecting the respective button for **host group** update:

- *Add* - specify additional host groups from the existing ones or enter completely new host groups for the hosts;
- *Replace* - remove the host from any existing host groups and replace them with the one(s) specified in this field (existing or new host groups);
- *Remove* - remove specific host groups from hosts.

These fields are auto-complete - starting to type in them offers a dropdown of matching host groups. If the host group is new, it also appears in the dropdown and it is indicated by *(new)* after the string. Just scroll down to select.

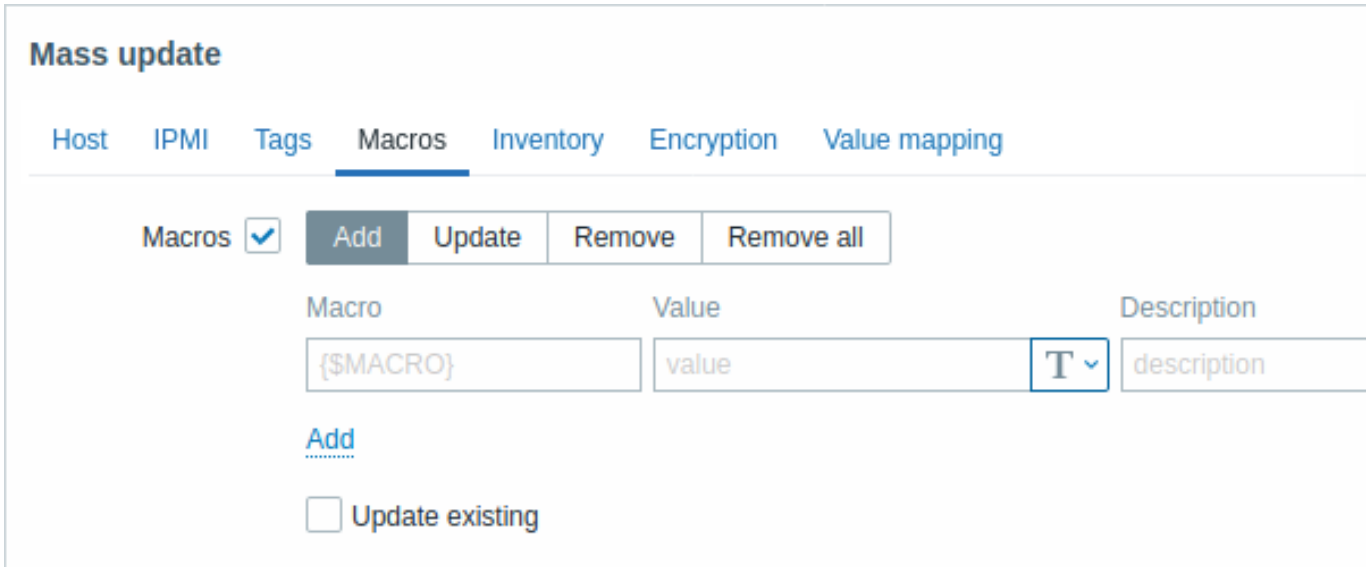




The following options are available when selecting the respective button for tag update:

- *Add* - add new tags for the hosts;
- *Replace* - remove any existing tags from the host and replace them with the one(s) specified below;
- *Remove* - remove specified tags from hosts.

User macros, `{INVENTORY.*}` macros, `{HOST.HOST}`, `{HOST.NAME}`, `{HOST.CONN}`, `{HOST.DNS}`, `{HOST.IP}`, `{HOST.PORT}` and `{HOST.ID}` macros are supported in tags. Note that tags with the same name but different values are not considered 'duplicates' and can be added to the same host.

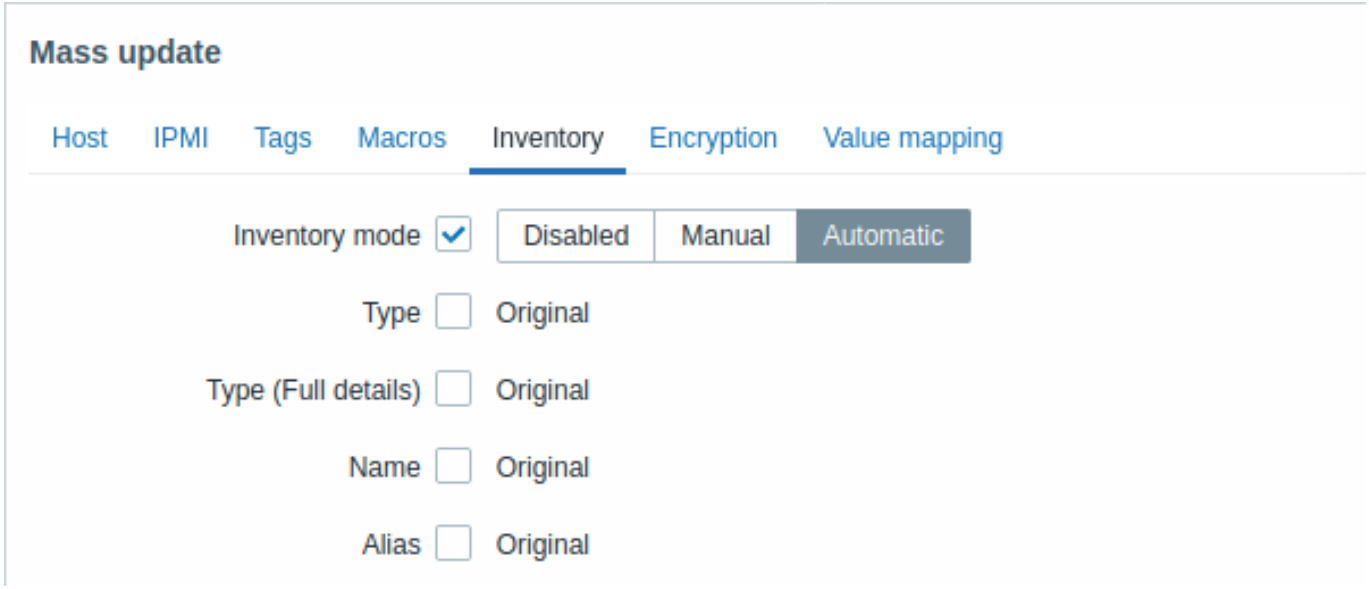


The following options are available when selecting the respective button for macros update:

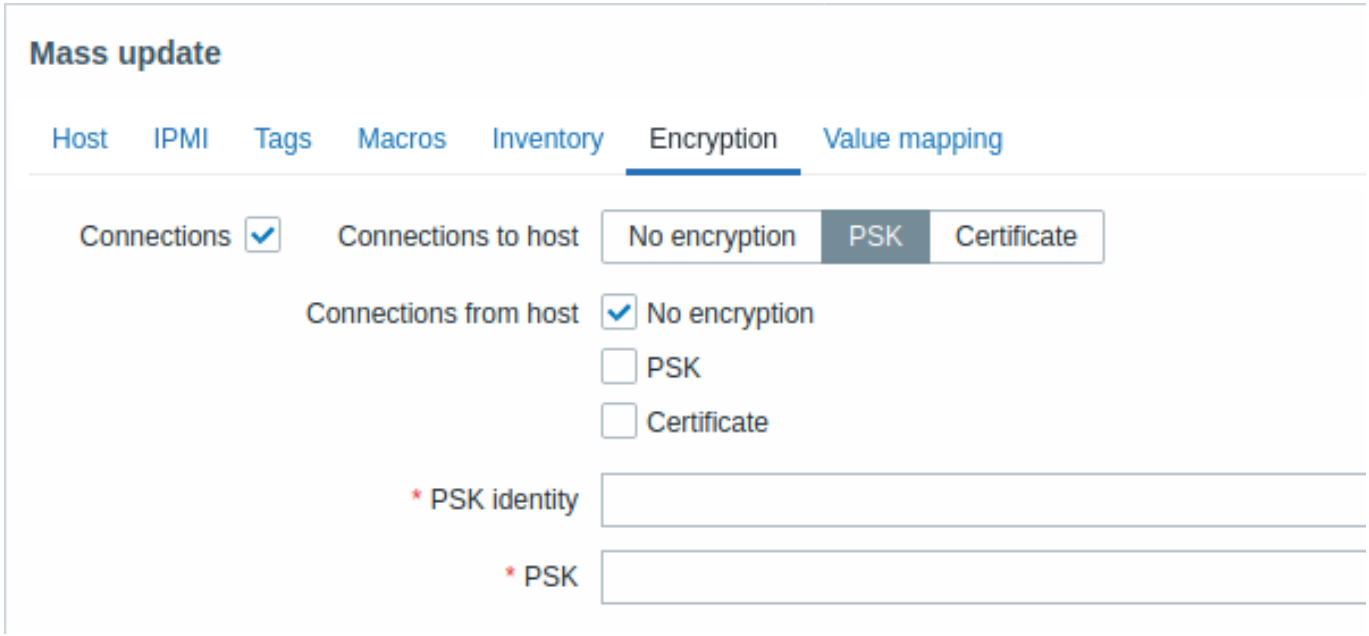
- *Add* - specify additional user macros for the hosts. If the *Update existing* checkbox is checked, value, type and description for the specified macro name will be updated. If unchecked, if a macro with that name already exist on the host(s), it will

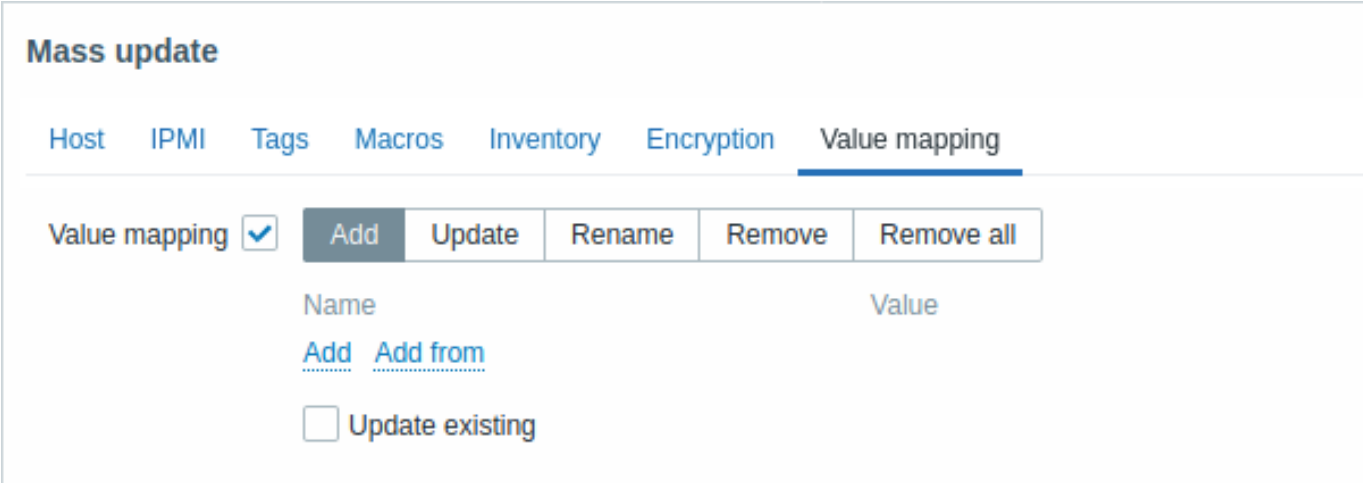
not be updated.

- *Update* - replace values, types and descriptions of macros specified in this list. If the *Add missing* checkbox is checked, macro that didn't previously exist on a host will be added as new macro. If unchecked, only macros that already exist on a host will be updated.
- *Remove* - remove specified macros from hosts. If the *Except selected* checkbox is checked, all macros except specified in the list will be removed. If unchecked, only macros specified in the list will be removed.
- *Remove all* - remove all user macros from hosts. If the *I confirm to remove all macros* checkbox is not checked, a new popup window will open asking to confirm removal of all macros.



To be able to mass update inventory fields, the *Inventory mode* should be set to 'Manual' or 'Automatic'.





Buttons with the following options are available for value map update:

- *Add* - add value maps to the hosts. If you mark *Update existing*, all properties of the value map with this name will be updated. Otherwise, if a value map with that name already exists, it will not be updated.
- *Update* - update existing value maps. If you mark *Add missing*, a value map that didn't previously exist on a host will be added as a new value map. Otherwise, only the value maps that already exist on a host will be updated.
- *Rename* - give new name to an existing value map
- *Remove* - remove the specified value maps from the hosts. If you mark *Except selected*, all value maps will be removed **except** the ones that are specified.
- *Remove all* - remove all value maps from the hosts. If the *I confirm to remove all value maps* checkbox is not marked, a new popup window will open asking to confirm the removal.

When done with all required changes, click *Update*. The attributes will be updated accordingly for all the selected hosts.

Items

Overview

An item is an individual metric.

Items are used for collecting data. Once you have configured a host, you must add items to get actual data. One way of quickly adding many items is to attach one of the predefined templates to a host. However, for optimized system performance, you may need to fine-tune the templates to have as many items and as frequent monitoring as necessary.

To specify what sort of data to collect from a host, use the item key. For example, an item with the key name **system.cpu.load** will collect processor load data, while an item with the key name **net.if.in** will collect incoming traffic information.

Additional parameters can be specified in square brackets after the key name. For example, system.cpu.load**[avg5]** will return the processor load average for the last 5 minutes, while net.if.in**[eth0]** will show incoming traffic in the interface "eth0".

Note:
See individual sections of item types for all supported item types and item keys.

Proceed to creating and configuring an item.

Creating an item

Overview

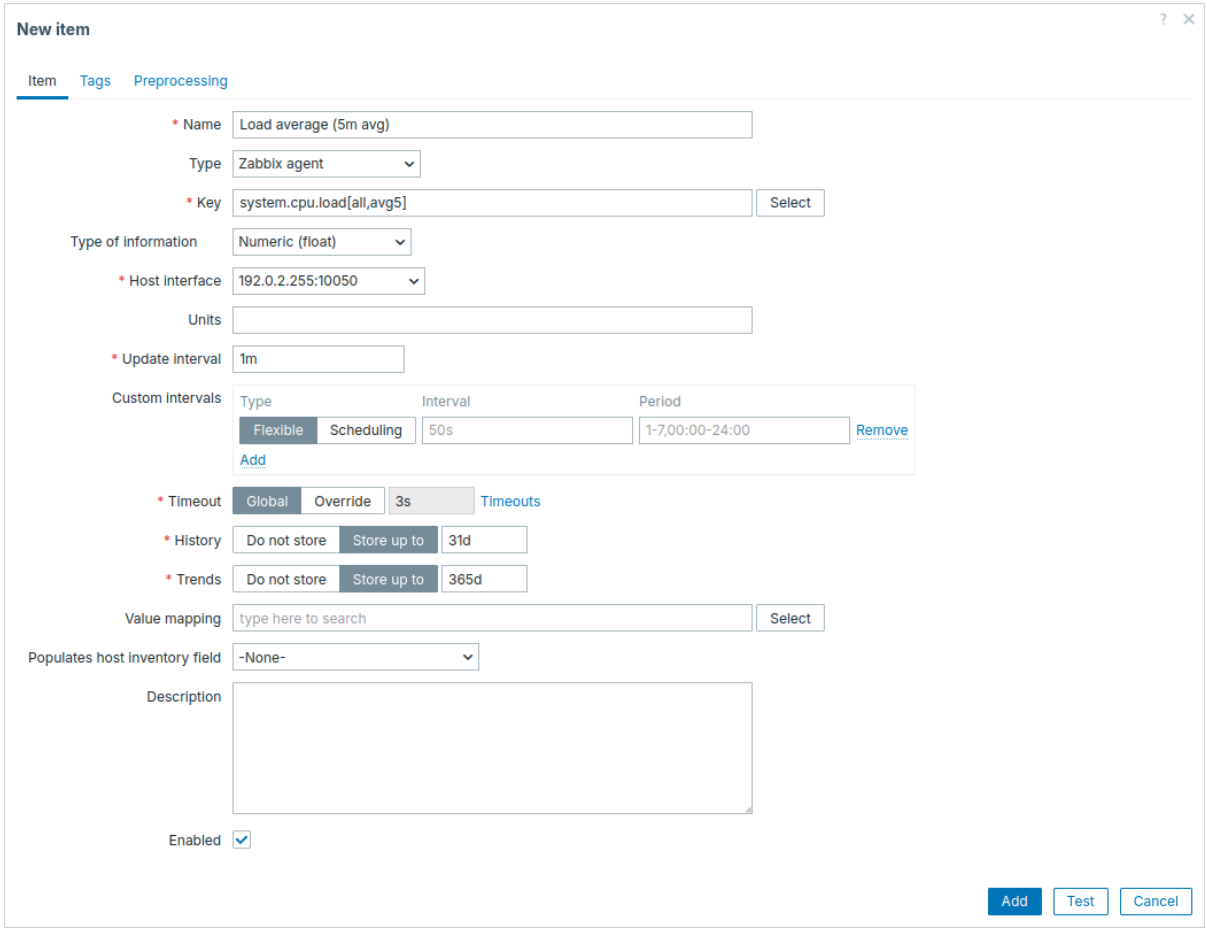
To create an item in Zabbix frontend, do the following:

1. Go to *Data collection* > *Hosts*.
2. Click *Items* in the row of the host.
3. Click *Create item* in the upper-right corner of the screen.
4. Enter parameters of the item in the form.

Alternatively, you can create an item by opening an existing one, clicking the *Clone* button, and saving it under a different name.

Configuration

The **Item** tab contains general item attributes.



All mandatory input fields are marked with a red asterisk.

Parameter	Description
Name	Item name. User macros are supported.
Type	Item type. See individual item type sections.
Key	Item key (up to 2048 characters). The supported item keys can be found in individual item type sections. The key must be unique within a single host. If key type is 'Zabbix agent', 'Zabbix agent (active)' or 'Simple check', the key value must be supported by Zabbix agent or Zabbix server. See also: the correct key format.
Type of information	The data type used to validate the item value and store it in the database, both after any conversions: **Numeric (unsigned)** - 64-bit unsigned integer; **Numeric (float)** - 64-bit floating-point number; **Character** - short text; **Log** - long text with or without log properties (timestamp, source, severity, log event ID); **Text** - long text; **Binary** - binary number (supported only for dependent items); **JSON** - structured JSON data, stored natively in the database (not supported for calculated items). Items returning values with the binary or JSON data type are not supported in calculated item formulas or trigger expressions. See item data limits for storage limits per data type and database backend. For items that return data in a single specific format, a matching data type is selected automatically—for example, `system.cpu.load` and `system.cpu.util` both return a 64-bit floating-point number, so *Numeric (float)* is selected automatically.
Host interface	Select the host interface. This field is available when editing an item on the host level.

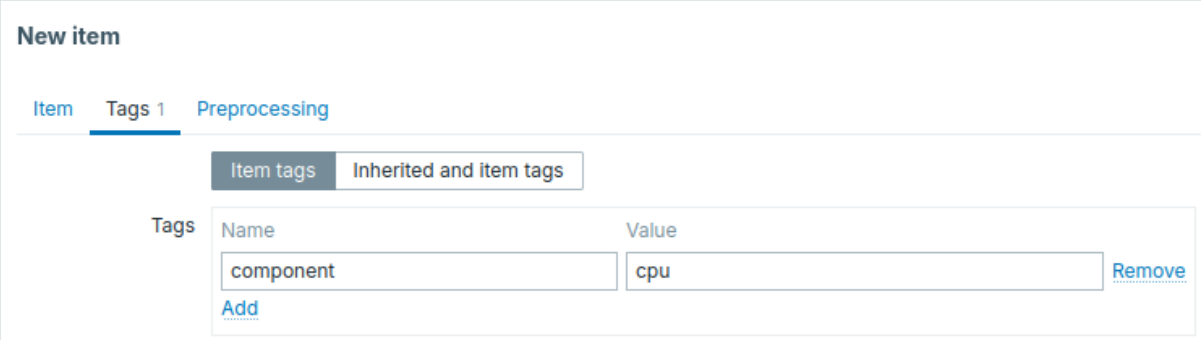
Parameter	Description
Units	If a unit symbol is set, Zabbix applies postprocessing to the received item value and displays it with the specified unit postfix. Supported unit symbols with special formatting (and examples of received item value → displayed value): **B** - bytes (1024 → 1 KB) **Bps** - bytes per second (1024 → 1 KBps) **s** - seconds, displayed using up to three largest non-zero time units (881764 → 10d 4h 56m) **uptime** - elapsed time in hh:mm:ss or N days, hh:mm:ss (881764 → 10 days, 04:56:04) **unixtime** - Unix timestamp, formatted as yyyy.mm.dd hh:mm:ss (881764 → 1970-01-11 04:56:04 AM); for correct formatting, the received item value must be *Numeric (unsigned)*. For other units (such as Hz, W, etc.), if the received value exceeds 1000, it is divided by 1000 and displayed with a corresponding prefix (5000 → 5 KHz, 881764 → 881.76 KW). If the unit symbol is prefixed with !, unit conversion and prefixing are disabled (1024 !B → 1024 B, 61 !s → 61 s). For more examples and details on unit symbols and unit conversion, see Item value suffixes.
Update interval	Retrieve a new value for this item every N seconds. Maximum allowed update interval is 86400 seconds (1 day). Time suffixes are supported, e.g., 30s, 1m, 2h, 1d. User macros are supported. A single macro has to fill the whole field. Multiple macros in a field or macros mixed with text are not supported. *Note*: The update interval can only be set to '0' if custom intervals exist with a non-zero value. If set to '0', and a custom interval (flexible or scheduled) exists with a non-zero value, the item will be polled during the custom interval duration. *Note* that the first item poll after the item became active or after update interval change might occur earlier than the configured value. New items will be checked within 60 seconds of their creation, unless they have Scheduling or Flexible update interval and the *Update interval* is set to 0. An existing passive item can be polled for value immediately by pushing the *Execute now* button.
Custom intervals	You can create custom rules for checking the item: **Flexible** - create an exception to the *Update interval* (interval with different frequency). **Scheduling** - create a custom polling schedule. For detailed information see Custom intervals. Time suffixes are supported in the *Interval* field, e.g., 30s, 1m, 2h, 1d. User macros are supported. A single macro has to fill the whole field. Multiple macros in a field or macros mixed with text are not supported.
Timeout	Set the item check timeout (available for supported item types). Select the timeout option: **Global** - proxy/global timeout is used (displayed in the grayed out *Timeout* field). **Override** - custom timeout is used (set in the *Timeout* field; allowed range: 1 - 600s). Time suffixes, e.g., 30s, 1m, and user macros are supported. Clicking the *Timeouts* link allows you to configure proxy timeouts or global timeouts (if a proxy is not used). Note that the *Timeouts* link is visible only to users of *Super admin* type with permissions to *Administration* > *General* or *Administration* > *Proxies* frontend sections.
History	Select either: **Do not store** - item history is not stored. Useful for master items if only dependent items need to keep history. This setting cannot be overridden by global housekeeper settings. **Store up to** - specify the duration of keeping detailed history in the database (1 hour to 25 years). Older data will be removed by the housekeeper. Stored in seconds. Time suffixes are supported, e.g., 2h, 1d. User macros are supported. The *Store up to* value can be overridden globally in *Administration* > *Housekeeping*. If a global overriding setting exists, an orange  info icon is displayed. If you position your mouse on it, a warning message is displayed, e.g., *Overridden by global housekeeper settings (1d)*. It is recommended to keep the recorded values for the smallest possible time to reduce the size of value history in the database. Instead of storing a long history of values, you can store longer data of trends. See also History and trends.

Parameter	Description
Trends	Select either: **Do not store** - trends are not stored. This setting cannot be overridden by global housekeeper settings. **Store up to** - specify the duration of keeping aggregated (hourly min, max, avg, count) history in the database (1 day to 25 years). Older data will be removed by the housekeeper. Stored in seconds. Time suffixes are supported, e.g., 24h, 1d. User macros are supported. The *Store up to* value can be overridden globally in *Administration* > *Housekeeping*. If a global overriding setting exists, an orange  info icon is displayed. If you position your mouse on it, a warning message is displayed, e.g., *Overridden by global housekeeper settings (7d)*. *Note:* Keeping trends is not available for non-numeric data - character, log and text. See also History and trends.
Value mapping	Apply value mapping to this item. Value mapping does not change received values, it is for displaying data only. It works with *Numeric(unsigned)*, *Numeric(float)* and *Character* items. For example, "Windows service states".
Log time format	Available for items of type **Log** only. Supported placeholders: **y**: *Year (1970-2038)*. **M**: *Month (01-12)*. **d**: *Day (01-31)*. **h**: *Hour (00-23)*. **m**: *Minute (00-59)*. **s**: *Second (00-59)*. If left blank, the timestamp will be set to 0 in Unix time, representing January 1, 1970. For example, consider the following line from the Zabbix agent log file: " 23480:20100328:154718.045 Zabbix agent started. Zabbix 1.8.2 (revision 11211)." It begins with six character positions for PID, followed by date, time, and the rest of the message. The log time format for this line would be "ppppp:yyyyMMdd:hhmmss". Note that "p" and ":" characters are placeholders and can be any characters except "yMdhms".
Populates host inventory field	You can select a host inventory field that the value of item will populate. This will work if automatic inventory population is enabled for the host and with *Numeric(unsigned)*, *Numeric(float)*, *Character*, and *Text* items.
Description	Enter an item description. User macros are supported.
Enabled	Mark the checkbox to enable the item so it will be processed. When you disable the item, it is immediately removed from the history cache (except for its last value, which is kept for logs).
Latest data	Click the link to view the latest data for the item. This link is only available when editing an already existing item.

Note:
Item type specific fields are described on corresponding pages.

Note:
When editing an existing template level item on a host level, a number of fields are read-only. You can use the link in the form header and go to the template level and edit them there, keeping in mind that the changes on a template level will change the item for all hosts that the template is linked to.

The **Tags** tab allows you to define item-level tags.



Item value preprocessing

The **Preprocessing** tab allows you to define transformation rules for the received values.

Item testing

Attention:
To perform item testing, ensure that the system time on the server and the proxy is synchronized. In the case when the server time is behind, item testing may return an error message "The task has been expired." Having set different time zones on the server and the proxy, however, won't affect the testing result.

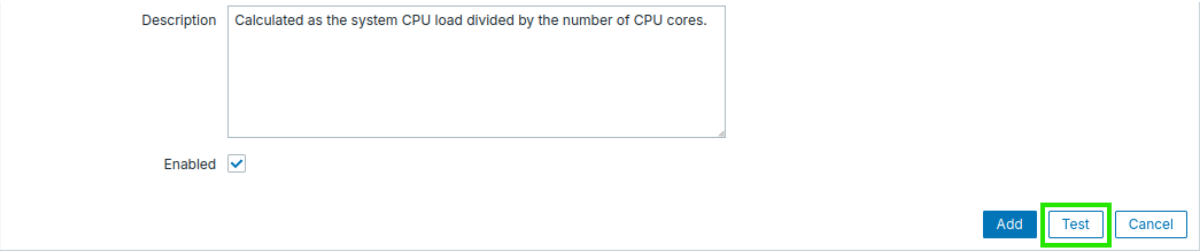
It is possible to test an item and, if configured correctly, get a real value in return. Testing can occur even before an item is saved.

Testing is available for host and template items, item prototypes and low-level discovery rules. Testing is not available for active items.

Item testing is available for the following passive item types:

- Zabbix agent
- SNMP agent (v1, v2, v3)
- IPMI agent
- SSH checks
- Telnet checks
- JMX agent
- Simple checks (except `icmpping*`, `vmware.*` items)
- Zabbix internal
- Calculated items
- External checks
- Database monitor
- HTTP agent
- Script
- Browser

To test an item, click the *Test* button at the bottom of the item configuration form. Note that the *Test* button will be disabled for items that cannot be tested (like active checks, excluded simple checks).



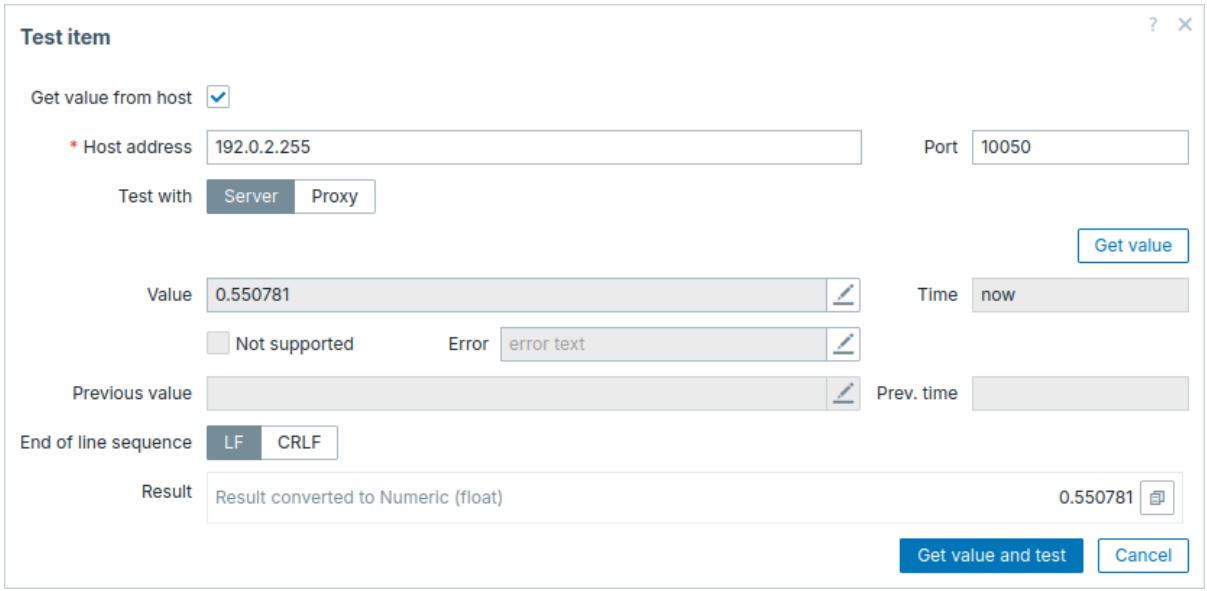
The item testing form has fields for the required host parameters (host address, port, test with server/proxy (proxy name)) and item-specific details (such as SNMPv2 community or SNMPv3 security credentials). These fields are context aware:

- The values are pre-filled when possible, i.e., for items requiring an agent, by taking the information from the selected agent interface of the host.
- The values have to be filled manually for template items.
- Plain-text macro values are resolved.
- Fields where the value (or part of the value) is a secret or Vault macro are empty and have to be entered manually. If any item parameter contains a secret macro value, the following warning message is displayed: "Item contains user-defined macros with secret values. Values of these macros should be entered manually."

- The fields are disabled when not needed in the context of the item type (e.g., the host address field and the proxy field are disabled for calculated items)

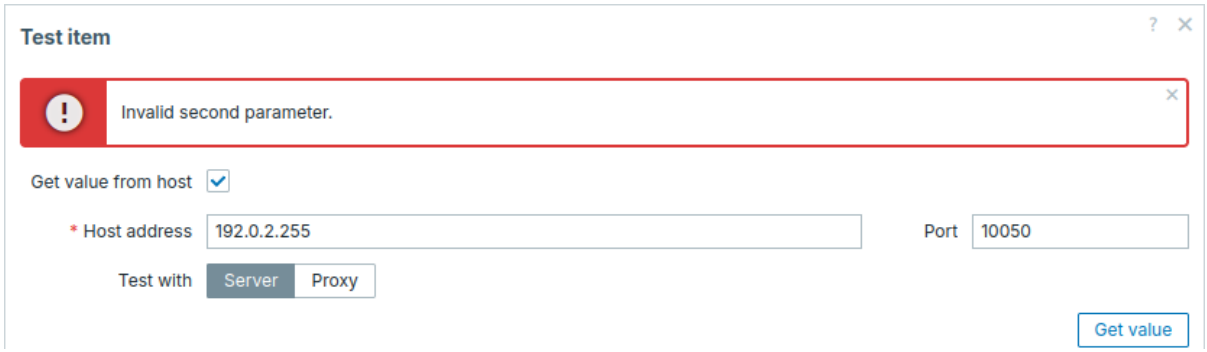
To test the item, click *Get value*. If the value is retrieved successfully, it will fill the *Value* field, moving the current value (if any) to the *Previous value* field while also calculating the *Prev. time* field, i.e., the time difference between the two values (clicks) and trying to detect an EOL sequence and switch to CRLF if detecting "\n\r" in retrieved value.

Click *Get value and test* to test the preprocessing.



Values retrieved from a host and test results are truncated to a maximum size of 512KB when sent to the frontend. If a test result is truncated, a warning icon is displayed; hover over it for details. If a value is truncated, it may fail data type validation in consecutive tests comparing large values (e.g., JSON); in such cases, an error message is displayed. Note that data larger than 512KB is still processed fully by Zabbix server.

If the configuration is incorrect, an error message is displayed describing the possible cause.



A successfully retrieved value from host can also be used to test preprocessing steps.

Form buttons

Buttons at the bottom of the form allow you to perform several operations.

Add	Add an item. This button is only available for new items.
Update	Update the properties of an item.
Clone	Create another item based on the properties of the current item.
Execute now	Execute a check for a new item value immediately. Supported for **passive** checks only (see more details). *Note* that when checking for a value immediately, configuration cache is not updated, thus the value will not reflect very recent changes to item configuration.
Test	Test if item configuration is correct by getting a value.
Clear history and trends	Delete the item history and trends.
Delete	Delete the item.
Cancel	Cancel the editing of item properties.

Item data limits

Item data limits depend on the item data type and database backend.

Numeric (unsigned) values are stored as-is. If a floating-point value is received, only the integer part is stored (e.g., 1.23 → 1).

Numeric (float) values are stored as-is and support approximately 15 to 17 digits of precision, with a range of approximately -1.79E+308 to 1.79E+308. Values in scientific notation are also supported (e.g., 1.23E+7, 1e308, 1.1E-4).

Text values are truncated before being stored to match the database value type limit:

Database	Type of information		
	Character	**Log**	**Text**
MySQL	255 characters	65536 bytes	65536 bytes
PostgreSQL	255 characters	65536 characters	65536 characters
SQLite (only Zabbix proxy)	255 characters	65536 characters	65536 characters

Binary values are limited to 16MiB (16777216 bytes). Values that exceed this limit are discarded, and a corresponding error is displayed in the frontend.

Values with the JSON data type are limited to 128MiB (134217728 bytes). Values that exceed this limit are discarded, and a corresponding error is displayed in the frontend. For MySQL/MariaDB, consider setting up daily partitions for the `history_json` table (used to store JSON values), as it can grow large quickly and lock the housekeeper for an extended period. If your use case involves large JSON values (1MiB or more), see how to configure your system for large JSON value support.

When Zabbix writes any item data to the database, it uses INSERT queries and logs them (if `DebugLevel` is set to 4 or 5). Large log entries are truncated to 64KB.

Custom script limit

Available custom script length depends on the database used:

Database	Limit in characters	Limit in bytes
MySQL	65535	65535
PostgreSQL	65535	not limited
SQLite (only Zabbix proxy)	65535	not limited

Item timeout

Item timeout specifies how long Zabbix should wait before aborting the check as failed.

When the timeout is reached, the check is aborted even if data retrieval is unfinished. If data is received partially, the item becomes unsupported and an error message is logged (for example, if data is successfully collected for only one of multiple OIDs in an SNMP check).

For many item types, you can set **flexible** item timeouts:

- per single item
- per item type (on a proxy level)
- per item type (on a global level)

A custom timeout *per single item* is useful if you want to specify a longer timeout for a specific item, while keeping timeouts for other items low.

Timeout precedence

1. Single item timeout overrides any other timeouts.
2. Proxy-level timeouts override global timeouts.

The timeout from Zabbix server or proxy configuration plays **no** role in checks for which flexible timeouts are set.

Flexible timeout support

Flexible item timeouts are supported for these item types:

- Zabbix agent (both passive and active checks)
- SNMP agent (except legacy SNMP checks[1])
- Simple check (except `icmpping*` and VMware items[2])
- SSH agent
- Telnet agent
- External check

- Database monitor
- Script
- HTTP agent
- Browser

[1] For legacy SNMP checks (single OID polling), timeout settings from the server or proxy apply.
 [2] For `icmpping*` items, the timeout value is specified directly in the item key; also a hardcoded maximum possible timeout of 600 seconds applies. VMware monitoring items use their own VMwareTimeout parameter from server or proxy configuration.

Unsupported items

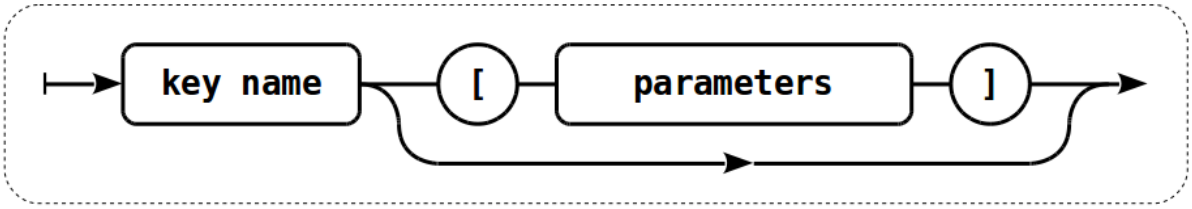
Unsupported items are reported as having a *Not supported* status and are still rechecked at their standard *Update interval*.

An item becomes unsupported if its value cannot be retrieved for some reason (e.g., connection error, no pollers configured to process the item). In addition, items that do not receive historical data remain in the *Not supported* status. This is expected behavior—an item transitions to a *Normal* status only when it receives new, valid historical data.

Like problems, unsupported items are re-evaluated only when new data is received—even if there is no historical data available for that item anymore. In other words, items and triggers change their state solely upon receiving new data. However, if a trigger expression includes a date and time and/or nodata() function, the trigger will be recalculated every 30 seconds (see triggers calculation time for details).

Item key format

Item key format, including key parameters, must follow syntax rules. The following illustrations depict the supported syntax. Allowed elements and characters at each point can be determined by following the arrows - if some block can be reached through the line, it is allowed, if not - it is not allowed.



To construct a valid item key, one starts with specifying the key name, then there's a choice to either have parameters or not - as depicted by the two lines that could be followed.

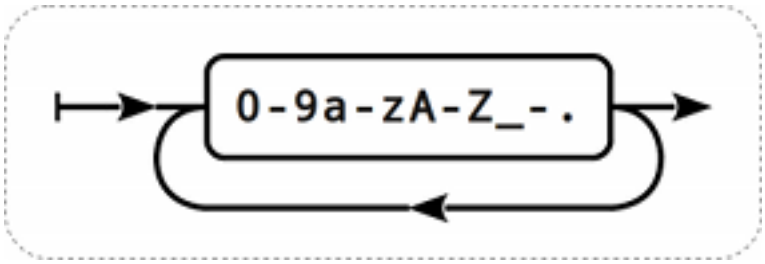
Key name

The key name itself has a limited range of allowed characters, which just follow each other. Allowed characters are:

`0-9a-zA-Z_-.`

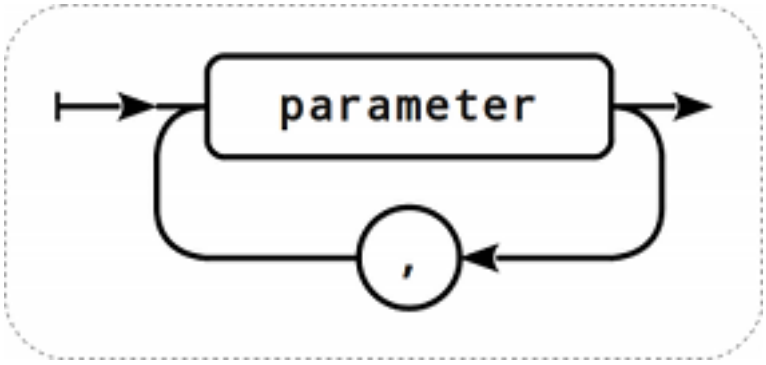
Which means:

- all numbers;
- all lowercase letters;
- all uppercase letters;
- underscore;
- dash;
- dot.

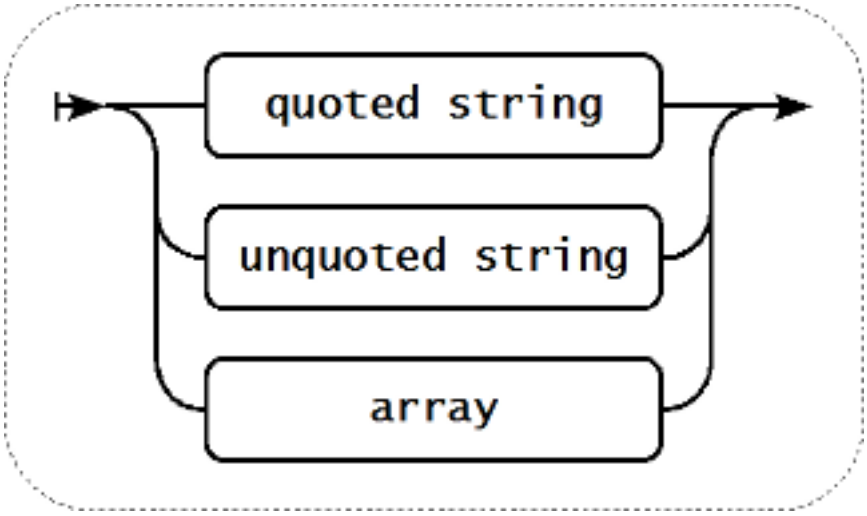


Key parameters

An item key can have multiple parameters that are comma separated.



Each key parameter can be either a quoted string, an unquoted string or an array.

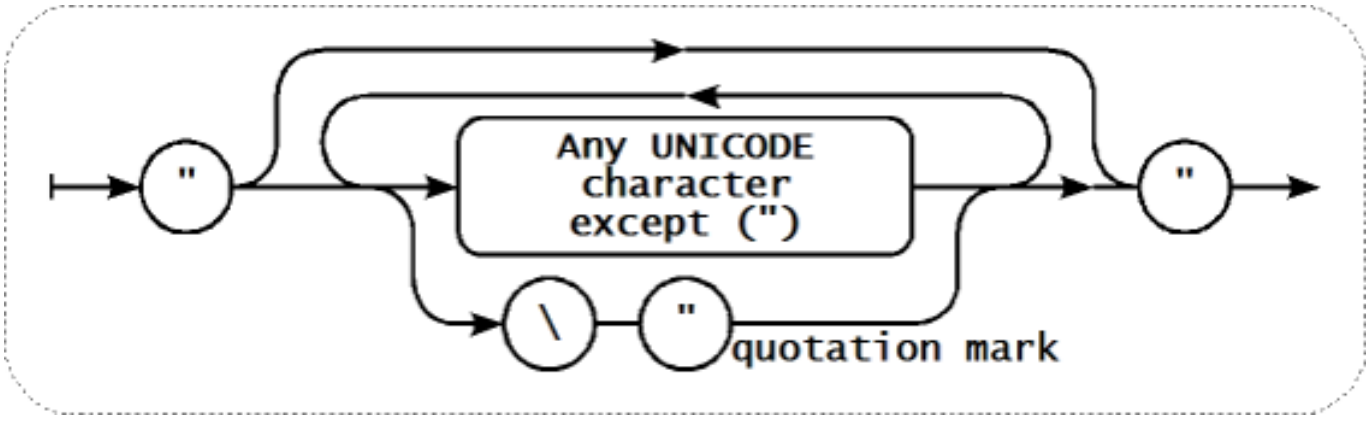


The parameter can also be left empty, thus using the default value. In that case, the appropriate number of commas must be added if any further parameters are specified. For example, item key **icmpping[„200„500]** would specify that the interval between individual pings is 200 milliseconds, timeout - 500 milliseconds, and all other parameters are left at their defaults.

It is possible to include macros in the parameters. Those can be user macros or some of the built-in macros. To see what particular built-in macros are supported in item key parameters, search the page Supported macros for ”item key parameters”.

Parameter - quoted string

If the key parameter is a quoted string, any Unicode character is allowed. If the key parameter string contains a quotation mark, this parameter has to be quoted, and each quotation mark which is a part of the parameter string has to be escaped with a backslash (\) character. If the key parameter string contains comma, this parameter has to be quoted.

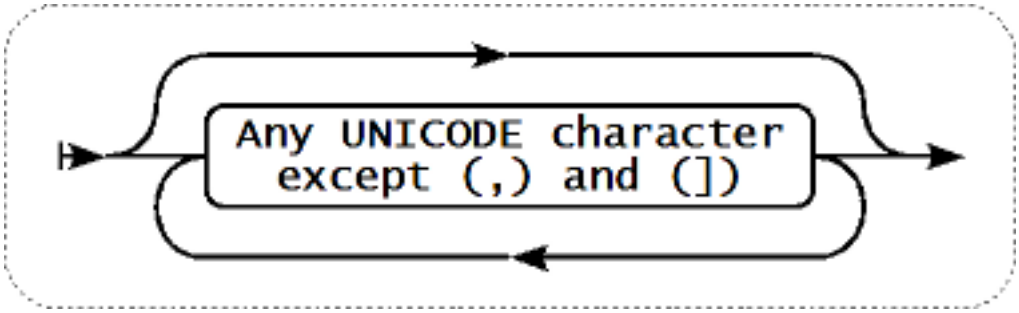


Warning:
To quote item key parameters, use double quotes only. Single quotes are not supported.

Attention:
Multi-level parameter arrays, e.g., `[a,[b,[c,d]],e]`, are not allowed.

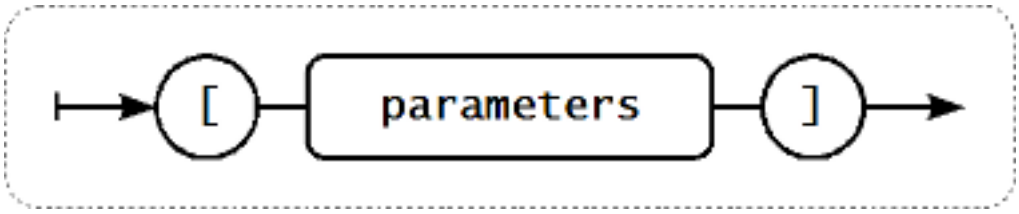
Parameter - unquoted string

If the key parameter is an unquoted string, any Unicode character is allowed except comma and right square bracket (]). Unquoted parameter cannot start with left square bracket ([).



Parameter - array

If the key parameter is an array, it is again enclosed in square brackets, where individual parameters come in line with the rules and syntax of specifying multiple parameters.



Custom intervals

Overview

It is possible to create custom rules regarding the times when an item is checked. The two methods for that are *Flexible intervals*, which allow you to redefine the default update interval, and *Scheduling*, whereby an item check can be executed at a specific time or sequence of times.

Flexible intervals

Flexible intervals allow you to redefine the default update interval for specific time periods. A flexible interval is defined with *Interval* and *Period* where:

- *Interval* - the update interval for the specified time period. Time suffixes are supported, e.g., 30s, 1m, 2h, 1d.
- *Period* - the time period when the flexible interval is active (see the time periods for detailed description of the *Period* format)

If multiple flexible intervals overlap, the smallest *Interval* value is used for the overlapping period. Note that if the smallest value of overlapping flexible intervals is '0', no polling will take place. Outside the flexible intervals the default update interval is used.

Note that if the flexible interval equals the length of the period, the item will be checked exactly once. If the flexible interval is greater than the period, the item might be checked once or it might not be checked at all (thus such configuration is not recommended). If the flexible interval is less than the period, the item will be checked at least once.

If the flexible interval is set to '0', the item is not polled during the flexible interval period and resumes polling according to the default *Update interval* once the period is over. Examples:

Interval	Period	Description
10	1-5,09:00-18:00	Item will be checked every 10 seconds during working hours.
0	1-7,00:00-7:00	Item will not be checked during the night.
0	7-7,00:00-24:00	Item will not be checked on Sundays.
60	1-7,12:00-12:01	Item will be checked at 12:00 every day. Note that this was used as a workaround for scheduled checks and it is recommended to use scheduling intervals for such checks.

Scheduling intervals

Scheduling intervals are used to check items at specific times. While flexible intervals are designed to redefine the default item update interval, the scheduling intervals are used to specify an independent checking schedule, which is executed in parallel.

A scheduling interval is defined as: `md<filter>wd<filter>h<filter>m<filter>s<filter>` where:

- **md** - month days

- **wd** - week days
- **h** - hours
- **m** - minutes
- **s** - seconds

`<filter>` is used to specify values for its prefix (days, hours, minutes, seconds) and is defined as: `[<from>[-<to>]][/<step>][,<filter`
where:

- `<from>` and `<to>` define the range of matching values (included). If `<to>` is omitted then the filter matches a `<from> - <from>` range. If `<from>` is also omitted then the filter matches all possible values.
- `<step>` defines the skips of the number value through the range. By default, `<step>` has the value of 1, which means that all values of the defined range are matched.

While the filter definitions are optional, at least one filter must be used. A filter must either have a range or the *<step>* value defined.

An empty filter matches either '0' if no lower-level filter is defined or all possible values otherwise. For example, if the hour filter is omitted then only '0' hour will match, provided minute and seconds filters are omitted too; otherwise, an empty hour filter will match all hour values.

Valid `<from>` and `<to>` values for their respective filter prefix are:

Prefix	Description	*<from>*	*<to>*
md	Month days	1-31	1-31
wd	Week days	1-7	1-7
h	Hours	0-23	0-23
m	Minutes	0-59	0-59
s	Seconds	0-59	0-59

The `<from>` value must be less or equal to `<to>` value. The `<step>` value must be greater or equal to 1 and less or equal to `<to> - <from>`.

Single digit month days, hours, minutes and seconds values can be prefixed with 0. For example, `md01-31` and `h/02` are valid intervals, but `md01-031` and `wd01-07` are not.

In Zabbix frontend, multiple scheduling intervals are entered in separate rows. In Zabbix API, they are concatenated into a single string with a semicolon `;` as a separator.

If a time is matched by several intervals it is executed only once. For example, `wd1h9;h9` will be executed only once on Monday at 9am.

Examples:

Interval	Will be executed
m0-59	every minute
h9-17/2	every 2 hours starting with 9:00 (9:00, 11:00 ...)
m0,30 or m/30	hourly at hh:00 and hh:30
m0,5,10,15,20,25,30,35,40,45,50,55 or m/5	every five minutes
wd1-5h9	every Monday till Friday at 9:00
wd1-5h9-18	every Monday till Friday at 9:00,10:00,...,18:00
h9,10,11 or h9-11	every day at 9:00, 10:00 and 11:00
md1h9m30	every 1st day of each month at 9:30
md1wd1h9m30	every 1st day of each month at 9:30 if it is Monday
h9m/30	every day at 9:00, 9:30
h9m0-59/30	every day at 9:00, 9:30
h9,10m/30	every day at 9:00, 9:30, 10:00, 10:30
h9-10m30	every day at 9:30, 10:30
h9m10-40/30	every day at 9:10, 9:40
h9,10m10-40/30	every day at 9:10, 9:40, 10:10, 10:40
h9-10m10-40/30	every day at 9:10, 9:40, 10:10, 10:40
h9m10-40	every day at 9:10, 9:11, 9:12, ... 9:40
h9m10-40/1	every day at 9:10, 9:11, 9:12, ... 9:40
h9-12,15	every day at 9:00, 10:00, 11:00, 12:00, 15:00
h9-12,15m0	every day at 9:00, 10:00, 11:00, 12:00, 15:00

Interval	Will be executed
h9-12,15m0s30	every day at 9:00:30, 10:00:30, 11:00:30, 12:00:30, 15:00:30
h9-12s30	every day at 9:00:30, 9:01:30, 9:02:30 ... 12:58:30, 12:59:30
h9m/30;h10 (*API-specific syntax*)	every day at 9:00, 9:30, 10:00
h9m/30 h10 (*add this as another row in frontend*)	every day at 9:00, 9:30, 10:00

Aligning time zones for proxies and agent

Note that Zabbix proxies and agent use their local time zones when processing scheduling intervals.

For this reason, when scheduling intervals are applied to items monitored by Zabbix proxy or agent active items, it is recommended to set the time zone of the respective proxies or agent the same as Zabbix server; otherwise, the queue may report item delays incorrectly.

The time zone for Zabbix proxy or agent can be set using the environment variable TZ in the `systemd` unit file:

```
[Service]
...
Environment="TZ=Europe/Amsterdam"
```

Item value preprocessing

Overview

Preprocessing allows you to apply transformations to the received item values before saving them to the database. These transformations/preprocessing steps are performed by Zabbix server or proxy (if items are monitored by proxy).

This feature supports a variety of use cases, such as:

- converting bytes to bits (e.g., multiplying network traffic values by ”8”);
- calculating per-second statistics for incrementally increasing values;
- applying regular expressions to extract or modify values;
- executing custom scripts on values;
- discarding unchanged values to optimize database storage.

One or more preprocessing steps can be configured for an item. These steps are executed in the order they are configured.

Attention:
If a preprocessing step fails, an item becomes unsupported. This can be avoided by *Custom on fail* error-handling (available for most transformations), allowing you to discard values or set custom values.

 For log items, log metadata (without value) will always reset the item unsupported state, making it supported again. This happens even if the initial error occurred after receiving a log value from agent.

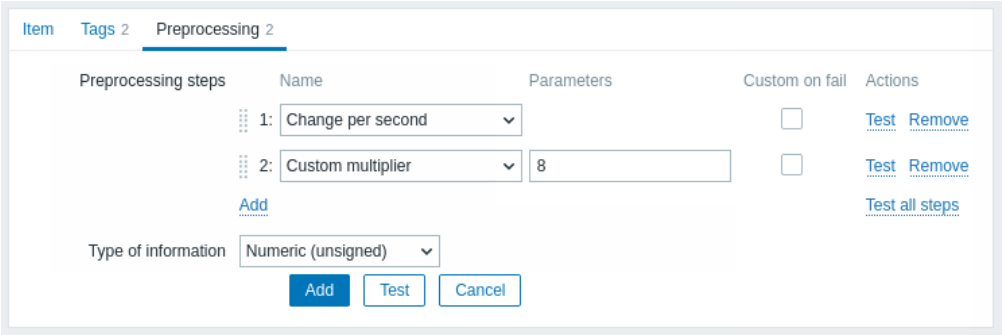
All values passed to preprocessing are initially treated as strings. Conversion to the desired value type (as defined in item configuration) is performed at the end of the preprocessing pipeline. However, specific preprocessing steps may trigger earlier conversions, if required. For detailed technical information, see Preprocessing details.

To ensure that your preprocessing configuration works as expected, you can test it.

See also: Preprocessing examples

Configuration

Preprocessing steps are defined in the **Preprocessing** tab of the item configuration form.



Click *Add* to select a supported transformation.

The *Type of information* field is displayed at the bottom of the tab when at least one preprocessing step is defined. If required, it is possible to change the type of information without leaving the *Preprocessing* tab. See Creating an item for the detailed parameter description.

Supported transformations

All supported transformations are listed below. Click the transformation name to see full details about it.

Name	Description	Type
Regular expression	Match the value to the regular expression and replace with the required output.	Text
Replace	Find the search string and replace it with another (or nothing).	
Trim	Remove specified characters from the beginning and end of the value.	
Right trim	Remove specified characters from the end of the value.	
Left trim	Remove specified characters from the beginning of the value.	
XML XPath	Extract value or fragment from XML data using XPath functionality.	Structured data
JSON Path	Extract value or fragment from JSON data using JSONPath functionality.	
CSV to JSON	Convert CSV file data into JSON format.	
XML to JSON	Convert data in XML format to JSON.	
SNMP walk value	Extract value by the specified OID/MIB name and apply formatting options.	SNMP
SNMP walk to JSON	Convert SNMP values to JSON.	
SNMP get value	Apply formatting options to the SNMP get value.	
Custom multiplier	Multiply the value by the specified integer or floating-point value.	Arithmetic
Simple change	Calculate the difference between the current and previous value.	Change
Change per second	Calculate the value change (difference between the current and previous value) speed per second.	
Boolean to decimal	Convert the value from boolean format to decimal.	Numeral systems
Octal to decimal	Convert the value from octal format to decimal.	
Hexadecimal to decimal	Convert the value from hexadecimal format to decimal.	
JavaScript	Enter JavaScript code.	Custom scripts
In range	Define a range that a value should be in.	Validation
Matches regular expression	Specify a regular expression that a value must match.	
Does not match regular expression	Specify a regular expression that a value must not match.	
Check for error in JSON	Check for an application-level error message located at JSONPath.	
Check for error in XML	Check for an application-level error message located at XPath.	
Check for error using a regular expression	Check for an application-level error message using a regular expression.	
Check for not supported value	Check if no item value could be retrieved.	
Discard unchanged	Discard a value if it has not changed.	Throttling
Discard unchanged with heartbeat	Discard a value if it has not changed within the defined time period.	
Prometheus pattern	Use the following query to extract the required data from Prometheus metrics.	Prometheus
Prometheus to JSON	Convert the required Prometheus metrics to JSON.	

If the *Custom on fail* checkbox is enabled for the selected transformation, you can mark it and specify custom error-handling

options: *Discard value* (discard the value), *Set value to* (set a specified value), or *Set error to* (set a specified error message). If a preprocessing step fails, the item will not become unsupported when either the *Discard value* or *Set value to* option is selected.

Note that for *Change* and *Throttling* preprocessing steps, Zabbix has to remember the last value to calculate/compare the new value as required. These previous values are handled by the preprocessing manager. If Zabbix server or proxy is restarted or there is any change made to preprocessing steps, the last value of the corresponding item is reset, resulting in:

- for *Simple change*, *Change per second* steps - the next value will be ignored because there is no previous value to calculate the change from;
- for *Discard unchanged*, *Discard unchanged with heartbeat* steps - the next value will never be discarded, even if it should have been because of discarding rules.

Regular expression

Match the value to the regular expression and replace with the required output.

Parameters:

- **pattern** - the regular expression;
- **output** - the output formatting template. An \N (where N=1…9) escape sequence is replaced with the Nth matched group. A \0 escape sequence is replaced with the matched text.

Comments:

- A failure to match the input value will make the item unsupported;
- The regular expression supports extraction of maximum 10 captured groups with the \N sequence;
- Please refer to regular expressions section for some existing examples.

Replace

Find the search string and replace it with another (or nothing).

Parameters:

- **search string** - the string to find and replace, case-sensitive (required);
- **replacement** - the string to replace the search string with. The replacement string may also be empty, effectively allowing the search string to be deleted when found.

Comments:

- All occurrences of the search string will be replaced;
- It is possible to use escape sequences to search for or replace line breaks, carriage return, tabs and spaces "\n \r \t \s"; backslash can be escaped as "\\" and escape sequences can be escaped as "\\n";
- Escaping of line breaks, carriage return, tabs is automatically done during low-level discovery.

Trim

Remove specified characters from the beginning and end of the value.

Right trim

Remove specified characters from the end of the value.

Left trim

Remove specified characters from the beginning of the value.

XML XPath

Extract value or fragment from XML data using XPath functionality.

Comments:

- For this option to work, Zabbix server (or Zabbix proxy) must be compiled with libxml support;
- Namespaces are not supported.

Examples:

```
number(/document/item/value) #will extract '10' from <document><item><value>10</value></item></document>
number(/document/item/@attribute) #will extract '10' from <document><item attribute="10"></item></document
/document/item #will extract '<item><value>10</value></item>' from <document><item><value>10</value></item
```

JSON Path

Extract value or fragment from JSON data using JSONPath functionality.

CSV to JSON

Convert CSV file data into JSON format.

For more information, see: CSV to JSON preprocessing.

XML to JSON

Convert data in XML format to JSON.

For more information, see: Serialization rules.

SNMP walk value

Extract value by the specified OID/MIB name and apply formatting options:

- **Unchanged** - return hex-string as unescaped hex string (*note* that display hints are still applied);
- **UTF-8 from hex-STRING** - convert hex-string to UTF-8 string;
- **MAC from hex-STRING** - validate hex-string as MAC address and return a proper MAC address string (where ' ' are replaced by ':');
- **Integer from BITS** - convert the first 8 bytes of a bit string expressed as a sequence of hex characters (e.g., "1A 2B 3C 4D") into a 64-bit unsigned integer. In bit strings longer than 8 bytes, consequent bytes will be ignored.

SNMP walk to JSON

Convert SNMP values to JSON.

Specify a field name in the JSON and the corresponding SNMP OID path. Field values will be populated by values in the specified SNMP OID path.

Comments:

- Similar value formatting options as in the *SNMP walk value* step are available;
- You may use this preprocessing step for SNMP OID discovery.

SNMP get value

Apply formatting options to the SNMP get value:

- **UTF-8 from Hex-STRING** - convert hex-string to UTF-8 string;
- **MAC from Hex-STRING** - validate hex-string as MAC address and return a proper MAC address string (where ' ' are replaced by ':');
- **Integer from BITS** - convert the first 8 bytes of a bit string expressed as a sequence of hex characters (e.g., "1A 2B 3C 4D") into a 64-bit unsigned integer. In bit strings longer than 8 bytes, consequent bytes will be ignored.

Custom multiplier

Multiply the value by the specified integer or floating-point value.

Comments:

- Use this option to convert values received in KB, MBps, etc., into B, Bps. Otherwise, Zabbix cannot correctly set prefixes (K, M, G, etc.).
- Note that if the item type of information is *Numeric (unsigned)*, incoming values with a fractional part will be trimmed (i.e., '0.9' will become '0') before the custom multiplier is applied;
- If you use a custom multiplier or store value as *Change per second* for items with the type of information set to *Numeric (unsigned)* and the resulting calculated value is actually a float number, the calculated value is still accepted as a correct one by trimming the decimal part and storing the value as an integer;
- Supported: scientific notation, for example, 1e+70; user macros and LLD macros; strings that include macros, for example, `{#MACRO}e+10`, `{$MACRO1}e+{$MACRO2}`. The macros must resolve to an integer or a floating-point number.

Simple change

Calculate the difference between the current and previous value.

Comments:

- This step can be useful to measure a constantly growing value;
- Evaluated as **value**-**prev_value**, where *value* - the current value; *prev_value* - the previously received value;
- Only one change operation per item ("Simple change" or "Change per second") is allowed;
- If the current value is smaller than the previous value, Zabbix discards that difference (stores nothing) and waits for another value.

Change per second

Calculate the value change (difference between the current and previous value) speed per second.

Comments:

- This step is useful for calculating the speed per second of a constantly growing value;
- As this calculation may produce floating-point numbers, it is recommended to set the 'Type of information' to *Numeric (float)*, even if the incoming raw values are integers. This is especially relevant for small numbers where the decimal part matters. If the floating-point values are large and may exceed the 'float' field length in which case the entire value may be lost, it is actually recommended to use *Numeric (unsigned)* and thus trim only the decimal part.
- Evaluated as (**value**-**prev_value**)/(**time**-**prev_time**), where *value* - the current value; *prev_value* - the previously received value; *time* - the current timestamp; *prev_time* - the timestamp of the previous value;
- Only one change operation per item ("Simple change" or "Change per second") is allowed;
- If the current value is smaller than the previous value, Zabbix discards that difference (stores nothing) and waits for another value. This helps to work correctly with, for instance, a wrapping (overflow) of 32-bit SNMP counters.

Boolean to decimal

Convert the value from boolean format to decimal.

Comments:

- The textual representation is translated into either 0 or 1. Thus, 'TRUE' is stored as 1 and 'FALSE' is stored as 0. All values are matched in a case-insensitive way. Currently recognized values are, for *TRUE* - true, t, yes, y, on, up, running, enabled, available, ok, master; for *FALSE* - false, f, no, n, off, down, unused, disabled, unavailable, err, slave. Additionally, any non-zero numeric value is considered to be TRUE and zero is considered to be FALSE.

Octal to decimal

Convert the value from octal format to decimal.

Hexadecimal to decimal

Convert the value from hexadecimal format to decimal.

JavaScript

Enter JavaScript code in the modal editor that opens when clicking in the parameter field or on the pencil icon next to it.

> **Attention:**
> Do not use undeclared assignments in preprocessing JavaScript. Use `var` to declare local variables.

Comments:

- The available JavaScript length depends on the database used;
- For more information, see: Javascript preprocessing.

In range

Define a range that a value should be in by specifying minimum/maximum values (inclusive).

Comments:

- Numeric values are accepted (including any number of digits, optional decimal part and optional exponential part, negative values);
- The minimum value should be less than the maximum;
- At least one value must exist;
- User macros and low-level discovery macros can be used.

Matches regular expression

Specify a regular expression that a value must match.

Does not match regular expression

Specify a regular expression that a value must not match.

Check for error in JSON

Check for an application-level error message located at JSONPath. Stop processing if succeeded and the message is not empty; otherwise, continue processing with the value that was before this preprocessing step.

Comments:

- These external service errors are reported to the user as is, without adding preprocessing step information;
- No error will be reported in case of failing to parse invalid JSON.

Check for error in XML

Check for an application-level error message located at XPath. Stop processing if succeeded and the message is not empty; otherwise, continue processing with the value that was before this preprocessing step.

Comments:

- These external service errors are reported to the user as is, without adding preprocessing step information;
- No error will be reported in case of failing to parse invalid XML.

Check for error using a regular expression

Check for an application-level error message using a regular expression. Stop processing if succeeded and the message is not empty; otherwise, continue processing with the value that was before this preprocessing step.

Parameters:

- **pattern** - the regular expression;
- **output** - the output formatting template. An \N (where N=1…9) escape sequence is replaced with the Nth matched group. A \0 escape sequence is replaced with the matched text.

Comments:

- These external service errors are reported to the user as is, without adding preprocessing step information.

Check for not supported value

Check if no item value could be retrieved. Specify how the failure should be processed, based on inspecting the returned error message.

Parameters:

- **scope** - select the error processing scope:
any error - any error;
error matches - only the error that matches the regular expression specified in *pattern*;
error does not match - only the error that does not match the regular expression specified in *pattern*
- **pattern** - the regular expression to match the error to. If *any error* is selected in the scope parameter, this field is not displayed. If displayed, this field is mandatory.

Comments:

- Normally, the absence/failure to retrieve a value would lead to the item becoming unsupported. This preprocessing step allows you to modify this behavior.
- This preprocessing step only checks if no item value could be retrieved. It does not check, for example, if the type of the retrieved value (e.g., string) matches the item's type of information (e.g., numeric); for details, see Preprocessing examples. If there is a type mismatch, the item may still become unsupported after all preprocessing steps are executed. To check for a type mismatch, you may use, for example, the *Custom multiplier* preprocessing step; see Preprocessing examples.
- Capturing regular expression groups is supported in the *Set value to* or *Set error to* fields. Use \N (where N=1…9) to retrieve the Nth matched group; use \0 to retrieve the matched text;
- These steps are always executed as the first preprocessing steps and are placed above all others after saving changes to the item;
- Multiple *Check for not supported value* steps are supported, in the specified order. A step for *any error* will be automatically placed as the last step in this group.

Discard unchanged

Discard a value if it has not changed.

Comments:

- If a value is discarded, it is not saved in the database and Zabbix server has no knowledge that this value was received. No trigger expressions will be evaluated, as a result, no problems for related triggers will be created/resolved. Functions will work only based on data that is actually saved in the database. As trends are built based on data in the database, if there is no value saved for an hour then there will also be no trends data for that hour.
- Only one throttling option can be specified per item.

Discard unchanged with heartbeat

Discard a value if it has not changed within the defined time period (in seconds).

Comments:

- Positive integer values are supported to specify the seconds (minimum - 1 second);
- Time suffixes can be used (e.g., `30s`, `1m`, `2h`, `1d`);
- User macros and low-level discovery macros can be used;

- If a value is discarded, it is not saved in the database and Zabbix server has no knowledge that this value was received. No trigger expressions will be evaluated, as a result, no problems for related triggers will be created/resolved. Functions will work only based on data that is actually saved in the database. As trends are built based on data in the database, if there is no value saved for an hour then there will also be no trends data for that hour.
- Only one throttling option can be specified per item.

Prometheus pattern

Use the following query to extract the required data from Prometheus metrics.

See Prometheus checks for more details.

Prometheus to JSON

Convert the required Prometheus metrics to JSON.

See Prometheus checks for more details.

Macro support

User macros and user macros with context are supported in:

- preprocessing step parameters, including JavaScript code;
- custom error-handling parameters (*Set value to* and *Set error to* fields).

Note:
The macro context is ignored when a macro is replaced with its value. The macro value is inserted in the code as is, it is not possible to add additional escaping before placing the value in the JavaScript code. Please be advised that this can cause JavaScript errors in some cases.

Testing

See preprocessing testing.

Preprocessing testing

Testing

Testing preprocessing steps is useful to make sure that complex preprocessing pipelines yield the results that are expected from them, without waiting for the item value to be received and preprocessed.

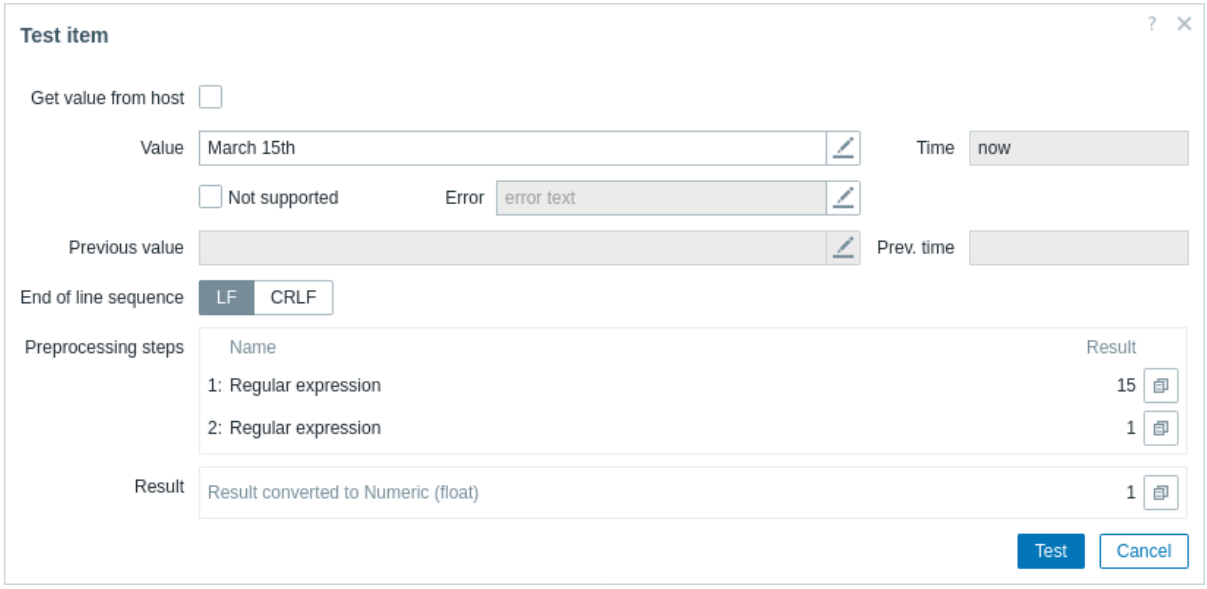


It is possible to test:

- against a hypothetical value
- against a real value from a host

Each preprocessing step can be tested individually as well as all steps can be tested together. When you click the *Test* or *Test all steps* button respectively in the Actions block, a testing window is opened.

Testing hypothetical value



Parameter	Description
Get value from host	If you want to test a hypothetical value, leave this checkbox unmarked. See also: Testing real value.
Value	Enter the input value to test. Clicking in the parameter field or on the view/edit button will open a text area window for entering the value or code block.
Not supported	Mark this checkbox to test an unsupported value. This option is useful to test the *Check for not supported value* preprocessing step.
Error	Enter the error text. This field is enabled when *Get value from host* is unchecked, but *Not supported* is checked. If *Get value from host* is checked, this field gets filled with the actual error message (read-only) from the host.
Time	Time of the input value is displayed: `now` (read-only).
Previous value	Enter a previous input value to compare to. Only for *Change* and *Throttling* preprocessing steps.
Previous time	Enter the previous input value time to compare to. Only for *Change* and *Throttling* preprocessing steps. The default value is based on the 'Update interval' field value of the item (if '1m', then this field is filled with `now-1m`). If nothing is specified or the user has no access to the host, the default is `now-30s`.
Macros	If any macros are used, they are listed along with their values. The values are editable for testing purposes, but the changes will only be saved within the testing context.
End of line sequence	Select the end of line sequence for multiline input values: **LF** - LF (line feed) sequence **CRLF** - CRLF (carriage-return line-feed) sequence.
Preprocessing steps	Preprocessing steps are listed; the testing result is displayed for each step after the *Test* button is clicked. Test results are truncated to a maximum size of 512KB when sent to the frontend. Test results can be copied (not more than the truncated 512KB). If a result is truncated, a warning icon is displayed. The warning description is displayed on mouseover. Note that data larger than 512KB is still processed fully by Zabbix server. If the step failed in testing, an error icon is displayed. The error description is displayed on mouseover. In case "Custom on fail" is specified for the step and that action is performed, a new line appears right after the preprocessing test step row, showing what action was done and what outcome it produced (error or value).

Parameter	Description
Result	The final result of testing preprocessing steps is displayed in all cases when all steps are tested together (when you click the *Test all steps* button). The type of conversion to the value type of the item is also displayed, for example, `Result converted to Numeric (unsigned)`. Test results are truncated to a maximum size of 512KB when sent to the frontend. The final result can be copied (not more than the truncated 512KB). If a result is truncated, a warning icon is displayed. The warning description is displayed on mouseover. Note that data larger than 512KB is still processed fully by Zabbix server.

Click *Test* to see the result after each preprocessing step.

Test values are stored between test sessions for either individual steps or all steps, allowing the user to change preprocessing steps or item configuration and then return to the testing window without having to re-enter information. Values are lost on a page refresh though.

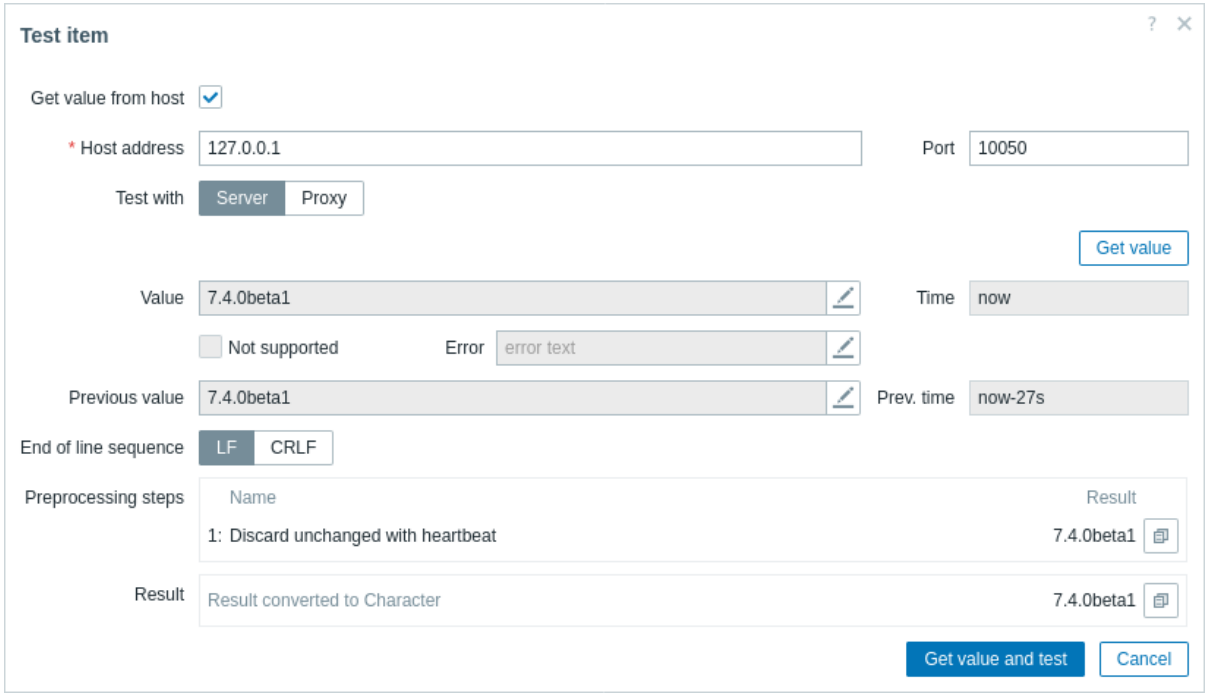
The testing is done by Zabbix server. The frontend sends a corresponding request to the server and waits for the result. The request contains the input value and preprocessing steps (with expanded user macros). For *Change* and *Throttling* steps, an optional previous value and time can be specified. The server responds with results for each preprocessing step.

All technical errors or input validation errors are displayed in the error box at the top of the testing window.

Testing real value

To test preprocessing against a real value:

- Mark the *Get value from host* checkbox
- Enter or verify host parameters (host address, port, proxy name/no proxy) and item-specific details (such as SNMPv2 community or SNMPv3 security credentials). These fields are context-aware:
  - The values are pre-filled when possible, i.e., for items requiring an agent, by taking the information from the selected agent interface of the host
  - The values have to be filled manually for template items
  - Plain-text macro values are resolved
  - Where the field value (or part of the value) is a secret or Vault macro, the field will be empty and has to be filled out manually. If any item parameter contains a secret macro value, the following warning message is displayed: "Item contains user-defined macros with secret values. Values of these macros should be entered manually."
  - The fields are disabled when not needed in the context of the item type (e.g., the host address and the proxy fields are disabled for calculated items)
- Click *Get value and test* to test the preprocessing



If you have specified a value mapping in the item configuration form (the *Show value* field), the item test dialog will show another line after the final result, named *Result with value map applied*.

Parameters that are specific to getting a real value from a host:

Parameter	Description
Get value from host	Mark this checkbox to get a real value from the host.
Host address	Enter the host address. This field is automatically filled by the address of the item host interface.
Port	Enter the host port. This field is automatically filled by the port of item host interface.
Additional fields for SNMP interfaces (SNMP version, SNMP community, Context name, etc.)	See Configuring SNMP monitoring for additional details on configuring an SNMP interface (v1, v2 and v3). These fields are automatically filled from the item host interface.
Proxy	Specify the proxy if the host is monitored by a proxy. This field is automatically filled by the proxy of the host (if any).
Value	Value retrieved from the host. Clicking in the parameter field or on the view/edit button will open a text area window of the value or code block. Values are truncated to a maximum size of 512KB and only in the frontend. If a result is truncated, a warning icon is displayed. The warning description is displayed on mouseover. Note that data larger than 512KB is still processed fully by Zabbix server.

For the rest of the parameters, see Testing hypothetical value above.

Preprocessing details

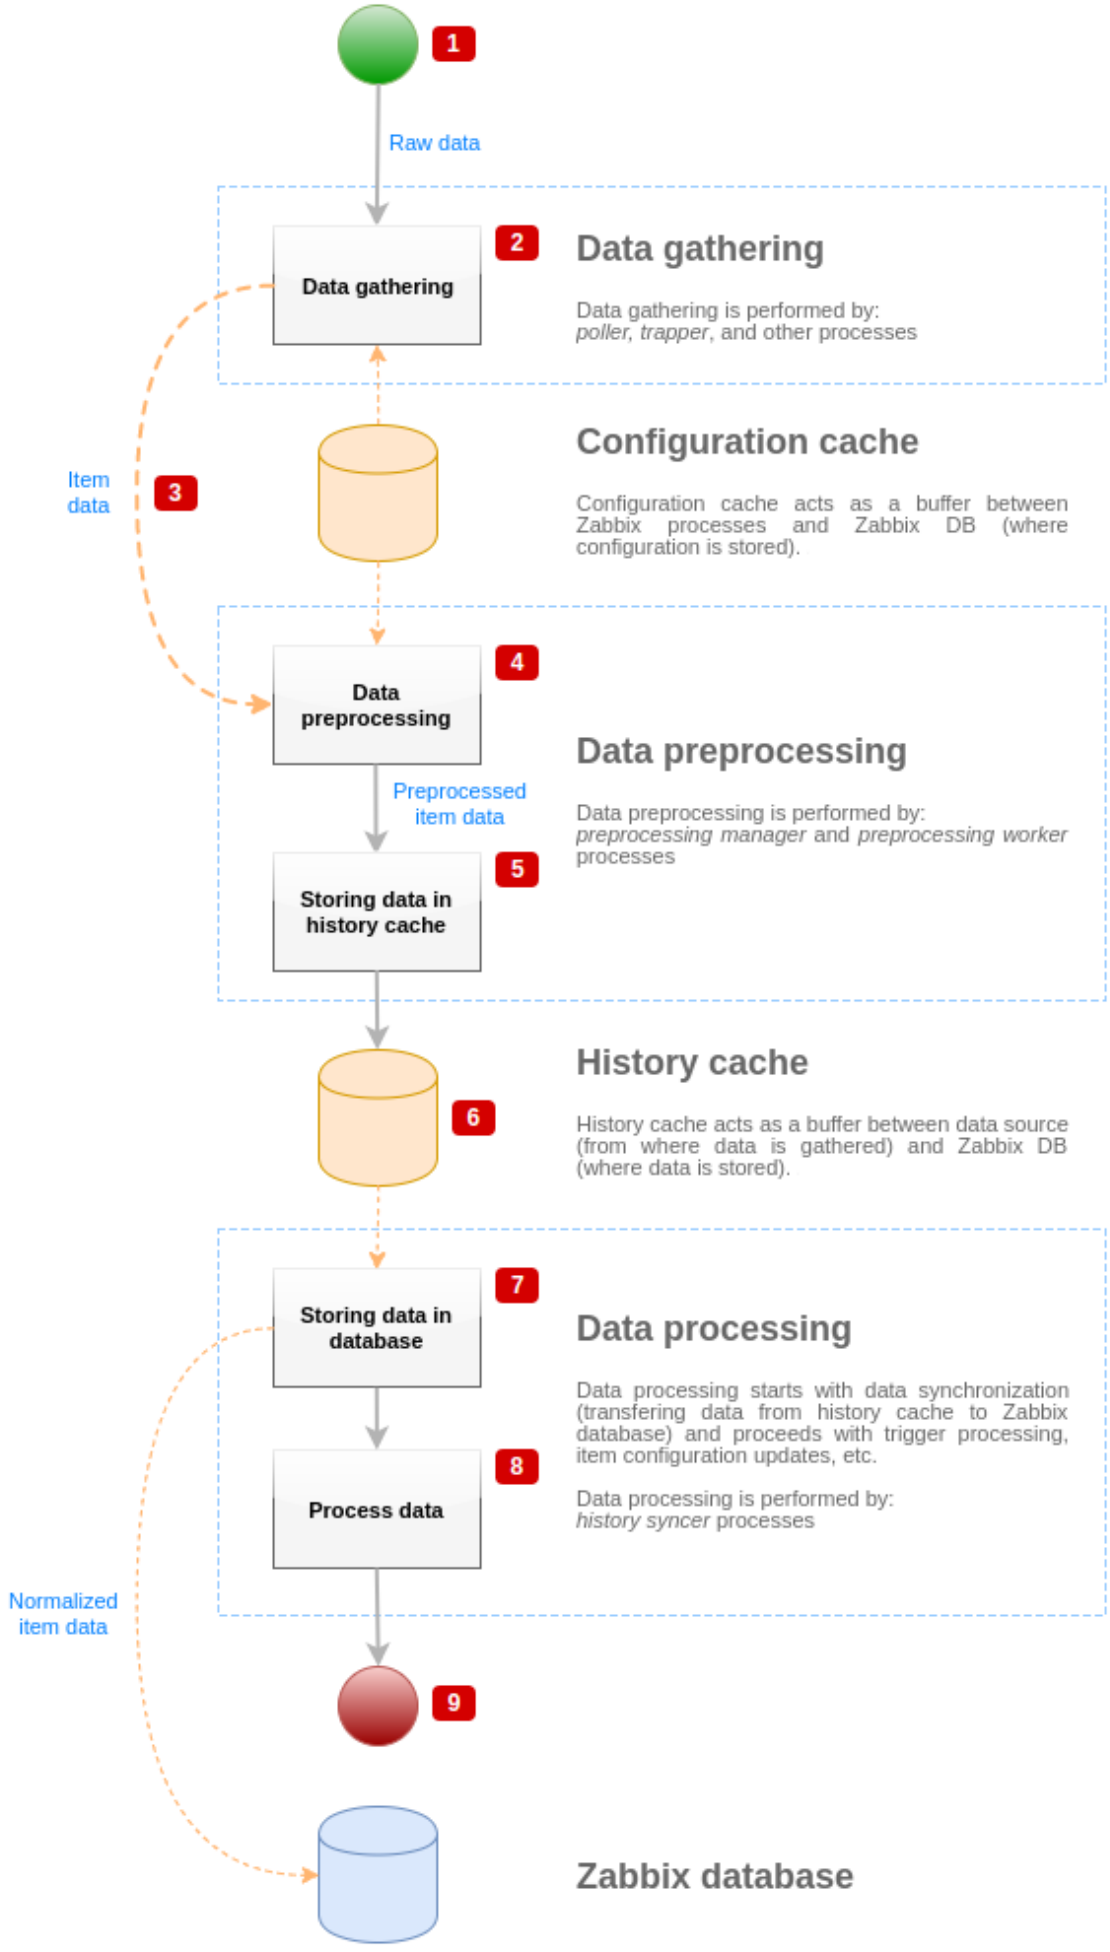
Overview

This section provides item value preprocessing details. The item value preprocessing allows you to define and execute transformation rules for the received item values.

Preprocessing is managed by the preprocessing manager process along with preprocessing workers that perform the preprocessing steps. All values with preprocessing, received from different data gatherers, pass through the preprocessing manager before being added to the history cache. Socket-based IPC communication is used between data gatherers (pollers, trappers, etc.) and the preprocessing process. Either Zabbix server or Zabbix proxy (for the items monitored by the proxy) performs the preprocessing steps.

Item value processing

To visualize the data flow from data source to the Zabbix database, we can use the following simplified diagram:


1
Raw data
Data gathering
2
Data gathering
Data gathering is performed by:
poller, trapper, and other processes
Item data
3
Configuration cache
Configuration cache acts as a buffer between Zabbix processes and Zabbix DB (where configuration is stored).
Data preprocessing
4
Data preprocessing
Data preprocessing is performed by:
preprocessing manager and preprocessing worker processes
Preprocessed item data
Storing data in history cache
5
History cache
6
History cache acts as a buffer between data source (from where data is gathered) and Zabbix DB (where data is stored).
Storing data in database
7
Data processing
Data processing starts with data synchronization (transfering data from history cache to Zabbix database) and proceeds with trigger processing, item configuration updates, etc.
Process data
8
Data processing is performed by:
history syncer processes
Normalized item data
9
Zabbix database

The diagram above shows only processes, objects and actions related to item value processing in a **simplified** form. The diagram does not show conditional direction changes, error handling or loops. The local data cache of the preprocessing manager is not shown either because it doesn't affect the data flow directly. The aim of this diagram is to show processes involved in the item value processing and the way they interact.

- Data gathering starts with raw data from a data source. At this point, the data contains only ID, timestamp and value (can be multiple values as well).
- No matter what type of data gatherer is used, the idea is the same for active or passive checks, for trapper items, etc., as it only changes the data format and the communication starter (either data gatherer is waiting for a connection and data, or data gatherer initiates the communication and requests the data). The raw data is validated, the item configuration is retrieved from the configuration cache (data is enriched with the configuration data).
- A socket-based IPC mechanism is used to pass data from data gatherers to the preprocessing manager. At this point the data gatherer continues to gather data without waiting for the response from preprocessing manager.
- Data preprocessing is performed. This includes the execution of preprocessing steps and dependent item processing.

Note:
An item can change its state to NOT SUPPORTED while preprocessing is performed if any of preprocessing steps fail.

- The history data from the local data cache of the preprocessing manager is being flushed into the history cache.
- At this point the data flow stops until the next synchronization of history cache (when the history syncer process performs data synchronization).
- The synchronization process starts with data normalization before storing data in Zabbix database. The data normalization performs conversions to the desired item type (type defined in item configuration), including truncation of textual data based on predefined sizes allowed for those types (HISTORY_STR_VALUE_LEN for string, HISTORY_TEXT_VALUE_LEN for text and HISTORY_LOG_VALUE_LEN for log values). The data is being sent to the Zabbix database after the normalization is done.

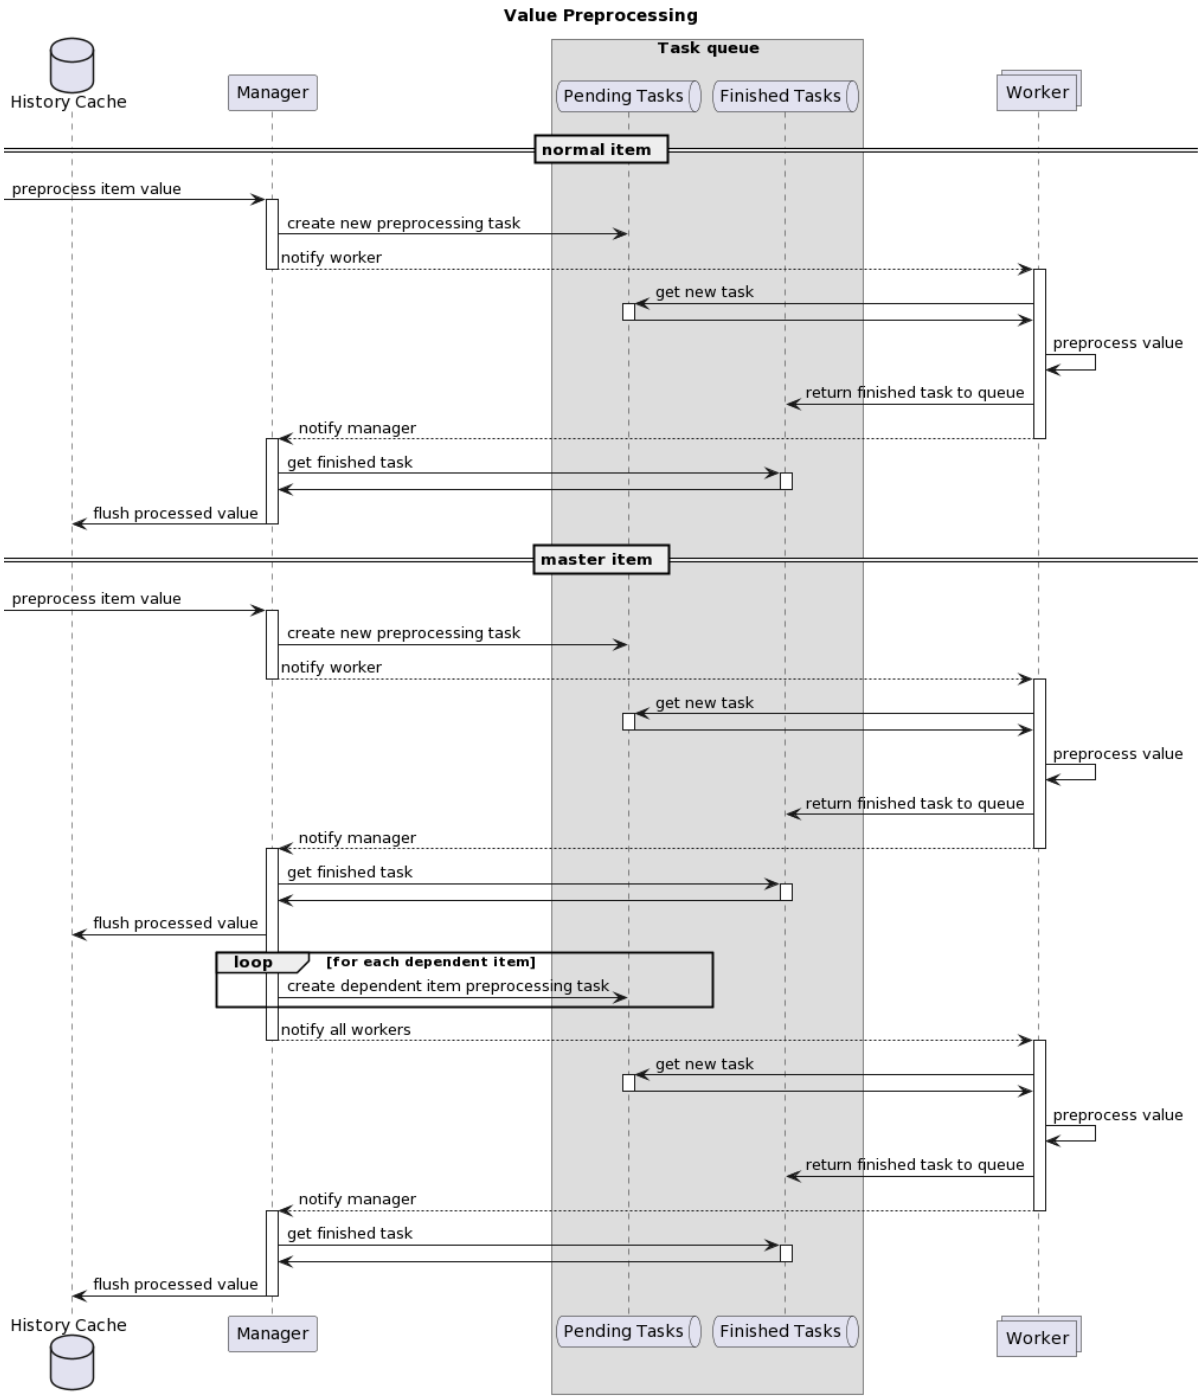
Note:
An item can change its state to NOT SUPPORTED if data normalization fails (for example, when a textual value cannot be converted to number).

- The gathered data is being processed - triggers are checked, the item configuration is updated if item becomes NOT SUPPORTED, etc.
- This is considered the end of data flow from the point of view of item value processing.

Item value preprocessing

Data preprocessing is performed in the following steps:

- If the item has neither preprocessing nor dependent items, its value is either added to the history cache or sent to the LLD manager. Otherwise, the item value is passed to the preprocessing manager using a UNIX socket-based IPC mechanism.
- A preprocessing task is created and added to the queue and preprocessing workers are notified about the new task.
- At this point the data flow stops until there is at least one unoccupied (i.e., not executing any tasks) preprocessing worker.
- When a preprocessing worker is available, it takes the next task from the queue.
- After the preprocessing is done (both failed and successful execution of preprocessing steps), the preprocessed value is added to the finished task queue and the manager is notified about a new finished task.
- The preprocessing manager converts the result to desired format (defined by item value type) and either adds it to the history cache or sends to the LLD manager.
- If there are dependent items for the processed item, then dependent items are added to the preprocessing queue with the preprocessed master item value. Dependent items are enqueued bypassing the normal value preprocessing requests, but only for master items with the value set and not in a NOT SUPPORTED state.



Note that in the diagram the master item preprocessing is slightly simplified by skipping the preprocessing caching.

Preprocessing queue

The preprocessing queue is organized as:

- the list of pending tasks:
  - tasks created directly from value preprocessing requests in the order they were received
- the list of immediate tasks (processed before pending tasks):
  - testing tasks (created in response to item/preprocessing testing requests by the frontend)
  - dependent item tasks
  - sequence tasks (tasks that must be executed in a strict order):
    - having preprocessing steps using the last value:
      - change
      - throttling
      - JavaScript (bytecode caching)
    - dependent item preprocessing caching
- the list of finished tasks

Preprocessing caching

Preprocessing caching was introduced to improve the preprocessing performance for multiple dependent items having similar preprocessing steps (which is a common LLD outcome).

Caching is done by preprocessing one dependent item and reusing some of the internal preprocessing data for the rest of the dependent items. The preprocessing cache is supported only for the first preprocessing step of the following types:

- Prometheus pattern (indexes input by metrics)
- JSONPath (parses the data into object tree and indexes the first expression `[?(@.path == "value")]`)

Preprocessing workers

The Zabbix server configuration file allows users to set the count of preprocessing worker threads. The StartPreprocessors configuration parameter should be used to set the number of pre-started instances of preprocessing workers, which should at least match the number of available CPU cores.

If preprocessing tasks are not CPU-bound and involve frequent network requests, configuring additional workers is recommended. The optimal number of preprocessing workers can be determined by many factors, including the count of "preprocessable" items (items that require to execute any preprocessing steps), the count of data gathering processes, the average step count for item preprocessing, etc. Insufficient workers can lead to high memory usage. For troubleshooting excessive memory usage on your Zabbix installation, see Profiling excessive memory usage with tcmalloc.

But assuming that there are no heavy preprocessing operations like parsing large XML/JSON chunks, the number of preprocessing workers can match the total number of data gatherers. This way, there will mostly (except for the cases when data from the gatherer comes in bulk) be at least one unoccupied preprocessing worker for collected data.

Warning:
Too many data gathering processes (pollers, unreachable pollers, ODBC pollers, HTTP pollers, Java pollers, pingers, trappers, proxypollers) together with IPMI manager, SNMP trapper and preprocessing workers can exhaust the per-process file descriptor limit for the preprocessing manager.

 Exhausting the per-process file descriptor limit will cause Zabbix server to stop, typically shortly after startup but sometimes taking longer. To avoid such issues, review the Zabbix server configuration file to optimize the number of concurrent checks and processes. Additionally, if necessary, ensure that the file descriptor limit is set sufficiently high by checking and adjusting system limits.

Value processing pipeline

Item value processing is executed in multiple steps (or phases) by multiple processes. This can cause:

- A dependent item can receive values, while THE master value cannot. This can be achieved by using the following use case:
  - Master item has value type `UINT` (trapper item can be used), dependent item has value type `TEXT`.
  - No preprocessing steps are required for both master and dependent items.
  - Textual value (for example, "abc") should be passed to master item.
  - As there are no preprocessing steps to execute, preprocessing manager checks if master item is not in NOT SUPPORTED state and if value is set (both are true) and enqueues dependent item with the same value as master item (as there are no preprocessing steps).
  - When both master and dependent items reach history synchronization phase, master item becomes NOT SUPPORTED because of the value conversion error (textual data cannot be converted to unsigned integer).

As a result, the dependent item receives a value, while the master item changes its state to NOT SUPPORTED.

- A dependent item receives value that is not present in the master item history. The use case is very similar to the previous one, except for the master item type. For example, if `CHAR` type is used for master item, then master item value will be truncated at the history synchronization phase, while dependent items will receive their values from the initial (not truncated) value of the master item.

Preprocessing examples

Overview

This section presents examples of using preprocessing steps to accomplish some practical tasks.

Filtering VMware event log records

This example uses the Matches regular expression preprocessing step to filter unnecessary events from the VMware event log.

1. On a working VMware Hypervisor host, check that the vmware.eventlog item is present and working properly. Note that the event log item could already be present on the hypervisor if a VMware template has been linked during the host creation.

2. On the VMware Hypervisor host, create a dependent item of *Log* type and set the event log item as its master item.

3. In the *Preprocessing* tab of the dependent item, click *Add* to create a preprocessing step and select *Matches regular expression* from the drop-down. Then, specify one of the following patterns:

- For filtering all log events:

```
.* logged in .*
```

- For filtering lines containing usernames after "User":

```
\bUser\s+\K\S+
```

Attention:
If the regular expression is not matched, then the dependent item becomes unsupported with a corresponding error message. To avoid this, mark the *Custom on fail* checkbox and select an option such as discarding the value or setting a custom one. Please note that discarded values are not stored in the database; as a result, triggers are not evaluated and trend data is not generated.

Alternatively, you may use the Regular expression preprocessing step to extract matching groups and control output:

- For extracting and outputting the entire log event containing "logged in", specify the following parameters:

```
Pattern: .*logged in.*
Output: \0
```

- For extracting and outputting usernames following "User":

```
Pattern: User (.*?)(?=\ )
Output: \1
```

Checking retrieved value type

This example uses the Custom multiplier preprocessing step to check if the retrieved item value type is numeric.

In the *Preprocessing* tab of an item, select the *Custom multiplier* preprocessing step and specify the following parameter (multiplies the retrieved value by 1):

```
1
```

Attention:
If preprocessing fails (e.g., input is not numeric), then the item becomes unsupported with a corresponding error message. To avoid this, mark the *Custom on fail* checkbox and select an option such as discarding the value or setting a custom one. Please note that discarded values are not stored in the database; as a result, triggers are not evaluated and trend data is not generated.

Checking for not supported value

This example uses the Check for not supported value preprocessing step to check if the item value could not be retrieved.

When a Zabbix server/proxy poller process attempts to collect an item value, it may:

- Return a valid result.
- Return a result that initially seems valid but may become unsupported later (e.g., due to a value type mismatch after preprocessing).
- Return an error of collecting the value, causing the item to become unsupported. Common causes include:
  - Unknown item key (for Zabbix agent, Simple check, or Zabbix internal items)
  - Unknown OID (SNMP agent), unknown sensor (IPMI agent), or no JMX metric (JMX agent)
  - Cannot read trap file (SNMP trap)
  - Script not found (External check)
  - No such URL (HTTP agent, Browser)
  - Login failed (SSH agent, TELNET agent)
  - Invalid formula syntax (Calculated), JavaScript syntax error (Script), or invalid SQL (Database monitor)

To detect and handle errors of collecting item values, you can use the *Check for not supported value* preprocessing step. Note that this step is always executed first and only detects errors that occur before preprocessing begins.

In the *Preprocessing* tab of an item, select the *Check for not supported value* preprocessing step and specify one of the following parameters:

- For any errors:

```
Parameter: any error
```

- For errors containing ”cannot connect”:

```
Parameter: error matches
Pattern: (?i)cannot connect
```

Then, use the *Custom on fail* option to discard the value (in this case, the error), set a custom value, or return a custom error message. Please note that discarded values are not stored in the database; as a result, triggers are not evaluated and trend data is not generated.

JSONPath functionality

Overview

This section outlines the supported JSONPath functionality within item value preprocessing steps.

JSONPath is composed of segments separated by dots. A segment can take the form of a simple word, representing a JSON value name, the wildcard character (*), or a more intricate construct enclosed within square brackets. The dot before a bracketed segment is optional and can be omitted.

JSONPath example	Description
`$.object.name`	Return `object.name` contents.
`$.object['name']`	Return `object.name` contents.
`$.object.['name']`	Return `object.name` contents.
`$["object"]['name']`	Return `object.name` contents.
`$.['object'].["name"]`	Return `object.name` contents.
`$.object.history.length()`	Return the number of `object.history` array elements.
`$[?(@.name == 'Object')].price.first()`	Return the value of the `price` property from the first object named ”Object”.
`$[?(@.name == 'Object')].history.first().length()`	Return the number of history array elements from the first object named ”Object”.
`$[?(@.price > 10)].length()`	Return the number of objects with a price greater than 10.

See also: Escaping special characters from LLD macro values in JSONPath.

Supported segments

Segment	Description
`<name>`	Match object property by name.
`*`	Match all object properties.
`['<name>']`	Match object property by name.
`['<name>', '<name>', ...]`	Match object property by any of the listed names.
`[<index>]`	Match array element by index.
`[<number>, <number>, ...]`	Match array element by any of the listed indexes.
`[*]`	Match all object properties or array elements.
`[<start>:<end>]`	Match array elements by the defined range: **<start>** - the first index to match (including); if not specified, matches all array elements from the beginning; if negative, specifies starting offset from the end of array; **<end>** - the last index to match (excluding); if not specified, matches all array elements to the end; if negative, specifies starting offset from the end of array.
`[?(<expression>)]`	Match objects/array elements by applying a filter expression.

To find a matching segment ignoring its ancestry (detached segment), it must be prefixed with two dots (`..`). For example, `$..name` or `$..['name']` return values of all `name` properties.

Matched element names can be extracted by adding a tilde (~) suffix to the JSONPath. It returns the name of the matched object or an index in string format of the matched array item. The output format follows the same rules as other JSONPath queries - definite path results are returned ’as is’, and indefinite path results are returned in an array. However, there is minimal value in extracting the name of an element that matches a definitive path, as it is already known.

Filter expression

The filter expression is an arithmetical expression in infix notation.

Supported operands:

Operand	Description
"<text>" '<text>'	Text constant. Example: 'value: \'1\'' "value: '1'"
<number>	Numeric constant supporting scientific notation. Example: 123
<jsonpath starting with $>	Value referred to by the JSONPath from the input document root node; only definite paths are supported. Example: $.object.name
<jsonpath starting with @>	Value referred to by the JSONPath from the current object/element; only definite paths are supported. Example: @.name

Supported operators:

Operator	Type	Description	Result
-	Binary	Subtraction	Number
+	Binary	Addition	Number
/	Binary	Division	Number
*	Binary	Multiplication	Number
==	Binary	Equality	Boolean (1/0)
!=	Binary	Inequality	Boolean (1/0)
	Binary	Less than	Boolean (1/0)
<=	Binary	Less than or equal to	Boolean (1/0)
>	Binary	Greater than	Boolean (1/0)
>=	Binary	Greater than or equal to	Boolean (1/0)
=~	Binary	Matches regular expression	Boolean (1/0)
!	Unary	Boolean NOT	Boolean (1/0)
||	Binary	Boolean OR	Boolean (1/0)
&&	Binary	Boolean AND	Boolean (1/0)

Functions

Functions can be used at the end of JSONPath. Multiple functions can be chained if the preceding function returns value that is accepted by the following function.

Supported functions:

Function	Description	Input	Output
avg	Average value of numbers in an input array	Array of numbers	Number
min	Minimum value of numbers in an input array	Array of numbers	Number
max	Maximum value of numbers in an input array	Array of numbers	Number
sum	Sum of numbers in an input array	Array of numbers	Number
length	Number of elements in an input array	Array	Number
first	The first element of an array	Array	A JSON construct (object, array, value) depending on the contents of the input array

JSONPath aggregate functions accept quoted numeric values. These values are automatically converted from strings to numeric types when aggregation is needed. Incompatible input will cause the function to generate an error.

Output value

JSONPaths can be divided into definite and indefinite paths. A definite path can return only null or a single match. An indefinite path can return multiple matches: JSONPaths with detached, multiple name/index lists, array slices, or expression segments. However, when a function is used, the JSONPath becomes definite, as functions always output a single value.

A definite path returns the object/array/value it is referencing. In contrast, an indefinite path returns an array of the matched objects/arrays/values.

> **Attention:**
> The property order in JSONPath query results may not align with the original JSON property order due to internal optimization methods. For example, the JSONPath `$.books[1]["author", "title"]` may return `["title", "author"]`. If preserving the original property order is essential, alternative post-query processing methods should be considered.

Path formatting rules

Whitespaces (space, tab character) can be used in bracket notation segments and expressions, for example: `$[ 'a' ][ 0 ][ ?( $.b == 'c' ) ][ : -1 ].first( )`.

Strings should be enclosed with single (') or double (") quotes. Inside the strings, single or double quotes (depending on which are used to enclose it) and backslashes (\) are escaped with the backslash (\) character.

Example

```
{
  "books": [
    {
      "category": "reference",
      "author": "Nigel Rees",
      "title": "Sayings of the Century",
      "price": 8.95,
      "id": 1
    },
    {
      "category": "fiction",
      "author": "Evelyn Waugh",
      "title": "Sword of Honour",
      "price": 12.99,
      "id": 2
    },
    {
      "category": "fiction",
      "author": "Herman Melville",
      "title": "Moby Dick",
      "isbn": "0-553-21311-3",
      "price": 8.99,
      "id": 3
    },
    {
      "category": "fiction",
      "author": "J. R. R. Tolkien",
      "title": "The Lord of the Rings",
      "isbn": "0-395-19395-8",
      "price": 22.99,
      "id": 4
    }
  ],
  "services": {
    "delivery": {
      "servicegroup": 1000,
      "description": "Next day delivery in local town",
      "active": true,
      "price": 5
    },
    "bookbinding": {
      "servicegroup": 1001,
```

```
      "description": "Printing and assembling book in A5 format",
      "active": true,
      "price": 154.99
    },
    "restoration": {
      "servicegroup": 1002,
      "description": "Various restoration methods",
      "active": false,
      "methods": [
        {
          "description": "Chemical cleaning",
          "price": 46
        },
        {
          "description": "Pressing pages damaged by moisture",
          "price": 24.5
        },
        {
          "description": "Rebinding torn book",
          "price": 99.49
        }
      ]
    }
  },
  "filters": {
    "price": 10,
    "category": "fiction",
    "no filters": "no \"filters\""
  },
  "closed message": "Store is closed",
  "tags": [
    "a",
    "b",
    "c",
    "d",
    "e"
  ]
}
```


JSONPath	Type	Result
`$.filters.price`	definite	10
`$.filters.category`	definite	fiction
`$.filters['no filters']`	definite	no ”filters”
`$.filters`	definite	{ ”price”: 10, ”category”: ”fiction”, ”no filters”: ”no \”filters\”” }
`$.books[1].title`	definite	Sword of Honour
`$.books[-1].author`	definite	J. R. R. Tolkien
`$.books.length()`	definite	4
`$.tags[:]`	indefinite	[”a”, ”b”, ”c”, ”d”, ”e”]
`$.tags[2:]`	indefinite	[”c”, ”d”, ”e”]
`$.tags[:3]`	indefinite	[”a”, ”b”, ”c”]
`$.tags[1:4]`	indefinite	[”b”, ”c”, ”d”]
`$.tags[-2:]`	indefinite	[”d”, ”e”]
`$.tags[:-3]`	indefinite	[”a”, ”b”]
`$.tags[:-3].length()`	definite	2
`$.books[0, 2].title`	indefinite	[”Moby Dick”, ”Sayings of the Century”]
`$.books[1]['author', "title"]`	indefinite	[”Sword of Honour”, ”Evelyn Waugh”]
`$..id`	indefinite	[1, 2, 3, 4]

JSONPath	Type	Result
`$.services..price`	indefinite	[154.99, 5, 46, 24.5, 99.49]
`$.books[?(@.id == 4 - 0.4 * 5)].title`	indefinite	[”Sword of Honour”] Note: This query shows that arithmetical operations can be used in queries; it can be simplified to `$.books[?(@.id == 2)].title`
`$.books[?(@.id == 2 || @.id == 4)].title`	indefinite	[”Sword of Honour”, ”The Lord of the Rings”]
`$.books[?(!(@.id == 2))].title`	indefinite	[”Sayings of the Century”, ”Moby Dick”, ”The Lord of the Rings”]
`$.books[?(@.id != 2)].title`	indefinite	[”Sayings of the Century”, ”Moby Dick”, ”The Lord of the Rings”]
`$.books[?(@.title =~ " of ")].title`	indefinite	[”Sayings of the Century”, ”Sword of Honour”, ”The Lord of the Rings”]
`$.books[?(@.price > 12.99)].title`	indefinite	[”The Lord of the Rings”]
`$.books[?(@.author > "Herman Melville")].title`	indefinite	[”Sayings of the Century”, ”The Lord of the Rings”]
`$.books[?(@.price > $.filters.price)].title`	indefinite	[”Sword of Honour”, ”The Lord of the Rings”]
`$.books[?(@.category == $.filters.category)].title`	indefinite	[”Sword of Honour”,”Moby Dick”,”The Lord of the Rings”]
`$.books[?(@.category == "fiction" && @.price < 10)].title`	indefinite	[”Moby Dick”]

`$..[?(@.id)]`	indefinite	[
		{
		”price”: 8.95,
		”id”: 1,
		”category”: ”reference”,
		”author”: ”Nigel Rees”,
		”title”: ”Sayings of the Century”
		},
		{
		”price”: 12.99,
		”id”: 2,
		”category”: ”fiction”,
		”author”: ”Evelyn Waugh”,
		”title”: ”Sword of Honour”
		},
		{
		”price”: 8.99,
		”id”: 3,
		”category”: ”fiction”,
		”author”: ”Herman Melville”,
		”title”: ”Moby Dick”,
		”isbn”: ”0-553-21311-3”
		},
		{
		”price”: 22.99,
		”id”: 4,
		”category”: ”fiction”,
		”author”: ”J. R. R. Tolkien”,
		”title”: ”The Lord of the Rings”,
		”isbn”: ”0-395-19395-8”
		}
		]

`$.services..[?(@.price > 50)].description`	indefinite	[”Printing and assembling book in A5 format”, ”Rebinding torn book”]
`$..id.length()`	definite	4
`$.books[?(@.id == 2)].title.first()`	definite	Sword of Honour

JSONPath	Type	Result
`$..tags.first().length()`	definite	5 Note: `$..tags` is an indefinite path, so it returns an array of matched elements, i.e., `[["a", "b", "c", "d", "e" ]]`; `first()` returns the first element, i.e., `["a", "b", "c", "d", "e"]`; `length()` calculates the length of the element, i.e.,5.
`$.books[*].price.min()`	definite	8.95
`$..price.max()`	definite	154.99
`$.books[?(@.category == "fiction")].price.avg()`	definite	14.99
`$.books[?(@.category == $.filters.xyz)].title`	indefinite	Note: A query without match returns NULL for definite and indefinite paths.
`$.services[?(@.active=="true")].servicegroup`	indefinite	`[1001,1000]` Note: Text constants must be used in boolean value comparisons.
`$.services[?(@.active=="false")].servicegroup`	indefinite	`[1002]` Note: Text constants must be used in boolean value comparisons.
`$.services[?(@.servicegroup=="1002")].restore`	definite	`false`

Escaping special characters from LLD macro values in JSONPath

When low-level discovery macros are used in JSONPath preprocessing and their values are resolved, the following rules of escaping special characters are applied:

- only backslash (\) and double quote (”) characters are considered for escaping;
- if the resolved macro value contains these characters, each of them is escaped with a backslash;
- if they are already escaped with a backslash, it is not considered as escaping and both the backslash and the following special characters are escaped once again.

For example:

JSONPath	LLD macro value	After substitution
$.[?(@.value == ”{#MACRO}”)]	special ”value”	$.[?(@.value == ”special \”value\””)]
	c:\temp	$.[?(@.value == ”c:\\temp”)]
	a\\b	$.[?(@.value == ”a\\\\b”)]

When used in the expression, the macro that may have special characters should be enclosed in double quotes:

JSONPath	LLD macro value	After substitution	Result
$.[?(@.value == ”{#MACRO}”)]	special ”value”	$.[?(@.value == ”special \”value\””)]	OK
$.[?(@.value == {#MACRO})]		$.[?(@.value == special \”value\”)]	**Bad JSONPath expression**

When used in the path, the macro that may have special characters should be enclosed in square brackets **and** double quotes:

JSONPath	LLD macro value	After substitution	Result
$.[”{#MACRO}”].value	c:\temp	$.[”c:\\temp”].value	OK
$.{#MACRO}.value		$.c:\\temp.value	**Bad JSONPath expression**

JavaScript preprocessing

Overview

This section provides details of preprocessing by JavaScript.

> **Attention:**
> Do not use undeclared assignments in preprocessing JavaScript. Use `var` to declare local variables.

JavaScript preprocessing

JavaScript preprocessing is done by invoking JavaScript function with a single parameter 'value' and user-provided function body. The preprocessing step result is the value returned by this function, for example, to perform Fahrenheit to Celsius conversion, enter:

```
return (value - 32)  * 5 / 9
```

in JavaScript preprocessing parameters, which will be wrapped into a JavaScript function by server:

```
function (value)
{
   return (value - 32) * 5 / 9
}
```

The input parameter 'value' is always passed as a string. The return value is automatically coerced to string via toString() method (if it fails, then the error is returned as string value), with a few exceptions:

- returning undefined value will result in an error;
- returning null value will cause the input value to be discarded, much like 'Discard value' preprocessing on 'Custom on fail' action.

Errors can be returned by throwing values/objects (normally either strings or Error objects).

For example:

```
if (value == 0)
    throw "Zero input value"
return 1/value
```

Each script has a 10-second execution timeout (depending on the script, it might take longer for the timeout to trigger); exceeding it will return error. A 512-megabyte heap limit is enforced.

The JavaScript preprocessing step bytecode is cached and reused when the step is applied next time. Any changes to the item's preprocessing steps will cause the cached script to be reset and recompiled later.

Consecutive runtime failures (3 in a row) will cause the engine to be reinitialized to mitigate the possibility of one script breaking the execution environment for the next scripts (this action is logged with DebugLevel 4 and higher).

JavaScript preprocessing is implemented with Duktape JavaScript engine.

> **Note:**
> Built-in JavaScript object method definitions cannot be redefined in item preprocessing scripts.

See also: Additional JavaScript objects and global functions

Using macros in scripts

It is possible to use user macros (as well as LLD macros in the context of low-level discovery) in JavaScript code. If a script contains user macros, these macros are resolved by server/proxy before executing specific preprocessing steps. Note that when testing preprocessing steps in the frontend, macro values will not be pulled and need to be entered manually.

> **Note:**
> Context is ignored when a macro is replaced with its value. Macro value is inserted in the code as is, it is not possible to add additional escaping before placing the value in the JavaScript code. Please be advised that this can cause JavaScript errors in some cases.

In an example below, if received value exceeds a {$THRESHOLD} macro value, the threshold value (if present) will be returned instead:

```
var threshold = '{$THRESHOLD}';
return (!isNaN(threshold) && value > threshold) ? threshold : value;
```

Examples

The following examples illustrate how you can use JavaScript preprocessing.

Each example contains a brief description, a function body for JavaScript preprocessing parameters, and the preprocessing step result - value returned by the function.

Example 1: Convert a number (scientific notation to integer)

Convert the number "2.62128e+07" from scientific notation to an integer.

```
return (Number(value))
```

Value returned by the function: 26212800.

Example 2: Convert a number (binary to decimal)

Convert the binary number "11010010" to a decimal number.

```
return(parseInt(value,2))
```

Value returned by the function: 210.

Example 3: Round a number

Round the number "18.2345" to 2 digits.

```
return(Math.round(value* 100) / 100)
```

Value returned by the function: 18.23.

Example 4: Count letters in a string

Count the number of letters in the string "Zabbix".

```
return (value.length)
```

Value returned by the function: 6.

Example 5: Get time remaining

Get the remaining time (in seconds) until the expiration date of a certificate (Feb 12 12:33:56 2022 GMT).

```
var split = value.split(' '),
    MONTHS_LIST = ['Jan', 'Feb', 'Mar', 'Apr', 'May', 'Jun', 'Jul', 'Aug', 'Sep', 'Oct', 'Nov', 'Dec'],
    month_index = ('0' + (MONTHS_LIST.indexOf(split[0]) + 1)).slice(-2),
    ISOdate = split[3] + '-' + month_index + '-' + split[1] + 'T' + split[2],
    now = Date.now();

return parseInt((Date.parse(ISOdate) - now) / 1000);
```

Value returned by the function: 44380233.

Example 6: Remove JSON properties

Modify the JSON data structure by removing any properties with the key "data_size" or "index_size".

```
var obj=JSON.parse(value);

for (i = 0; i < Object.keys(obj).length; i++) {
    delete obj[i]["data_size"];
    delete obj[i]["index_size"];
}

return JSON.stringify(obj)
```

Value accepted by the function:

```
[
    {
        "table_name":"history",
        "data_size":"326.05",
        "index_size":"174.34"
    },
    {
        "table_name":"history_log",
        "data_size":"6.02",
        "index_size":"3.45"
```

```
    }
]
```

Value returned by the function:

```
[
    {
        "table_name":"history"
    },
    {
        "table_name":"history_log"
    }
]
```

Example 7: Convert Apache status to JSON

Convert the value received from a web.page.get Zabbix agent item (e.g., web.page.get[http://127.0.0.1:80/server-status?auto]) to a JSON object.

```
// Convert Apache status to JSON

// Split the value into substrings and put these substrings into an array
var lines = value.split('\n');

// Create an empty object "output"
var output = {};

// Create an object "workers" with predefined properties
var workers = {
    '_': 0, 'S': 0, 'R': 0, 'W': 0,
    'K': 0, 'D': 0, 'C': 0, 'L': 0,
    'G': 0, 'I': 0, '.': 0
};

// Add the substrings from the "lines" array to the "output" object as properties (key-value pairs)
for (var i = 0; i < lines.length; i++) {
    var line = lines[i].match(/([A-z0-9 ]+): (.*)/);

    if (line !== null) {
        output[line[1]] = isNaN(line[2]) ? line[2] : Number(line[2]);
    }
}

// Multiversion metrics
output.ServerUptimeSeconds = output.ServerUptimeSeconds || output.Uptime;
output.ServerVersion = output.ServerVersion || output.Server;

// Parse "Scoreboard" property to get the worker count
if (typeof output.Scoreboard === 'string') {
    for (var i = 0; i < output.Scoreboard.length; i++) {
        var char = output.Scoreboard[i];

        workers[char]++;
    }
}

// Add worker data to the "output" object
output.Workers = {
    waiting: workers['_'], starting: workers['S'], reading: workers['R'],
    sending: workers['W'], keepalive: workers['K'], dnslookup: workers['D'],
    closing: workers['C'], logging: workers['L'], finishing: workers['G'],
    cleanup: workers['I'], slot: workers['.']
};

// Return JSON string
```

```
return JSON.stringify(output);
```

Value accepted by the function:

```
HTTP/1.1 200 OK
Date: Mon, 27 Mar 2023 11:08:39 GMT
Server: Apache/2.4.52 (Ubuntu)
Vary: Accept-Encoding
Content-Encoding: gzip
Content-Length: 405
Content-Type: text/plain; charset=ISO-8859-1

127.0.0.1
ServerVersion: Apache/2.4.52 (Ubuntu)
ServerMPM: prefork
Server Built: 2023-03-08T17:32:01
CurrentTime: Monday, 27-Mar-2023 14:08:39 EEST
RestartTime: Monday, 27-Mar-2023 12:19:59 EEST
ParentServerConfigGeneration: 1
ParentServerMPMGeneration: 0
ServerUptimeSeconds: 6520
ServerUptime: 1 hour 48 minutes 40 seconds
Load1: 0.56
Load5: 0.33
Load15: 0.28
Total Accesses: 2476
Total kBytes: 8370
Total Duration: 52718
CPUUser: 8.16
CPUSystem: 3.44
CPUChildrenUser: 0
CPUChildrenSystem: 0
CPULoad: .177914
Uptime: 6520
ReqPerSec: .379755
BytesPerSec: 3461.58
BytesPerReq: 3461.58
DurationPerReq: 21.2916
BusyWorkers: 2
IdleWorkers: 6
Scoreboard: ____KW__.....................................................................................
```

Value returned by the function:

```
{
    "Date": "Mon, 27 Mar 2023 11:08:39 GMT",
    "Server": "Apache/2.4.52 (Ubuntu)",
    "Vary": "Accept-Encoding",
    "Encoding": "gzip",
    "Length": 405,
    "Type": "text/plain; charset=ISO-8859-1",
    "ServerVersion": "Apache/2.4.52 (Ubuntu)",
    "ServerMPM": "prefork",
    "Server Built": "2023-03-08T17:32:01",
    "CurrentTime": "Monday, 27-Mar-2023 14:08:39 EEST",
    "RestartTime": "Monday, 27-Mar-2023 12:19:59 EEST",
    "ParentServerConfigGeneration": 1,
    "ParentServerMPMGeneration": 0,
    "ServerUptimeSeconds": 6520,
    "ServerUptime": "1 hour 48 minutes 40 seconds",
    "Load1": 0.56,
    "Load5": 0.33,
    "Load15": 0.28,
    "Total Accesses": 2476,
```

```
    "Total kBytes": 8370,
    "Total Duration": 52718,
    "CPUUser": 8.16,
    "CPUSystem": 3.44,
    "CPUChildrenUser": 0,
    "CPUChildrenSystem": 0,
    "CPULoad": 0.177914,
    "Uptime": 6520,
    "ReqPerSec": 0.379755,
    "BytesPerSec": 1314.55,
    "BytesPerReq": 3461.58,
    "DurationPerReq": 21.2916,
    "BusyWorkers": 2,
    "IdleWorkers": 6,
    "Scoreboard": "____KW__.....
    "Workers": {
        "waiting": 6,
        "starting": 0,
        "reading": 0,
        "sending": 1,
        "keepalive": 1,
        "dnslookup": 0,
        "closing": 0,
        "logging": 0,
        "finishing": 0,
        "cleanup": 0,
        "slot": 142
    }
}
```

Additional JavaScript objects

Overview

This section describes Zabbix additions to the JavaScript language implemented with Duktape, and supported global JavaScript functions.

Attention:
Do not use undeclared assignments in preprocessing JavaScript. Use `var` to declare local variables.

Built-in objects

Zabbix

The Zabbix object provides interaction with the internal Zabbix functionality.

Method	Description
`log(loglevel, message)`	Writes <message> into Zabbix log using <loglevel> log level (see configuration file DebugLevel parameter).

Example:

```
Zabbix.log(3, "this is a log entry written with 'Warning' log level")
```

You may use the following aliases:

Alias	Alias to
console.log(object)	Zabbix.log(4, JSON.stringify(object))
console.warn(object)	Zabbix.log(3, JSON.stringify(object))
console.error(object)	Zabbix.log(2, JSON.stringify(object))

Attention:
The total size of all logged messages is limited to 8 MB per script execution.

Method	Description
`sleep(delay)`	Delay JavaScript execution by `delay` milliseconds.

Example (delay execution by 15 seconds):

```
Zabbix.sleep(15000)
```

HttpRequest

This object encapsulates a cURL handle, allowing simple HTTP requests to be made. Errors are thrown as exceptions.

Attention:
The initialization of multiple `HttpRequest` objects is limited to 10 per script execution.

Method	Description
`addHeader(value)`	Adds HTTP header field. This field is used for all following requests until cleared with the `clearHeader()` method. The total length of header fields that can be added to a single `HttpRequest` object is limited to 128 Kbytes (special characters and header names included).
`clearHeader()`	Clears HTTP header. If no header fields are set, `HttpRequest` will set Content-Type to application/json if the data being posted is JSON-formatted; text/plain otherwise.
`connect(url)`	Sends HTTP CONNECT request to the URL and returns the response.
`customRequest(method, url, data)`	Allows specifying any HTTP method in the first parameter. Sends the method request to the URL with optional *data* payload and returns the response.
`delete(url, data)`	Sends HTTP DELETE request to the URL with optional *data* payload and returns the response.
`getHeaders(<asArray>)`	Returns the object of received HTTP header fields. The `asArray` parameter may be set to "true" (e.g., `getHeaders(true)`), "false" or be undefined. If set to "true", the received HTTP header field values will be returned as arrays; this should be used to retrieve the field values of multiple same-name headers. If not set or set to "false", the received HTTP header field values will be returned as strings.
`get(url, data)`	Sends HTTP GET request to the URL with optional *data* payload and returns the response.
`head(url)`	Sends HTTP HEAD request to the URL and returns the response.
`options(url)`	Sends HTTP OPTIONS request to the URL and returns the response.
`patch(url, data)`	Sends HTTP PATCH request to the URL with optional *data* payload and returns the response.
`put(url, data)`	Sends HTTP PUT request to the URL with optional *data* payload and returns the response.
`post(url, data)`	Sends HTTP POST request to the URL with optional *data* payload and returns the response.
`getStatus()`	Returns the status code of the last HTTP request.
`setProxy(proxy)`	Sets HTTP proxy to "proxy" value. If this parameter is empty, then no proxy is used.

`setHttpAuth(bitmask, username, password)`	Sets enabled HTTP authentication methods (HTTPAUTH_BASIC, HTTPAUTH_DIGEST, HTTPAUTH_NEGOTIATE, HTTPAUTH_NTLM, HTTPAUTH_NONE) in the 'bitmask' parameter.
	The HTTPAUTH_NONE flag allows you to disable HTTP authentication.
	Examples:
	`request.setHttpAuth(HTTPAUTH_NTLM | HTTPAUTH_BASIC, username, password)`
	`request.setHttpAuth(HTTPAUTH_NONE)`

`trace(url, data)`	Sends HTTP TRACE request to the URL with optional *data* payload and returns the response.

Example:

```
try {
    Zabbix.log(4, 'jira webhook script value='+value);

    var result = {
        'tags': {
            'endpoint': 'jira'
        }
    },
    params = JSON.parse(value),
```

```
    req = new HttpRequest(),
    fields = {},
    resp;

    req.addHeader('Content-Type: application/json');
    req.addHeader('Authorization: Basic '+params.authentication);

    fields.summary = params.summary;
    fields.description = params.description;
    fields.project = {"key": params.project_key};
    fields.issuetype = {"id": params.issue_id};
    resp = req.post('https://jira.example.com/rest/api/2/issue/',
        JSON.stringify({"fields": fields})
    );

    if (req.getStatus() != 201) {
        throw 'Response code: '+req.getStatus();
    }

    resp = JSON.parse(resp);
    result.tags.issue_id = resp.id;
    result.tags.issue_key = resp.key;
} catch (error) {
    Zabbix.log(4, 'jira issue creation failed json : '+JSON.stringify({"fields": fields}));
    Zabbix.log(4, 'jira issue creation failed : '+error);

    result = {};
}

return JSON.stringify(result);
```

XML

The XML object allows the processing of XML data in the item and low-level discovery preprocessing and webhooks.

Attention:
In order to use XML object, server/proxy must be compiled with libxml2 support.

Method	Description
`XML.query(data, expression)`	Retrieves node content using XPath. Returns null if node is not found. **expression** - an XPath expression; **data** - XML data as a string.
`XML.toJson(data)`	Converts data in XML format to JSON.
`XML.fromJson(object)`	Converts data in JSON format to XML.

Example:

Input:

```
<menu>
    <food type = "breakfast">
        <name>Chocolate</name>
        <price>$5.95</price>
        <description></description>
        <calories>650</calories>
    </food>
</menu>
```

Output:

```
{
    "menu": {
        "food": {
```

```
            "@type": "breakfast",
            "name": "Chocolate",
            "price": "$5.95",
            "description": null,
            "calories": "650"
        }
    }
}
```

Serialization rules

XML to JSON conversion will be processed according to the following rules (for JSON to XML conversions reversed rules are applied):

1. XML attributes will be converted to keys that have their names prepended with '@'.

Example:

Input:

```
<xml foo="FOO">
  <bar>
    <baz>BAZ</baz>
  </bar>
</xml>
```

Output:

```
{
  "xml": {
    "@foo": "FOO",
    "bar": {
      "baz": "BAZ"
    }
  }
}
```

2. Self-closing elements (<foo/>) will be converted as having 'null' value.

Example:

Input:

```
<xml>
  <foo/>
</xml>
```

Output:

```
{
  "xml": {
    "foo": null
  }
}
```

3. Empty attributes (with "" value) will be converted as having empty string ('') value.

Example:

Input:

```
<xml>
  <foo bar="" />
</xml>
```

Output:

```
{
  "xml": {
    "foo": {
      "@bar": ""
    }
```

```
  }
}
```

4. Multiple child nodes with the same element name will be converted to a single key that has an array of values as its value.

Example:

Input:

```
<xml>
  <foo>BAR</foo>
  <foo>BAZ</foo>
  <foo>QUX</foo>
</xml>
```

Output:

```
{
  "xml": {
    "foo": ["BAR", "BAZ", "QUX"]
  }
}
```

5. If a text element has no attributes and no children, it will be converted as a string.

Example:

Input:

```
<xml>
    <foo>BAZ</foo>
</xml>
```

Output:

```
{
  "xml": {
    "foo": "BAZ"
   }
}
```

6. If a text element has no children but has attributes, text content will be converted to an element with the key '#text' and content as a value; attributes will be converted as described in the serialization rule 1.

Example:

Input:

```
<xml>
  <foo bar="BAR">
    BAZ
  </foo>
</xml>
```

Output:

```
{
  "xml": {
    "foo": {
      "@bar": "BAR",
      "#text": "BAZ"
    }
  }
}
```

Global JavaScript functions

Additional global JavaScript functions have been implemented with Duktape:

- `btoa(data)` - encodes the data to base64 string.
- `atob(base64_string)` - decodes base64 string as Uint8Array buffer.

```
try {
    b64 = btoa("test string");
    buffer = atob(b64);

    // Note that decoding logic depends on the data format of the buffer.
    decoded = String.fromCharCode.apply(this, [].slice.call(buffer));
}
catch (error) {
    return {'error.name' : error.name, 'error.message' : error.message};
}
```

- `md5(data)` - calculates the MD5 hash of the data.
- `sha256(data)` - calculates the SHA256 hash of the data.
- `hmac('<hash type>',key,data)` - returns a HMAC hash as a hex-formatted string; `md5` and `sha256` are supported as the `hash type`; the `key` and `data` parameters support binary data.

  Examples:

  - `hmac('md5',key,data)`
  - `hmac('sha256',key,data)`
- `sign(hash,key,data)` - returns calculated signature (RSA signature with SHA-256) as a string, where:
 hash - only `sha256` is allowed; otherwise, an error is thrown.
 key - the private key. It should correspond to PKCS#1 or PKCS#8 standard. The key can be provided in different forms:

  - with spaces instead of newlines
  - with escaped or non-escaped `\n` instead of newlines
  - without any newlines as a single-line string
  - as a JSON-formatted string

  The key also can be loaded from a user macro/secret macro/vault.

  data - the data that will be signed. It can be a string (binary data also supported) or buffer (Uint8Array/ArrayBuffer).

  Example:

  - `sign('sha256',key,data)`

  OpenSSL or GnuTLS is used to calculate the signatures. If Zabbix was built without any of these encryption libraries, an error will be thrown ('missing OpenSSL or GnuTLS library').

Browser item JavaScript objects

Overview

This section describes Zabbix additions to the JavaScript language implemented with Duktape for use in the Browser item script. These additions supplement the JavaScript objects described on the Additional JavaScript objects page.

Browser

The `Browser` object manages WebDriver sessions, initializing a session upon creation and terminating it upon destruction. A single script can support up to four `Browser` objects.

To construct a `Browser` object, use the `new Browser(options)` syntax. The `options` (*JSON object*) parameter specifies browser options, usually the WebDriver options method result (for example, `Browser.chromeOptions()`).

The following methods are supported with the `Browser` object.

Method	Description
`navigate(url)`	Navigate to the specified URL. Parameters: `url` - (string) URL to navigate to.
`getUrl()`	Return a string of the opened page URL.
`getPageSource()`	Return a string of the opened page source.

Method	Description
`findElement(strategy, selector)`	Return an `Element` object with one element in the opened page (or return `null` if no elements match `strategy` and `selector`). Parameters: `strategy` - (string, css selector/link text/partial link text/tag name/xpath) Location strategy; `selector` - (string) Element selector using the specified location strategy.
`findElements(strategy, target)`	Return an array of `Element` objects with multiple elements in the opened page (or return an empty array if no elements match location strategy and target). Parameters: `strategy` - (string, css selector/link text/partial link text/tag name/xpath) Location strategy; `target` - (string) Element selector using the specified location strategy.
`getCookies()`	Return an array of `Cookie` objects.
`addCookie(cookie)`	Set cookies. Parameters: `cookie` - (`Cookie` object) Cookie to set.
`getScreenshot()`	Return a base64 encoded string representing an image of the browser's viewport.
`setScreenSize(x,y)`	Set the browser's viewport size. Parameters: `x` - (string) Width in pixels; `y` - (string) Height in pixels.
`setScriptTimeout(timeout)`	Set script loading timeout. Parameters: `timeout` - (integer) Timeout value in milliseconds.
`setSessionTimeout(timeout)`	Set session (page load) timeout. Parameters: `timeout` - (integer) Timeout value in milliseconds.
`setElementWaitTimeout(timeout)`	Set element location strategy (implicit) timeout. Parameters: `timeout` - (integer) Timeout value in milliseconds.
`collectPerfEntries(mark)`	Collect performance entries to retrieve with the `getResult()` method. Parameters: `mark` - (string, optional) performance snapshot mark.
`getRawPerfEntries()`	Return an array of performance entry objects.
`getResult()`	Return a `Result` object with browser session statistics (error information, performance snapshots, etc.).
`getError()`	Return a `BrowserError` object with browser errors (or return `null` if there are no browser errors).
`setError(message)`	Set a custom error message to be included in the `Result` object. Parameters: `message` - (string) Error message.
`discardError()`	Discard the error to be returned in the `Result` object.
`getAlert()`	Return an `Alert` object with browser alerts (or return `null` if there are no browser alerts).
`chromeOptions()`	Return a `chromeOptions` object with predefined Chrome browser options.
`firefoxOptions()`	Return a `firefoxOptions` object with predefined Firefox browser options.
`safariOptions()`	Return a `safariOptions` object with predefined Safari browser options.
`edgeOptions()`	Return an `edgeOptions` object with predefined Edge browser options.

Method	Description
`switchFrame(target)`	Switch to the specified frame. Parameters: `target` - (browser element or integer, optional) Target frame. To select a frame by element, pass the element. To select a frame by index, pass the number. If left empty, will switch to the top-level browsing context.

All Browser methods can throw the following errors:

- `BrowserError` - derived from the `Error` object that is thrown if the `Browser` constructor fails; contains an additional `browser` property with a `Browser` object that threw this `BrowserError`.
- `WebdriverError` - derived from `BrowserError`; contains the same properties as the `BrowserError` object, which indicate if the error was generated in response to an error in the WebDriver response.

Element

The `Element` object is returned by the `Browser` object `findElement()`/`findElements()` methods and cannot be constructed directly.

The `Element` object represents an element in the web page and provides methods to interact with it.

The following methods are supported with the `Element` object.

Method	Description
`getAttribute(name)`	Return an attribute value string of the element attribute (or return `null` if the specified attribute was not found). Parameters: `name` - (string) Attribute name.
`getProperty(name)`	Return a property value string of the element property (or return `null` if the specified property was not found). Parameters: `name` - (string) Property name.
`getText()`	Return a text value string of the element text.
`click()`	Click an element.
`clear()`	Clear the content of an editable element.
`sendKeys(keys)`	Send keys. Parameters: `keys` - (string) Keys to send.

Cookie

The `Cookie` object is returned by the `Browser` object `getCookies()` method and passed to the `addCookie()` method.

While the `Cookie` object does not have any methods, it can contain the following properties:

Property	Type	Description
name	string	Name of the cookie.
value	string	Value of the cookie.
path	string	Path for which the cookie is valid. Defaults to `"/"` if omitted when adding a cookie.
domain	string	Domain to which the cookie is visible. Defaults to the URL domain of the session's current browsing context's active document if omitted when adding a cookie.
secure	boolean	Indicates whether the cookie is secure. Defaults to `false` if omitted when adding a cookie.
httpOnly	boolean	Indicates whether the cookie is HTTP-only. Defaults to `false` if omitted when adding a cookie.
expiry	integer	Expiration time of the cookie (in seconds since the Unix Epoch). Must not be set if omitted when adding a cookie.

Property	Type	Description
sameSite	string	The `sameSite` attribute of the cookie, which controls whether the cookie should be restricted to a first-party or same-site context. Can be set to either `"Lax"` or `"Strict"`. Defaults to `"None"` if omitted when adding a cookie.

Alert

The `Alert` object represents a web page alert, is returned by `Browser` object `getAlert()` method, and cannot be constructed directly.

The `Alert` object contains the `text` property with the alert text (or `null` if there are no alerts).

The following methods are supported with the `Alert` object.

Method	Description
`accept()`	Accept the alert.
`dismiss()`	Dismiss the alert.

Result

The `Result` object contains session statistics and is returned by the `Browser` object `getResult()` method.

Typically, the `Result` object is stringified and returned from the script, and then parsed into dependent item values through preprocessing.

While the `Result` object does not have any methods, it can contain the following properties.

Property		Type	Description
duration		string	Session duration from session creation to result retrieval.
error		object	Error information.
	http_status	integer	HTTP status returned by the WebDriver (or 0 if there are no WebDriver errors).
	error_code	string	Error returned by the WebDriver (or empty string if there are no WebDriver errors).
	message	string	WebDriver error message (or empty string if there are no WebDriver errors).
performance_data		object	Performance statistics.
	summary	object	Performance summary.
		navigation	object	Navigation summary.
		resource	object	Resource summary.
	details	array of objects	Performance statistics after each operation that could have resulted in navigation.
		mark	string	*(optional)* Performance snapshot mark specified with the `collectPerfEntries()` method.
		navigation	object	Navigation statistics.
		resource	object	Resource summary for this step.
		user	array of objects	Array of mark/measure type statistics.
	marks	array of objects	Marked performance snapshot indexes.
		name	string	Performance snapshot mark name.
		index	integer	Performance snapshot index in details array.

CSV to JSON preprocessing

Overview

In this preprocessing step it is possible to convert CSV file data into JSON format. It's supported in:

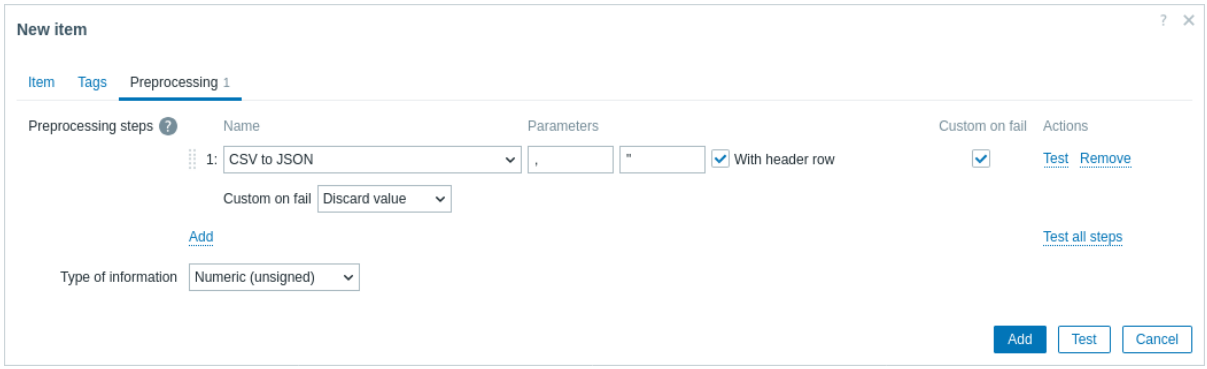
- items (item prototypes)
- low-level discovery rules

Configuration

To configure a CSV to JSON preprocessing step:

1. Go to the *Preprocessing* tab in item/discovery rule configuration.

2. Click *Add*.
3. Select the *CSV to JSON* option.



The first parameter allows you to set a custom delimiter. Note that if the first line of CSV input starts with "Sep=" and is followed by a single UTF-8 character then that character will be used as the delimiter in case the first parameter is not set. If the first parameter is not set and a delimiter is not retrieved from the "Sep=" line, then a comma is used as a separator.

The second optional parameter allows you to set a quotation symbol.

If the *With header row* checkbox is marked, the header line values will be interpreted as column names (see Header processing for more information).

If the *Custom on fail* checkbox is marked, the item will not become unsupported in case of a failed preprocessing step. Additionally, custom error handling options may be set: discard the value, set a specified value or set a specified error message.

Header processing

The CSV file header line can be processed in two different ways:

- If the *With header row* checkbox is marked - header line values are interpreted as column names. In this case the column names must be unique and the data row should not contain more columns than the header row.
- If the *With header row* checkbox is not marked - the header line is interpreted as data. Column names are generated automatically (1,2,3,4...).

CSV file example:

```
Nr,Item name,Key,Qty
1,active agent item,agent.hostname,33
"2","passive agent item","agent.version","44"
3,"active,passive agent items",agent.ping,55
```

Note:
A quotation character within a quoted field in the input must be escaped by preceding it with another quotation character.

Processing header line

JSON output when a header line is expected:

```
[
   {
      "Nr":"1",
      "Item name":"active agent item",
      "Key":"agent.hostname",
      "Qty":"33"
   },
   {
      "Nr":"2",
      "Item name":"passive agent item",
      "Key":"agent.version",
      "Qty":"44"
   },
   {
      "Nr":"3",
      "Item name":"active,passive agent items",
      "Key":"agent.ping",
      "Qty":"55"
```

```
    }
]
```

No header line processing

JSON output when a header line is not expected:

```
[
    {
        "1":"Nr",
        "2":"Item name",
        "3":"Key",
        "4":"Qty"
    },
    {
        "1":"1",
        "2":"active agent item",
        "3":"agent.hostname",
        "4":"33"
    },
    {
        "1":"2",
        "2":"passive agent item",
        "3":"agent.version",
        "4":"44"
    },
    {
        "1":"3",
        "2":"active,passive agent items",
        "3":"agent.ping",
        "4":"55"
    }
]
```

Item types

Overview

Item types cover various methods of acquiring data from your system. Each item type comes with its own set of supported item keys and required parameters.

The following item types are currently offered by Zabbix:

- Zabbix agent
  - Log file monitoring
- Simple check
  - VMware monitoring
- SNMP agent
- SNMP trap
- Zabbix internal
- Zabbix trapper
- External check
- Database monitor
- HTTP agent
  - Prometheus check
- IPMI agent
- SSH agent
- Telnet agent
- JMX agent
- Calculated item
- Dependent item
- Script item
- Browser item
- Telemetry query

Details for all item types are included in the subpages of this section. Even though item types offer a lot of options for data gathering, there are further options through user parameters or loadable modules.

Some checks are performed by Zabbix server alone (as agent-less monitoring) while others require Zabbix agent or even Zabbix Java gateway (with JMX monitoring).

> **Attention:**
> If a particular item type requires a particular interface (like an IPMI check needs an IPMI interface on the host) that interface must exist in the host definition.

Multiple interfaces can be set in the host definition: Zabbix agent, SNMP agent, JMX and IPMI. If an item can use more than one interface, it will search the available host interfaces (in the order: Agent→SNMP→JMX→IPMI) for the first appropriate one to be linked with.

All items that return text (character, log, text types of information) can return whitespace only as well (where applicable) setting the return value to an empty string (supported since 2.0).

Zabbix agent

Overview

This section provides details on the item keys that use communication with Zabbix agent for data gathering.

There are passive and active agent checks. When configuring an item, you can select the required type:

- *Zabbix agent* - for passive checks
- *Zabbix agent (active)* - for active checks

Note that all item keys supported by Zabbix agent are also supported by the new generation Zabbix agent 2. See the additional item keys that you can use with the agent 2 only.

Supported item keys

The item keys that you can use with Zabbix agent are listed below.

The item keys are listed without parameters and additional information. Click the item key to see the full details.

Item key	Description	Item group
kernel.maxfiles	The maximum number of opened files supported by OS.	Kernel
kernel.maxproc	The maximum number of processes supported by OS.	
kernel.openfiles	The number of currently open file descriptors.	
log	The monitoring of a log file.	Log monitoring
log.count	The count of matched lines in a monitored log file.	
logrt	The monitoring of a log file that is rotated.	
logrt.count	The count of matched lines in a monitored log file that is rotated.	
modbus.get	Reads Modbus data.	Modbus
net.dns	Checks the status of a DNS service.	Network
net.dns.perf	Checks the performance of a DNS service.	
net.dns.record	Performs a DNS query.	
net.if.collisions	The number of out-of-window collisions.	
net.if.discovery	The list of network interfaces.	
net.if.get	The list of network interfaces with detailed information.	
net.if.in	The incoming traffic statistics on a network interface.	
net.if.out	The outgoing traffic statistics on a network interface.	
net.if.total	The sum of incoming and outgoing traffic statistics on a network interface.	
net.tcp.listen	Checks if this TCP port is in LISTEN state.	
net.tcp.port	Checks if it is possible to make a TCP connection to the specified port.	
net.tcp.service	Checks if a service is running and accepting TCP connections.	
net.tcp.service.perf	Checks the performance of a TCP service.	
net.tcp.socket.count	Returns the number of TCP sockets that match parameters.	
net.udp.listen	Checks if this UDP port is in LISTEN state.	
net.udp.service	Checks if a service is running and responding to UDP requests.	
net.udp.service.perf	Checks the performance of a UDP service.	
net.udp.socket.count	Returns the number of UDP sockets that match parameters.	
proc.cpu.util	The process CPU utilization percentage.	Processes

Item key	Description	Item group
proc.get	The list of OS processes and their parameters.	
proc.mem	The memory used by the process in bytes.	
proc.num	The number of processes.	
sensor	Hardware sensor reading.	Sensors
system.boottime	The system boot time.	System
system.cpu.discovery	The list of detected CPUs/CPU cores.	
system.cpu.intr	The device interrupts.	
system.cpu.load	The CPU load.	
system.cpu.num	The number of CPUs.	
system.cpu.switches	The count of context switches.	
system.cpu.util	The CPU utilization percentage.	
system.hostname	The system host name.	
system.hw.chassis	The chassis information.	
system.hw.cpu	The CPU information.	
system.hw.devices	The listing of PCI or USB devices.	
system.hw.macaddr	The listing of MAC addresses.	
system.localtime	The system time.	
system.run	Run the specified command on the host.	
system.stat	The system statistics.	
system.sw.arch	The software architecture information.	
system.sw.os	The operating system information.	
system.sw.os.get	Detailed information about the operating system (version, type, distribution name, minor and major version, etc).	
system.sw.packages	The listing of installed packages.	
system.sw.packages.get	A detailed listing of installed packages.	
system.swap.in	The swap-in (from device into memory) statistics.	
system.swap.out	The swap-out (from memory onto device) statistics.	
system.swap.size	The swap space size in bytes or in percentage from total.	
system.uname	Identification of the system.	
system.uptime	The system uptime in seconds.	
system.users.num	The number of users logged in.	
vfs.dev.discovery	The list of block devices and their type.	Virtual file systems
vfs.dev.get	The list of block devices with detailed information.	
vfs.dev.read	The disk read statistics.	
vfs.dev.write	The disk write statistics.	
vfs.dir.count	The directory entry count.	
vfs.dir.get	The directory entry list.	
vfs.dir.size	The directory size.	
vfs.file.cksum	The file checksum, calculated by the UNIX cksum algorithm.	
vfs.file.contents	Retrieving the contents of a file.	
vfs.file.exists	Checks if the file exists.	
vfs.file.get	Returns information about a file.	
vfs.file.md5sum	The MD5 checksum of file.	
vfs.file.owner	Retrieves the owner of a file.	
vfs.file.permissions	Returns a 4-digit string containing the octal number with UNIX permissions.	
vfs.file.regexp	Retrieve a string in the file.	
vfs.file.regmatch	Find a string in the file.	
vfs.file.size	The file size.	
vfs.file.time	The file time information.	
vfs.fs.discovery	The list of mounted filesystems with their type and mount options.	
vfs.fs.get	The list of mounted filesystems with their type and mount options and, in the *full* mode, available disk space and inode statistics.	
vfs.fs.inode	The number or percentage of inodes.	
vfs.fs.size	The disk space in bytes or in percentage from total.	
vm.memory.size	The memory size in bytes or in percentage from total.	Virtual memory
web.page.get	Get the content of a web page.	Web monitoring
web.page.perf	The loading time of a full web page.	
web.page.regexp	Find a string on the web page.	

Item key	Description	Item group
agent.hostmetadata	The agent host metadata.	Zabbix
agent.hostname	The agent host name.	
agent.ping	The agent availability check.	
agent.variant	The variant of Zabbix agent (Zabbix agent or Zabbix agent 2).	
agent.version	The version of Zabbix agent.	
zabbix.stats	Returns a set of Zabbix server or proxy internal metrics. If `ip` and `port` are specified, metrics are retrieved from the remote instance; otherwise, from the local instance.	
zabbix.stats (queue)	Returns the number of monitored items in the queue which are delayed on Zabbix server or proxy. If `ip` and `port` are specified, metrics are retrieved from the remote instance; otherwise, from the local instance.	

Supported platforms

Except where specified differently in the item details, the agent items (and all parameters) are supported on:

- **Linux**
- **FreeBSD**
- **Solaris**
- **HP-UX**
- **AIX**
- **MacOS X**
- **OpenBSD**
- **NetBSD**

Many agent items are also supported on **Windows**. See the Windows agent item page for details.

Item key details

Parameters without angle brackets are mandatory. Parameters marked with angle brackets **< >** are optional.

kernel.maxfiles

 The maximum number of opened files supported by OS.
 Return value: *Integer*.
 Supported platforms: Linux, FreeBSD, MacOS X, OpenBSD, NetBSD.

kernel.maxproc

 The maximum number of processes supported by OS.
 Return value: *Integer*.
 Supported platforms: Linux 2.6 and later, FreeBSD, Solaris, MacOS X, OpenBSD, NetBSD.

kernel.openfiles

 The number of currently open file descriptors.
 Return value: *Integer*.
 Supported platforms: Linux (the item may work on other UNIX-like platforms).

log[file,<regexp>,<encoding>,<maxlines>,<mode>,<output>,<maxdelay>,<options>,<persistent dir>]

 The monitoring of a log file.
 Return value: *Log*.
 See supported platforms.

Parameters:

- **file** - the full path and name of a log file;
- **regexp** - a regular expression describing the required pattern;
- **encoding** - the code page identifier;
- **maxlines** - the maximum number of new lines per second the agent will send to Zabbix server or proxy. This parameter overrides the value of `MaxLinesPerSecond` in `zabbix_agentd.conf`;
- **mode** - possible values:
  - *all* (default);
  - *skip* - skip processing of older data (affects only newly created items);
- **output** - an optional output formatting template. The **\0** escape sequence is replaced with the matched part of text (from the first character where match begins until the character where match ends) while an **\N** (where N=1...9) escape sequence is replaced with Nth matched group (or an empty string if the N exceeds the number of captured groups);
- **maxdelay** - the maximum delay in seconds. Type: float. Values:
  - 0 - (default) never ignore log file lines;
  - > 0.0 - ignore older lines in order to get the most recent lines analyzed within "maxdelay" seconds. Read the maxdelay notes before using it!

- **options** - additional options:
mtime-noreread - non-unique records, reread only if the file size changes (ignore modification time change). (This parameter is deprecated since 5.0.2, because now mtime is ignored.)
- **persistent dir** (only in zabbix_agentd on Unix systems; not supported in Zabbix agent 2) - the absolute pathname of directory where to store persistent files. See also additional notes on persistent files.

Comments:

- The item must be configured as an active check.
- If the file is missing or permissions do not allow access, the item turns unsupported.
- If `output` is left empty - the whole line containing the matched text is returned. Note that all global regular expression types except *Matches regular expression* always return the whole matched line and the `output` parameter is ignored.
- Content extraction using the `output` parameter takes place on the agent.

Examples:

```
log[/var/log/syslog]
log[/var/log/syslog,error]
log[/home/zabbix/logs/logfile,,,100]
```

Example of using the `output` parameter for extracting a number from log record:

```
log[/app1/app.log,"task run [0-9.]+ sec, processed ([0-9]+) records, [0-9]+ errors",,,,\1] #this item will
```

Example of using the `output` parameter for rewriting a log record before sending to server:

```
log[/app1/app.log,"([0-9 :-]+) task run ([0-9.]+) sec, processed ([0-9]+) records, ([0-9]+) errors",,,,"\1
```

log.count[file,<regexp>,<encoding>,<maxproclines>,<mode>,<maxdelay>,<options>,<persistent dir>]

 The count of matched lines in a monitored log file.
 Return value: *Integer*.
 See supported platforms.

Parameters:

- **file** - the full path and name of log file;
- **regexp** - a regular expression describing the required pattern;
- **encoding** - the code page identifier;
- **maxproclines** - the maximum number of new lines per second the agent will analyze (cannot exceed 10000). The default value is 10*`MaxLinesPerSecond` in `zabbix_agentd.conf`.
- **mode** - possible values:
  - *all* (default);
  - *skip* - skip processing of older data (affects only newly created items).
- **maxdelay** - the maximum delay in seconds. Type: float. Values:
  - 0 - (default) never ignore log file lines;
  - > 0.0 - ignore older lines in order to get the most recent lines analyzed within "maxdelay" seconds. Read the maxdelay notes before using it!
- **options** - additional options:
mtime-noreread - non-unique records, reread only if the file size changes (ignore modification time change). (This parameter is deprecated since 5.0.2, because now mtime is ignored.)
- **persistent dir** (only in `zabbix_agentd` on Unix systems; not supported in Zabbix agent 2) - the absolute pathname of directory where to store persistent files. See also additional notes on persistent files.

Comments:

- The item must be configured as an active check.
- Matching lines are counted in the new lines since the last log check by the agent, and thus depend on the item update interval.
- If the file is missing or permissions do not allow access, the item turns unsupported.

logrt[file regexp,<regexp>,<encoding>,<maxlines>,<mode>,<output>,<maxdelay>,<options>,<persistent dir>]

 The monitoring of a log file that is rotated.
 Return value: *Log*.
 See supported platforms.

Parameters:

- **file regexp** - the absolute path to file, with the file name specified using a regular expression. Note that the regular expression applies only to the file name and does not need to match the entire name (e.g., `/path/to/agent` will match `zabbix_agentd.log`).
- **regexp** - a regular expression describing the required content pattern.
- **encoding** - the code page identifier.
- **maxlines** - the maximum number of new lines per second the agent will send to Zabbix server or proxy. This parameter overrides the value of `MaxLinesPerSecond` in `zabbix_agentd.conf`.
- **mode** - possible values:

    - *all* (default);
    - *skip* - skip processing of older data (affects only newly created items).
- **output** - an optional output formatting template. The **\0** escape sequence is replaced with the matched part of text (from the first character where match begins until the character where match ends) while an **\N** (where N=1...9) escape sequence is replaced with Nth matched group (or an empty string if the N exceeds the number of captured groups).
- **maxdelay** - the maximum delay in seconds. Type: float. Values:
    - 0 - (default) never ignore log file lines;
    - > 0.0 - ignore older lines in order to get the most recent lines analyzed within "maxdelay" seconds. Read the maxdelay notes before using it!
- **options** - the type of log file rotation and other options. Possible values:
    - *rotate* (default).
    - *copytruncate* - note that *copytruncate* cannot be used together with *maxdelay*. In this case *maxdelay* must be 0 or not specified; see copytruncate notes.
    - *mtime-reread* - non-unique records, reread if modification time or size changes (used by default if no `mtime-*` option is explicitly set).
    - *mtime-noreread* - non-unique records, reread only if the size changes (ignore modification time change).
- **persistent dir** (only in `zabbix_agentd` on Unix systems; not supported in Zabbix agent 2) - the absolute pathname of directory where to store persistent files. See also additional notes on persistent files.

Comments:

- The item must be configured as an active check.
- Log rotation is based on the last modification time of files.
- Note that `logrt` is designed to work with one currently active log file, with several other matching inactive files rotated. If, for example, a directory has many active log files, a separate `logrt` item should be created for each one. Otherwise, if one `logrt` item picks up too many files, it may lead to exhausted memory and a crash of monitoring.
- If `output` is left empty - the whole line containing the matched text is returned. Note that all global regular expression types except *Matches regular expression* always return the whole matched line and the `output` parameter is ignored.
- Content extraction using the `output` parameter takes place on the agent.
- In the `file regexp` parameter, the log directory path and log file regular expression must be separated by the correct directory separator:
    - On Windows, the separator must be a backslash (\); forward slashes may be tolerated at other positions, except the one that separates the log directory path and the log file regular expression (see examples below).
    - On other systems, it must be a forward slash (/).

Examples for Windows:

```
logrt["c:/dir1/dir2/dir3\filename.*\.log","pattern_to_match"] #this item will collect data from log files
logrt["//example.com/share/dir1/dir2/dir3\filename.*\.log","pattern_to_match"] #this item will collect dat
```

Examples for other systems:

```
logrt["/home/zabbix/logs/^logfile[0-9]{1,3}$",,,100] #this item will match a file like "logfile1" (will no
logrt["/home/user/^logfile_.*_[0-9]{1,3}$","pattern_to_match","UTF-8",100] #this item will collect data fr
```

Example of using the `output` parameter for extracting a number from log record:

```
logrt[/app1/^test.*log$,"task run [0-9.]+ sec, processed ([0-9]+) records, [0-9]+ errors",,,,\1] #this ite
```

Example of using the `output` parameter for rewriting a log record before sending to server:

```
logrt[/app1/^test.*log$,"([0-9 :-]+) task run ([0-9.]+) sec, processed ([0-9]+) records, ([0-9]+) errors",
```

logrt.count[file regexp,<regexp>,<encoding>,<maxproclines>,<mode>,<maxdelay>,<options>,<persistent dir>]

 The count of matched lines in a monitored log file that is rotated.
 Return value: *Integer*.
 See supported platforms.

Parameters:

- **file regexp** - the absolute path to file, with the file name specified using a regular expression. Note that the regular expression applies only to the file name and does not need to match the entire name (e.g., `/path/to/agent` will match `zabbix_agentd.log`).
- **regexp** - a regular expression describing the required pattern.
- **encoding** - the code page identifier.
- **maxproclines** - the maximum number of new lines per second the agent will analyze (cannot exceed 10000). The default value is 10*`MaxLinesPerSecond` in `zabbix_agentd.conf`.
- **mode** - possible values:
    - *all* (default);
    - *skip* - skip processing of older data (affects only newly created items).

- **maxdelay** - the maximum delay in seconds. Type: float. Values:
  - 0 - (default) never ignore log file lines;
  - > 0.0 - ignore older lines in order to get the most recent lines analyzed within "maxdelay" seconds. Read the `maxdelay` notes before using it!
- **options** - the type of log file rotation and other options. Possible values:
  - *rotate* (default).
  - *copytruncate* - note that `copytruncate` cannot be used together with `maxdelay`. In this case `maxdelay` must be 0 or not specified; see `copytruncate` notes.
  - *mtime-reread* - non-unique records, reread if modification time or size changes (default).
  - *mtime-noreread* - non-unique records, reread only if the size changes (ignore modification time change).
- **persistent dir** (only in `zabbix_agentd` on Unix systems; not supported in Zabbix agent 2) - the absolute pathname of directory where to store persistent files. See also additional notes on persistent files.

Comments:

- The item must be configured as an active check.
- Matching lines are counted in the new lines since the last log check by the agent, and thus depend on the item update interval.
- Log rotation is based on the last modification time of files.
- In the `file regexp` parameter, the log directory path and log file regular expression must be separated by the correct directory separator:
  - On Windows, the separator must be a backslash (\); forward slashes may be tolerated at other positions, except the one that separates the log directory path and the log file regular expression (see examples below).
  - On other systems, it must be a forward slash (/).

Examples for Windows:

```
logrt.count["c:/dir1/dir2/dir3\filename.*\.log","pattern_to_match"] #this item will count the number of ma
logrt.count["//example.com/share/dir1/dir2/dir3\filename.*\.log","pattern_to_match"] #this item will count
```

Examples for other systems:

```
logrt.count["/home/zabbix/logs/.*","^logfile[0-9]{1,3}$",,100] #this item will count the number of matches
logrt.count["/home/user/^logfile_.*_[0-9]{1,3}$","pattern_to_match","UTF-8",100] #this item will count the
```

modbus.get[endpoint,<slave id>,<function>,<address>,<count>,<type>,<endianness>,<offset>]

 Reads Modbus data.
 Return value: *JSON string*.
 Supported platforms: Linux, Windows.

Parameters:

- **endpoint** - the endpoint defined as `protocol://connection_string`;
- **slave id** - the slave ID;
- **function** - the Modbus function;
- **address** - the address of first registry, coil or input;
- **count** - the number of records to read;
- **type** - the type of data;
- **endianness** - the endianness configuration;
- **offset** - the number of registers, starting from `address`, the results of which will be discarded.

On Windows, this item follows UNC paths in the `endpoint` parameter, which can be a security risk.

See a detailed description of parameters.

net.dns[<ip>,name,<type>,<timeout>,<count>,<protocol>]

 Checks the status of a DNS service.
 Return values: - 0 - DNS resolution failed (DNS server did not respond or returned an error); - 1 - DNS resolution succeeded.
 See supported platforms.

Parameters:

- **ip** (ignored on Windows unless using Zabbix agent 2) - the IP address of DNS server (leave empty for the default DNS server);
- **name** - the DNS name to query;
- **type** - the record type to be queried (default is `SOA`);
- **timeout** (ignored on Windows unless using Zabbix agent 2) - the timeout for the request in seconds (default is 1 second);
- **count** (ignored on Windows unless using Zabbix agent 2) - the number of tries for the request (default is 2);
- **protocol** - the protocol used to perform DNS queries: *udp* (default) or *tcp*.

Comments:

- The possible values for `type` are: *ANY*, *A*, *NS*, *CNAME*, *MB*, *MG*, *MR*, *PTR*, *MD*, *MF*, *MX*, *SOA*, *NULL*, *WKS* (not supported for Zabbix agent on Windows, Zabbix agent 2 on all OS), *HINFO*, *MINFO*, *TXT*, *SRV*.
- For reverse DNS lookups (when `type` is set to *PTR*), you can provide the DNS name in both reversed and non-reversed format (see examples below). Note that when PTR record is requested, the DNS name is actually an IP address.
- Internationalized domain names are not supported, please use IDNA encoded names instead.

Examples:

```
net.dns[198.51.100.1,example.com,MX,2,1]

net.dns[,198.51.100.1,PTR]
net.dns[,1.100.51.198.in-addr.arpa,PTR]

net.dns[,2a00:1450:400f:800::200e,PTR]
net.dns[,e.0.0.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.8.0.f.0.0.4.0.5.4.1.0.0.a.2.ip6.arpa,PTR]
```

net.dns.perf[<ip>,name,<type>,<timeout>,<count>,<protocol>]

 Checks the performance of a DNS service.
 Return value: *Float* (0 - service is down; seconds - the number of seconds spent waiting for a response from the service).
 See supported platforms.

Parameters:

- **ip** (ignored on Windows unless using Zabbix agent 2) - the IP address of DNS server (leave empty for the default DNS server);
- **name** - the DNS name to query;
- **type** - the record type to be queried (default is `SOA`);
- **timeout** (ignored on Windows unless using Zabbix agent 2) - the timeout for the request in seconds (default is 1 second);
- **count** (ignored on Windows unless using Zabbix agent 2) - the number of tries for the request (default is 2);
- **protocol** - the protocol used to perform DNS queries: *udp* (default) or *tcp*.

Comments:

- The possible values for `type` are:
ANY, *A*, *NS*, *CNAME*, *MB*, *MG*, *MR*, *PTR*, *MD*, *MF*, *MX*, *SOA*, *NULL*, *WKS* (not supported for Zabbix agent on Windows, Zabbix agent 2 on all OS), *HINFO*, *MINFO*, *TXT*, *SRV*.
- For reverse DNS lookups (when `type` is set to *PTR*), you can provide the DNS name in both reversed and non-reversed format (see examples below). Note that when PTR record is requested, the DNS name is actually an IP address.
- Internationalized domain names are not supported, please use IDNA encoded names instead.
- The item returns a response time instead of 0 when the DNS server responds with an error code (for example, `NXDOMAIN` or `SERVFAIL`).

Examples:

```
net.dns.perf[198.51.100.1,example.com,MX,2,1]

net.dns.perf[,198.51.100.1,PTR]
net.dns.perf[,1.100.51.198.in-addr.arpa,PTR]

net.dns.perf[,2a00:1450:400f:800::200e,PTR]
net.dns.perf[,e.0.0.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.8.0.f.0.0.4.0.5.4.1.0.0.a.2.ip6.arpa,PTR]
```

net.dns.record[<ip>,name,<type>,<timeout>,<count>,<protocol>]

 Performs a DNS query.

Zabbix agent 2 also provides net.dns.get, which offers additional features such as more record types and greater control over DNS monitoring. Return value: a character string with the required type of information.
 See supported platforms.

Parameters:

- **ip** (ignored on Windows unless using Zabbix agent 2) - the IP address of DNS server (leave empty for the default DNS server);
- **name** - the DNS name to query;
- **type** - the record type to be queried (default is `SOA`);
- **timeout** (ignored on Windows unless using Zabbix agent 2) - the timeout for the request in seconds (default is 1 second);
- **count** (ignored on Windows unless using Zabbix agent 2) - the number of tries for the request (default is 2);
- **protocol** - the protocol used to perform DNS queries: *udp* (default) or *tcp*.

Comments:

- The possible values for `type` are:
ANY, *A*, *NS*, *CNAME*, *MB*, *MG*, *MR*, *PTR*, *MD*, *MF*, *MX*, *SOA*, *NULL*, *WKS* (not supported for Zabbix agent on Windows, Zabbix agent 2 on all OS), *HINFO*, *MINFO*, *TXT*, *SRV*.

- For reverse DNS lookups (when `type` is set to *PTR*), you can provide the DNS name in reversed or non-reversed format (see examples below). Note that when PTR record is requested, the DNS name is actually an IP address.
- Internationalized domain names are not supported, please use IDNA encoded names instead.

Examples:

```
net.dns.record[198.51.100.1,example.com,MX,2,1]

net.dns.record[,198.51.100.1,PTR]
net.dns.record[,1.100.51.198.in-addr.arpa,PTR]

net.dns.record[,2a00:1450:400f:800::200e,PTR]
net.dns.record[,e.0.0.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.8.0.f.0.0.4.0.5.4.1.0.0.a.2.ip6.arpa,PTR]
```

net.if.collisions[if]

 The number of out-of-window collisions.
 Return value: *Integer*.
 Supported platforms: Linux, FreeBSD, Solaris, AIX, MacOS X, OpenBSD, NetBSD. Root privileges are required on NetBSD.

Parameter:

- **if** - network interface name

net.if.discovery

 The list of network interfaces. Used for low-level discovery.
 Return value: *JSON string*.
 Supported platforms: Linux, FreeBSD, Solaris, HP-UX, AIX, OpenBSD, NetBSD, Windows.

net.if.get[<regexp>]

 The list of network interfaces with detailed information.
 Used for low-level discovery.
 This item is an extended version of the `net.if.discovery` Zabbix agent item.
 Return value: *JSON object*.
 Supported platforms: Linux.

Parameters:

- **regexp** - a regular expression matching network interface names.

Comments:

- Returns a JSON object with network interfaces whose names match the given regular expression (see more details).

net.if.in[if,<mode>]

 The incoming traffic statistics on a network interface.
 Return value: *Integer*.
 Supported platforms: Linux, FreeBSD, Solaris[5], HP-UX, AIX, MacOS X, OpenBSD, NetBSD, Windows. Root privileges are required on NetBSD.

Parameters:

- **if** - network interface name (Unix); network interface full description or IPv4 address; or, if in braces, network interface GUID (Windows);
- **mode** - possible values:
  - *bytes* - number of bytes (default);
  - *packets* - number of packets;
  - *errors* - number of errors;
  - *dropped* - number of dropped packets;
  - *overruns (fifo)* - the number of FIFO buffer errors;
  - *frame* - the number of packet framing errors;
  - *compressed* - the number of compressed packets received by the device driver;
  - *multicast* - the number of multicast frames received by the device driver.

Comments:

- You may use this key with the *Change per second* preprocessing step in order to get the bytes-per-second statistics.
- The *dropped* mode is supported only on Linux, FreeBSD, HP-UX, MacOS X, OpenBSD, NetBSD.
- The *overruns*, *frame*, *compressed*, *multicast* modes are supported only on Linux.
- On HP-UX this item does not provide details on loopback interfaces (e.g., lo0).

Examples:

```
net.if.in[eth0]
net.if.in[eth0,errors]
```

net.if.out[if,<mode>]

 The outgoing traffic statistics on a network interface.
 Return value: *Integer*.
 Supported platforms: Linux, FreeBSD, Solaris[5], HP-UX, AIX, MacOS X, OpenBSD, NetBSD, Windows. Root privileges are required on NetBSD.

Parameters:

- **if** - network interface name (Unix); network interface full description or IPv4 address; or, if in braces, network interface GUID (Windows);
- **mode** - possible values:
  - *bytes* - number of bytes (default);
  - *packets* - number of packets;
  - *errors* - number of errors;
  - *dropped* - number of dropped packets;
  - *overruns (fifo)* - the number of FIFO buffer errors;
  - *collisions (colls)* - the number of collisions detected on the interface;
  - *carrier* - the number of carrier losses detected by the device driver;
  - *compressed* - the number of compressed packets transmitted by the device driver.

Comments:

- You may use this key with the *Change per second* preprocessing step in order to get the bytes-per-second statistics.
- The *dropped* mode is supported only on Linux, HP-UX.
- The *overruns*, *collision*, *carrier*, *compressed* modes are supported only on Linux.
- On HP-UX this item does not provide details on loopback interfaces (e.g., lo0).

Examples:

```
net.if.out[eth0]
net.if.out[eth0,errors]
```

net.if.total[if,<mode>]

 The sum of incoming and outgoing traffic statistics on a network interface.
 Return value: *Integer*.
 Supported platforms: Linux, FreeBSD, Solaris[5], HP-UX, AIX, MacOS X, OpenBSD, NetBSD, Windows. Root privileges are required on NetBSD.

Parameters:

- **if** - network interface name (Unix); network interface full description or IPv4 address; or, if in braces, network interface GUID (Windows);
- **mode** - possible values:
  - *bytes* - number of bytes (default);
  - *packets* - number of packets;
  - *errors* - number of errors;
  - *dropped* - number of dropped packets;
  - *overruns (fifo)* - the number of FIFO buffer errors;
  - *collisions (colls)* - the number of collisions detected on the interface;
  - *compressed* - the number of compressed packets transmitted or received by the device driver.

Comments:

- You may use this key with the *Change per second* preprocessing step in order to get the bytes-per-second statistics.
- The *dropped* mode is supported only on Linux, HP-UX. Dropped packets are supported only if both `net.if.in` and `net.if.out` work for dropped packets on your platform.
- The *overruns*, *collision*, *compressed* modes are supported only on Linux.
- On HP-UX this item does not provide details on loopback interfaces (e.g., lo0).

Examples:

```
net.if.total[eth0]
net.if.total[eth0,errors]
```

net.tcp.listen[port]

 Checks if this TCP port is in LISTEN state.
 Return values: 0 - it is not in LISTEN state; 1 - it is in LISTEN state.
 Supported platforms: Linux, FreeBSD, Solaris, MacOS X, Windows.

Parameter:

- **port** - TCP port number

On Linux kernels 2.6.14 and above, the information about listening TCP sockets is obtained from the kernel's NETLINK interface, if possible. Otherwise, the information is retrieved from `/proc/net/tcp` and `/roc/net/tcp6` files.

Example:

```
net.tcp.listen[80]
```

net.tcp.port[<ip>,port]

 Checks if it is possible to make a TCP connection to the specified port.
 Return values: 0 - cannot connect; 1 - can connect.
 See supported platforms.

Parameters:

- **ip** - the IP address or DNS name (default is 127.0.0.1);
- **port** - the port number.

Comments:

- For simple TCP performance testing use net.tcp.service.perf[tcp,<ip>,<port>].
- These checks may result in additional messages in system daemon logfiles (SMTP and SSH sessions being logged usually).

Example:

```
net.tcp.port[,80] #this item can be used to test the web server availability running on port 80
```

net.tcp.service[service,<ip>,<port>]

 Checks if a service is running and accepting TCP connections.
 Return values: 0 - service is down; 1 - service is running.
 See supported platforms.

Parameters:

- **service** - *ssh*, *ldap*, *smtp*, *ftp*, *http*, *pop*, *nntp*, *imap*, *tcp*, *https*, or *telnet* (see details);
- **ip** - the IP address or DNS name (default is 127.0.0.1);
- **port** - the port number (by default the standard service port number is used).

Comments:

- These checks may result in additional messages in system daemon logfiles (SMTP and SSH sessions being logged usually).
- Checking of encrypted protocols (like IMAP on port 993 or POP on port 995) is currently not supported. As a workaround, please use net.tcp.port[] for checks like these.
- Checking of LDAP and HTTPS on Windows is only supported by Zabbix agent 2.
- The telnet check looks for a login prompt (':' at the end).

Example:

```
net.tcp.service[ftp,,45] #this item can be used to test the availability of FTP server on TCP port 45
```

net.tcp.service.perf[service,<ip>,<port>]

 Checks the performance of a TCP service.
 Return values: *Float* (0 - service is down; seconds - the number of seconds spent waiting for a response from the service).
 See supported platforms.

Parameters:

- **service** - *ssh*, *ldap*, *smtp*, *ftp*, *http*, *pop*, *nntp*, *imap*, *tcp*, *https*, or *telnet* (see details);
- **ip** - the IP address or DNS name (default is 127.0.0.1);
- **port** - the port number (by default the standard service port number is used).

Comments:

- Checking of encrypted protocols (like IMAP on port 993 or POP on port 995) is currently not supported. As a workaround, please use net.tcp.service.perf[tcp,<ip>,<port>] for checks like these.
- The telnet check looks for a login prompt (':' at the end).

Example:

```
net.tcp.service.perf[ssh] #this item can be used to test the speed of initial response from the SSH server
```

net.tcp.socket.count[<laddr>,<lport>,<raddr>,<rport>,<state>]

 Returns the number of TCP sockets that match parameters.
 Return value: *Integer*.
 Supported platforms: Linux, Windows.

Parameters:

- **laddr** - the local IPv4/6 address or CIDR subnet;
- **lport** - the local port number or service name;
- **raddr** - the remote IPv4/6 address or CIDR subnet;

- **rport** - the remote port number or service name;
- **state** - the connection state (*established*, *syn_sent*, *syn_recv*, *fin_wait1*, *fin_wait2*, *time_wait*, *close*, *close_wait*, *last_ack*, *listen*, *closing*).

Example:

```
net.tcp.socket.count[,80,,,established] #the number of connections to local TCP port 80 in the established
```

net.udp.listen[port]

 Checks if this UDP port is in LISTEN state.
 Return values: 0 - it is not in LISTEN state; 1 - it is in LISTEN state.
 Supported platforms: Linux, FreeBSD, Solaris, MacOS X.

Parameter:

- **port** - UDP port number

Example:

```
net.udp.listen[68]
```

net.udp.service[service,<ip>,<port>]

 Checks if a service is running and responding to UDP requests.
 Return values: 0 - service is down; 1 - service is running.
 See supported platforms.

Parameters:

- **service** - *ntp* (see details);
- **ip** - the IP address or DNS name (default is `127.0.0.1`);
- **port** - the port number (by default the standard service port number is used).

Example:

```
net.udp.service[ntp,,45] #this item can be used to test the availability of NTP service on UDP port 45
```

net.udp.service.perf[service,<ip>,<port>]

 Checks the performance of a UDP service.
 Return values: *Float* (0 - service is down; seconds - the number of seconds spent waiting for a response from the service).
 See supported platforms.

Parameters:

- **service** - *ntp* (see details);
- **ip** - the IP address or DNS name (default is `127.0.0.1`);
- **port** - the port number (by default the standard service port number is used).

Example:

```
net.udp.service.perf[ntp] #this item can be used to test response time from NTP service
```

net.udp.socket.count[<laddr>,<lport>,<raddr>,<rport>,<state>]

 Returns the number of UDP sockets that match parameters.
 Return value: *Integer*.
 Supported platforms: Linux, Windows.

Parameters:

- **laddr** - the local IPv4/6 address or CIDR subnet;
- **lport** - the local port number or service name;
- **raddr** - the remote IPv4/6 address or CIDR subnet;
- **rport** - the remote port number or service name;
- **state** - the connection state (*established*, *unconn*).

Example:

```
net.udp.socket.count[,,,,established] #returns the number of UDP sockets in the connected state
```

proc.cpu.util[<name>,<user>,<type>,<cmdline>,<mode>,<zone>]

 The process CPU utilization percentage.
 Return value: *Float*.
 Supported platforms: Linux, Solaris[6].

Parameters:

- **name** - the process name (default is *all processes*);
- **user** - the user name (default is *all users*);
- **type** - the CPU utilization type: *total* (default), *user*, or *system*;

- **cmdline** - filter by command line (it is a regular expression);
- **mode** - the data gathering mode: *avg1* (default), *avg5*, or *avg15*;
- **zone** - the target zone: *current* (default) or *all*. This parameter is supported on Solaris only.

Comments:

- The returned value is based on a single CPU core utilization percentage. For example, the CPU utilization of a process fully using two cores is 200%.
- The process CPU utilization data is gathered by a collector which supports the maximum of 1024 unique (by name, user and command line) queries. Queries not accessed during the last 24 hours are removed from the collector.
- When setting the `zone` parameter to *current* (or default) in case the agent has been compiled on a Solaris without zone support, but running on a newer Solaris where zones are supported, then the agent will return NOTSUPPORTED (the agent cannot limit results to only the current zone). However, *all* is supported in this case.

Examples:

```
proc.cpu.util[,root] #CPU utilization of all processes running under the "root" user
proc.cpu.util[zabbix_server,zabbix] #CPU utilization of all zabbix_server processes running under the zabb
```

proc.get[<name>,<user>,<cmdline>,<mode>]

 The list of OS processes and their parameters. Can be used for low-level discovery.
 Return value: *JSON string*.
 Supported platforms: Linux, FreeBSD, Windows, OpenBSD, NetBSD.

Parameters:

- **name** - the process name (default *all processes*).
- **user** - the user name (default *all users*).
- **cmdline** - filter by command line (it is a regular expression). This parameter is not supported for Windows; on other platforms it is not supported if mode is set to 'summary'.
- **mode** - possible values:
  - *process* (default);
  - *thread* (not supported for NetBSD);
  - *summary*. See a list of process parameters returned for each mode and OS.

Comments:

- If a value cannot be retrieved, for example, because of an error (process already died, lack of permissions, system call failure), -1 will be returned.
- See notes on selecting processes with `name` and `cmdline` parameters (Linux-specific).

Examples:

```
proc.get[zabbix_server,zabbix,,process] #list of all zabbix_server processes running under the zabbix user
proc.get[java,,,thread] #list of all Java processes, returns one entry per thread
proc.get[,zabbix,,summary] #combined data for processes of each type running under the zabbix user, return
```

proc.mem[<name>,<user>,<mode>,<cmdline>,<memtype>]

 The memory used by the process in bytes.
 Return value: - *Integer* - with `mode` as *max*, *min*, *sum*; - *Float* - with `mode` as *avg*
 Supported platforms: Linux, FreeBSD, Solaris, AIX, OpenBSD, NetBSD.

Parameters:

- **name** - the process name (default is *all processes*);
- **user** - the user name (default is *all users*);
- **mode** - possible values: *avg*, *max*, *min*, or *sum* (default);
- **cmdline** - filter by command line (it is a regular expression);
- **memtype** - the type of memory used by process

Comments:

- The `memtype` parameter is supported only on Linux, FreeBSD, Solaris[6], AIX.
- When several processes use shared memory, the sum of memory used by processes may result in large, unrealistic values.

See notes on selecting processes with `name` and `cmdline` parameters (Linux-specific).

When this item is invoked from the command line and contains a command line parameter (e.g., using the agent test mode: `zabbix_agentd -t proc.mem[,,,apache2]`), one extra process will be counted, as the agent will count itself.

Examples:

```
proc.mem[,root] #the memory used by all processes running under the "root" user
proc.mem[zabbix_server,zabbix] #the memory used by all zabbix_server processes running under the zabbix us
proc.mem[,oracle,max,oracleZABBIX] #the memory used by the most memory-hungry process running under Oracle
```

proc.num[<name>,<user>,<state>,<cmdline>,<zone>]

 The number of processes.
 Return value: *Integer*.
 Supported platforms: Linux, FreeBSD, Solaris[6], HP-UX, AIX, OpenBSD, NetBSD, Windows.

Parameters:

- **name** - the process name (default is *all processes*).
- **user** - the user name (default is *all users*).
- **state** - possible values:
  - *all* (default);
  - *disk* - uninterruptible sleep;
  - *run* - running;
  - *sleep* - interruptible sleep;
  - *trace* - stopped;
  - *zomb* - zombie.
- **cmdline** - filter by command line (it is a regular expression).
- **zone** - the target zone: *current* (default), or *all*. This parameter is supported on Solaris only.

Comments:

- The *disk* and *trace* state parameters are supported only on Linux, FreeBSD, OpenBSD, NetBSD.
- On Windows, only the `name` and `user` parameters are supported.
- When this item is invoked from the command line and contains a command line parameter (e.g., using the agent test mode: `zabbix_agentd -t proc.num[,,,apache2]`), one extra process will be counted, as the agent will count itself.
- When setting the `zone` parameter to `current` (or default) in case the agent has been compiled on a Solaris without zone support, but running on a newer Solaris where zones are supported, then the agent will return NOTSUPPORTED (the agent cannot limit results to only the current zone). However, *all* is supported in this case.
- See notes on selecting processes with `name` and `cmdline` parameters (Linux-specific).

Examples:

```
proc.num[,mysql] #the number of processes running under the mysql user
proc.num[apache2,www-data] #the number of apache2 processes running under the www-data user
proc.num[,oracle,sleep,oracleZABBIX] #the number of processes in sleep state running under Oracle having c
```

sensor[device,sensor,<mode>]

 Hardware sensor reading.
 Return value: *Float*.
 Supported platforms: Linux, OpenBSD.

Parameters:

- **device** - the device name;
- **sensor** - the sensor name;
- **mode** - possible values: *avg*, *max*, or *min*.

Comments:

- Reads `/proc/sys/dev/sensors` on Linux 2.4.
- Reads `/sys/class/hwmon` on Linux 2.6+. See a more detailed description of sensor item on Linux.
- Reads the `hw.sensors` MIB on OpenBSD.
- On Linux 2.4, if you specify mode then device and sensor are treated as regular expressions.
- On Linux 2.6+ and OpenBSD, or if mode is omitted, device and sensor are treated as verbatim.

Examples:

```
sensor[w83781d-i2c-0-2d,temp1]
sensor[cpu0,temp0] #the temperature of one CPU
sensor["cpu[0-2]$",temp,avg] #the average temperature of the first three CPUs
```

system.boottime

 The system boot time.
 Return value: *Integer (Unix timestamp)*.
 Supported platforms: Linux, FreeBSD, Solaris, MacOS X, OpenBSD, NetBSD.

system.cpu.discovery

 The list of detected CPUs/CPU cores. Used for low-level discovery.
 Return value: *JSON string*.
 See supported platforms.

system.cpu.intr

 The device interrupts.
 Return value: *Integer*.
 Supported platforms: Linux, FreeBSD, Solaris, AIX, OpenBSD, NetBSD.

system.cpu.load[<cpu>,<mode>]

 The CPU load.
 Return value: *Float*.
 See supported platforms.

Parameters:

- **cpu** - possible values: *all* (default) or *percpu* (the total load divided by online CPU count);
- **mode** - possible values: *avg1* (one-minute average, default), *avg5*, or *avg15*.

Example:

```
system.cpu.load[,avg5]
```

system.cpu.num[<type>]

 The number of CPUs.
 Return value: *Integer*.
 Supported platforms: Linux, FreeBSD, Solaris, HP-UX, AIX, MacOS X, OpenBSD, NetBSD, Windows.

Parameter:

- **type** - possible values: *online* (default) or *max*

The *max* type parameter is supported only on Linux, FreeBSD, Solaris, MacOS X.

Example:

```
system.cpu.num
```

system.cpu.switches

 The count of context switches.
 Return value: *Integer*.
 Supported platforms: Linux, FreeBSD, Solaris, AIX, OpenBSD, NetBSD.

system.cpu.util[<cpu>,<type>,<mode>,<logical or physical>]

 The CPU utilization percentage.
 Return value: *Float*.
 Supported platforms: Linux, FreeBSD, Solaris, HP-UX, AIX, OpenBSD, NetBSD, Windows.

Parameters:

- **cpu** - *<CPU number>* or *all* (default);
- **type** - possible values: *user* (default), *idle*, *nice*, *system*, *iowait*, *interrupt*, *softirq*, *steal*, *spin* (on OpenBSD), *guest* (on Linux kernels 2.6.24 and above), or *guest_nice* (on Linux kernels 2.6.33 and above);
- **mode** - possible values: *avg1* (one-minute average, default), *avg5*, or *avg15*;
- **logical or physical** - possible values: *logical* (default) or *physical*. This parameter is supported on AIX only.

Comments:

- The *nice* type parameter is supported only on Linux, FreeBSD, HP-UX, OpenBSD, NetBSD.
- The *iowait* type parameter is supported only on Linux 2.6 and later, Solaris, AIX.
- The *interrupt* type parameter is supported only on Linux 2.6 and later, FreeBSD, OpenBSD.
- The *softirq*, *steal*, *guest*, *guest_nice* type parameters are supported only on Linux 2.6 and later.
- The *avg5* and *avg15* mode parameters are supported on Linux, FreeBSD, Solaris, HP-UX, AIX, OpenBSD, NetBSD.

Example:

```
system.cpu.util[0,user,avg5]
```

system.hostname[<type>,<transform>]

 The system host name.
 Return value: *String*.
 See supported platforms.

Parameters:

- **type** - possible values:
  - *netbios* (default on Windows);
  - *host* (default on Linux);
  - *shorthost* (returns part of the hostname before the first dot, a full string for names without dots);
  - *fqdn* (returns Fully Qualified Domain Name).
- **transform** - possible values:

- *none* (default);
- *lower* (convert to lowercase).

The value is acquired by taking `nodename` from the uname() system API output.

Examples of returned values:

```
system.hostname → linux-w7x1
system.hostname → example.com
system.hostname[shorthost] → example
system.hostname → WIN-SERV2008-I6
system.hostname[host] → Win-Serv2008-I6LonG
system.hostname[host,lower] → win-serv2008-i6long
system.hostname[fqdn,lower] → blog.zabbix.com
```

system.hw.chassis[<info>]

 The chassis information.
 Return value: *String*.
 Supported platforms: Linux.

Parameter:

- **info** - possible values: *full* (default), *model*, *serial*, *type*, or *vendor*

Comments:

- This item key depends on the availability of the SMBIOS table.
- It will try to read the DMI table from sysfs, if sysfs access fails then try reading directly from memory.
- **Root permissions** are required because the value is acquired by reading from sysfs or memory.

Example:

```
system.hw.chassis[full] → Hewlett-Packard HP Pro 3010 Small Form Factor PC CZXXXXXXXX Desktop
```

system.hw.cpu[<cpu>,<info>]

 The CPU information.
 Return value: *String* or *Integer*.
 Supported platforms: Linux.

Parameters:

- **cpu** - <*CPU number*> or *all* (default);
- **info** - possible values: *full* (default), *curfreq*, *maxfreq*, *model* or *vendor*.

Comments:

- Gathers info from `/proc/cpuinfo` and `/sys/devices/system/cpu/[cpunum]/cpufreq/cpuinfo_max_freq`.
- If a CPU number and *curfreq* or *maxfreq* is specified, a numeric value is returned (Hz).

Example:

```
system.hw.cpu[0,vendor] → AuthenticAMD
```

system.hw.devices[<type>]

 The listing of PCI or USB devices.
 Return value: *Text*.
 Supported platforms: Linux.

Parameter:

- **type** - *pci* (default) or *usb*

Returns the output of either the lspci or lsusb utility (executed without any parameters).

Example:

```
system.hw.devices → 00:00.0 Host bridge: Advanced Micro Devices [AMD] RS780 Host Bridge
```

system.hw.macaddr[<interface>,<format>]

 The listing of MAC addresses.
 Return value: *String*.
 Supported platforms: Linux.

Parameters:

- **interface** - *all* (default) or a regular expression;
- **format** - *full* (default) or *short*.

Comments:

- Lists MAC addresses of the interfaces whose name matches the given `interface` regular expression (*all* lists for all interfaces).

- If `format` is specified as *short*, interface names and identical MAC addresses are not listed.

Example:

```
system.hw.macaddr["eth0$",full] → [eth0] 00:11:22:33:44:55
```

system.localtime[<type>]

 The system time.
 Return value: *Integer* - with `type` as *utc*; *String* - with `type` as *local*.
 See supported platforms.

Parameters:

- **type** - possible values:
  - *utc* - (default) the time since the Epoch (00:00:00 UTC, January 1, 1970), measured in seconds;
  - *local* - the time in the 'yyyy-mm-dd,hh:mm:ss.nnn,+hh:mm' format.

Must be used as a passive check for Zabbix agent; on Zabbix agent 2 it may be configured as an active check.

Example:

```
system.localtime[local] #create an item using this key and then use it to display the host time in the *Cl
```

system.run[command,<mode>]

 Run the specified command on the host.
 Return value: *Text* result of the command or 1 - with `mode` as *nowait* (regardless of the command result).
 See supported platforms.

Parameters:

- **command** - command for execution;
- **mode** - possible values:
  - *wait* - wait end of execution (default);
  - *nowait* - do not wait.

Comments:

- This item is disabled by default. Learn how to enable them.
- The return value of the item is a standard output together with a standard error produced by the command. Exit code checking is not performed.
- To be processed correctly, the return value of the command must be of `text` data type. An empty result is also allowed.
- The return value is limited to 16MB (including trailing whitespace that is truncated); database limits also apply.
- See also: Command execution.

Example:

```
system.run[ls -l /] #return a detailed file list of the root directory
```

system.stat[resource,<type>]

 The system statistics.
 Return value: *Integer* or *float*.
 Supported platforms: AIX.

Parameters:

- **ent** - the number of processor units this partition is entitled to receive (float).
- **kthr,<type>** - information about kernel thread states:
  - *r* - average number of runnable kernel threads (float);
  - *b* - average number of kernel threads placed in the Virtual Memory Manager wait queue (float).
- **memory,<type>** - information about the usage of virtual and real memory:
  - *avm* - active virtual pages (integer);
  - *fre* - size of the free list (integer).
- **page,<type>** - information about page faults and paging activity:
  - *fi* - file page-ins per second (float);
  - *fo* - file page-outs per second (float);
  - *pi* - pages paged in from paging space (float);
  - *po* - pages paged out to paging space (float);
  - *fr* - pages freed (page replacement) (float);
  - *sr* - pages scanned by page-replacement algorithm (float).
- **faults,<type>** - trap and interrupt rate:
  - *in* - device interrupts (float);
  - *sy* - system calls (float);
  - *cs* - kernel thread context switches (float).
- **cpu,<type>** - breakdown of percentage usage of processor time:
  - *us* - user time (float);

    - *sy* - system time (float);
    - *id* - idle time (float);
    - *wa* - idle time during which the system had outstanding disk/NFS I/O request(s) (float);
    - *pc* - number of physical processors consumed (float);
    - *ec* - the percentage of entitled capacity consumed (float);
    - *lbusy* - indicates the percentage of logical processor(s) utilization that occurred while executing at the user and system level (float);
    - *app* - indicates the available physical processors in the shared pool (float).
- **disk,<type>** - disk statistics:
    - *bps* - indicates the amount of data transferred (read or written) to the drive in bytes per second (integer);
    - *tps* - indicates the number of transfers per second that were issued to the physical disk/tape (float).

Comments:

- Take note of the following limitations in these items:
    - `system.stat[cpu,app]` - supported only on AIX LPAR of type "Shared";
    - `system.stat[cpu,ec]` - supported on AIX LPAR of type "Shared" and "Dedicated" ("Dedicated" always returns 100 (percent));
    - `system.stat[cpu,lbusy]` - supported only on AIX LPAR of type "Shared";
    - `system.stat[cpu,pc]` - supported on AIX LPAR of type "Shared" and "Dedicated";
    - `system.stat[ent]` - supported on AIX LPAR of type "Shared" and "Dedicated".

system.sw.arch

 The software architecture information.
 Return value: *String*.
 See supported platforms.

The info is acquired from the `uname()` function.

Example:

```
system.sw.arch → i686
```

system.sw.os[<info>]

 The operating system information.
 Return value: *String*.
 Supported platforms: Linux, Windows.

Parameter:

- **info** - possible values: *full* (default), *short*, or *name*

The info is acquired from (note that not all files and options are present in all distributions):

- `/proc/version` (*full*) on Linux;
- `/proc/version_signature` (*short*) on Linux;
- the PRETTY_NAME parameter from `/etc/os-release` on Linux-systems supporting it or `/etc/issue.net` (*name*);
- the `HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion` registry key on Windows.

Examples:

```
system.sw.os[short] → Ubuntu 2.6.35-28.50-generic 2.6.35.11
system.sw.os[full] → [s|Windows 10 Enterprise 22621.1.asd64fre.ni_release.220506-1250 Build 22621.963]
```

system.sw.os.get

 Detailed information about the operating system (version, type, distribution name, minor and major version, etc).
 Return value: *JSON string*.
 Supported platforms: Linux, Windows.

system.sw.packages[<regexp>,<manager>,<format>]

 The listing of installed packages.
 Return value: *Text*.
 Supported platforms: Linux.

Parameters:

- **regexp** - *all* (default) or a regular expression;
- **manager** - *all* (default) or a package manager;
- **format** - *full* (default) or *short*.

Comments:

- Lists (alphabetically) installed packages whose name matches the given regular expression (*all* lists them all).
- Supported package managers (executed command):
    - pkg (`dpkg --get-selections`);
    - pkgtool (`ls /var/log/packages`);
    - rpm (`rpm -qa`);

    - pacman (`pacman -Q`);
    - portage.
- If `format` is specified as *full*, packages are grouped by package managers (each manager on a separate line beginning with its name in square brackets).
- If `format` is specified as *short*, packages are not grouped and are listed on a single line.

Example:

```
system.sw.packages[mini,dpkg,short] → python-minimal, python2.6-minimal, ubuntu-minimal
```

system.sw.packages.get[<regexp>,<manager>]

 A detailed listing of installed packages.
 Return value: *JSON string*.
 Supported platforms: Linux.

Parameters:

- **regexp** - *all* (default) or a regular expression;
- **manager** - *all* (default) or a package manager (possible values: *rpm*, *dpkg*, *pkgtool*, *pacman*, or *portage*).

Comments:

- Returns unformatted JSON with the installed packages whose name matches the given regular expression.
- The output is an array of objects each containing the following keys: name, manager, version, size, architecture, buildtime and installtime (see more details).

system.swap.in[<device>,<type>]

 The swap-in (from device into memory) statistics.
 Return value: *Integer*.
 Supported platforms: Linux, FreeBSD, OpenBSD.

Parameters:

- **device** - specify the device used for swapping (Linux only) or *all* (default);
- **type** - possible values:
    - *count* (number of swapins, default on non-Linux platforms);
    - *sectors* (sectors swapped in);
    - *pages* (pages swapped in, default on Linux).

Comments:

- The source of this information is:
    - `/proc/swaps`, `/proc/partitions`, `/proc/stat` (Linux 2.4);
    - `/proc/swaps`, `/proc/diskstats`, `/proc/vmstat` (Linux 2.6).
- Note that *pages* will only work if device was not specified.
- The *sectors* type parameter is supported only on Linux.

Example:

```
system.swap.in[,pages]
```

system.swap.out[<device>,<type>]

 The swap-out (from memory onto device) statistics.
 Return value: *Integer*.
 Supported platforms: Linux, FreeBSD, OpenBSD.

Parameters:

- **device** - specify the device used for swapping (Linux only) or *all* (default);
- **type** - possible values:
    - *count* (number of swapouts, default on non-Linux platforms);
    - *sectors* (sectors swapped out);
    - *pages* (pages swapped out, default on Linux).

Comments:

- The source of this information is:
    - `/proc/swaps`, `/proc/partitions`, `/proc/stat` (Linux 2.4);
    - `/proc/swaps`, `/proc/diskstats`, `/proc/vmstat` (Linux 2.6).
- Note that *pages* will only work if device was not specified.
- The *sectors* type parameter is supported only on Linux.

Example:

```
system.swap.out[,pages]
```

system.swap.size[<device>,<type>]

 The swap space size in bytes or in percentage from total.
 Return value: *Integer* - for bytes; *Float* - for percentage.
 Supported platforms: Linux, FreeBSD, Solaris, AIX, OpenBSD, Windows.

Parameters:

- **device** - specify the device used for swapping (FreeBSD only) or *all* (default);
- **type** - possible values:
  - *free* (free swap space, default);
  - *pfree* (free swap space, in percent);
  - *pused* (used swap space, in percent);
  - *total* (total swap space);
  - *used* (used swap space).

Comments:

- Note that *pfree*, *pused* are not supported on Windows if swap size is 0.
- If device is not specified Zabbix agent will only take into account swap devices (files), the physical memory will be ignored. For example, on Solaris systems the `swap -s` command includes a portion of physical memory and swap devices (unlike `swap -l`).

Example:

```
system.swap.size[,pfree] → free swap space percentage
```

system.uname

 Identification of the system.
 Return value: *String*.
 See supported platforms.

Comments:

- On UNIX the value for this item is obtained with the uname() system call.
- On Windows the item returns the OS architecture, whereas on UNIX it returns the CPU architecture.

Examples:

```
system.uname → FreeBSD localhost 4.2-RELEASE FreeBSD 4.2-RELEASE #0: Mon Nov i386
system.uname → Windows ZABBIX-WIN 6.0.6001 Microsoft® Windows Server® 2008 Standard Service Pack 1 x86
```

system.uptime

 The system uptime in seconds.
 Return value: *Integer*.
 Supported platforms: Linux, FreeBSD, Solaris, AIX, MacOS X, OpenBSD, NetBSD, Windows.

In item configuration, use **s** or **uptime** units to get readable values.

system.users.num

 The number of users logged in.
 Return value: *Integer*.
 See supported platforms.

The **who** command is used on the agent side to obtain the value.

vfs.dev.discovery

 The list of block devices and their type. Used for low-level discovery.
 Return value: *JSON string*.
 Supported platforms: Linux.

vfs.dev.get[<mode>,<regexp>]

 The list of block devices with detailed information. Used for low-level discovery.
 Return value: *JSON string*.
 Supported platforms: Linux, AIX.

Parameters:

- **mode** - possible values:
  - *disks* (default) - returns list of disks;
  - *disk_stats* - returns statistics for each disk;
  - *devices* (linux only) - returns list of block devices (disks and partitions);
  - *device_stats* (linux only) - returns statistics for each device (disks and partitions).
- **regexp** - a regular expression used to filter block devices by their name.

vfs.dev.read[<device>,<type>,<mode>]

 The disk read statistics.
 Return value: - *Integer* - with `type` in *sectors*, *operations*, *bytes*; - *Float* - with `type` in *sps*, *ops*, *bps*.
 Supported platforms: Linux, FreeBSD, Solaris, AIX, OpenBSD.

Parameters:

- **device** - disk device (default is *all* [3]);
- **type** - possible values: *sectors*, *operations*, *bytes*, *sps*, *ops*, or *bps* (*sps*, *ops*, *bps* stand for: sectors, operations, bytes per second, respectively);
- **mode** - possible values: *avg1* (one-minute average, default), *avg5*, or *avg15*. This parameter is supported only with `type` in: sps, ops, bps.

Comments:

- If using an update interval of three hours or more[2], this item will always return 0.
- The *sectors* and *sps* type parameters are supported only on Linux.
- The *ops* type parameter is supported only on Linux and FreeBSD.
- The *bps* type parameter is supported only on FreeBSD.
- The *bytes* type parameter is supported only on FreeBSD, Solaris, AIX, OpenBSD.
- The `mode` parameter is supported only on Linux, FreeBSD.
- You may use relative device names (for example, `sda`) as well as an optional `/dev/` prefix (for example, `/dev/sda`).
- LVM logical volumes are supported.
- The default values of 'type' parameter for different OSes:
  - *AIX* - operations;
  - *FreeBSD* - bps;
  - *Linux* - sps;
  - *OpenBSD* - operations;
  - *Solaris* - bytes.
- *sps*, *ops* and *bps* on supported platforms is limited to 1024 devices (1023 individual and one for *all*).

Example:

```
vfs.dev.read[,operations]
```

vfs.dev.write[<device>,<type>,<mode>]

 The disk write statistics.
 Return value: *Integer* - with `type` in *sectors*, *operations*, *bytes*; *Float* - with `type` in *sps*, *ops*, *bps*.
 Supported platforms: Linux, FreeBSD, Solaris, AIX, OpenBSD.

Parameters:

- **device** - disk device (default is *all* [3]);
- **type** - possible values: *sectors*, *operations*, *bytes*, *sps*, *ops*, or *bps* (*sps*, *ops*, *bps* stand for: sectors, operations, bytes per second, respectively);
- **mode** - possible values: *avg1* (one-minute average, default), *avg5*, or *avg15*. This parameter is supported only with `type` in: sps, ops, bps.

Comments:

- If using an update interval of three hours or more[2], this item will always return 0.
- The *sectors* and *sps* type parameters are supported only on Linux.
- The *ops* type parameter is supported only on Linux and FreeBSD.
- The *bps* type parameter is supported only on FreeBSD.
- The *bytes* type parameter is supported only on FreeBSD, Solaris, AIX, OpenBSD.
- The `mode` parameter is supported only on Linux, FreeBSD.
- You may use relative device names (for example, `sda`) as well as an optional `/dev/` prefix (for example, `/dev/sda`).
- LVM logical volumes are supported.
- The default values of `type` parameter for different OSes:
  - *AIX* - operations
  - *FreeBSD* - bps
  - *Linux* - sps
  - *OpenBSD* - operations
  - *Solaris* - bytes
- *sps*, *ops* and *bps* on supported platforms is limited to 1024 devices (1023 individual and one for *all*).

Example:

```
vfs.dev.write[,operations]
```

vfs.dir.count[dir,<regex incl>,<regex excl>,<types incl>,<types excl>,<max depth>,<min size>,<max size>,<min age>,<max age>,<regex excl dir>]

 The directory entry count.
 Return value: *Integer*.
 See supported platforms.

Parameters:

- **dir** - the absolute path to directory;
- **regex incl** - a regular expression describing the name pattern of the entity (file, directory, symbolic link) to include; include all if empty (default value);
- **regex excl** - a regular expression describing the name pattern of the entity (file, directory, symbolic link) to exclude; don't exclude any if empty (default value);
- **types incl** - directory entry types to count, possible values:
  - *file* - regular file;
  - *dir* - subdirectory;
  - *sym* - symbolic link;
  - *sock* - socket;
  - *bdev* - block device;
  - *cdev* - character device;
  - *fifo* - FIFO;
  - *dev* - synonymous with "bdev,cdev";
  - *all* - all types (default), i.e. "file,dir,sym,sock,bdev,cdev,fifo". Multiple types must be separated with comma and quoted.
- **types excl** - directory entry types (see `types incl`) to NOT count. If some entry type is in both `types incl` and `types excl`, directory entries of this type are NOT counted.
- **max depth** - the maximum depth of subdirectories to traverse:
  - **-1** (default) - unlimited;
  - **0** - no descending into subdirectories.
- **min size** - the minimum size (in bytes) for file to be counted. Smaller files will not be counted. Memory suffixes can be used.
- **max size** - the maximum size (in bytes) for file to be counted. Larger files will not be counted. Memory suffixes can be used.
- **min age** - the minimum age (in seconds) of directory entry to be counted. More recent entries will not be counted. Time suffixes can be used.
- **max age** - the maximum age (in seconds) of directory entry to be counted. Entries so old and older will not be counted (modification time). Time suffixes can be used.
- **regex excl dir** - a regular expression describing the name pattern of the directory to exclude. All content of the directory will be excluded (in contrast to `regex_excl`)

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- Environment variables, e.g., `%APP_HOME%`, `$HOME` and `%TEMP%` are not supported.
- Pseudo-directories "." and ".." are never counted.
- Symbolic links are never followed for directory traversal.
- Both `regex incl` and `regex excl` are being applied to files and directories when calculating the entry count, but are ignored when picking subdirectories to traverse (if `regex incl` is `(?i)^.+\.zip$` and `max depth` is not set, then all subdirectories will be traversed, but only the files of type zip will be counted).
- The execution time is limited by the timeout value in the agent configuration (3 sec). Since large directory traversal may take longer than that, no data will be returned and the item will turn unsupported. Partial count will not be returned.
- When filtering by size, only regular files have meaningful sizes. Under Linux and BSD, directories also have non-zero sizes (a few Kb typically). Devices have zero sizes, e.g., the size of `/dev/sda1` does not reflect the respective partition size. Therefore, when using `<min_size>` and `<max_size>`, it is recommended to specify `<types_incl>` as "*file*", to avoid surprises.

Examples:

```
vfs.dir.count[/dev] #monitors the number of devices in /dev (Linux)
vfs.dir.count["C:\Users\ADMINI~1\AppData\Local\Temp"] #monitors the number of files in a temporary directo
```

vfs.dir.get[dir,<regex incl>,<regex excl>,<types incl>,<types excl>,<max depth>,<min size>,<max size>,<min age>,<max age>,<regex excl dir>]

 The directory entry list.
 Return value: *JSON string*.
 See supported platforms.

Parameters:

- **dir** - the absolute path to directory;
- **regex incl** - a regular expression describing the name pattern of the entity (file, directory, symbolic link) to include; include all if empty (default value);

- **regex excl** - a regular expression describing the name pattern of the entity (file, directory, symbolic link) to exclude; don't exclude any if empty (default value);
- **types incl** - directory entry types to list, possible values:
  - *file* - regular file;
  - *dir* - subdirectory;
  - *sym* - symbolic link;
  - *sock* - socket;
  - *bdev* - block device;
  - *cdev* - character device;
  - *fifo* - FIFO;
  - *dev* - synonymous with "bdev,cdev";
  - *all* - all types (default), i.e. "file,dir,sym,sock,bdev,cdev,fifo". Multiple types must be separated with comma and quoted.
- **types excl** - directory entry types (see `types incl`) to NOT list. If some entry type is in both `types incl` and `types excl`, directory entries of this type are NOT listed.
- **max depth** - the maximum depth of subdirectories to traverse:
  - **-1** (default) - unlimited;
  - **0** - no descending into subdirectories.
- **min size** - the minimum size (in bytes) for file to be listed. Smaller files will not be listed. Memory suffixes can be used.
- **max size** - the maximum size (in bytes) for file to be listed. Larger files will not be listed. Memory suffixes can be used.
- **min age** - the minimum age (in seconds) of directory entry to be listed. More recent entries will not be listed. Time suffixes can be used.
- **max age** - the maximum age (in seconds) of directory entry to be listed. Entries so old and older will not be listed (modification time). Time suffixes can be used.
- **regex excl dir** - a regular expression describing the name pattern of the directory to exclude. All content of the directory will be excluded (in contrast to `regex excl`)

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- Environment variables, e.g., `%APP_HOME%`, `$HOME` and `%TEMP%` are not supported.
- Pseudo-directories "." and ".." are never listed.
- Symbolic links are never followed for directory traversal.
- Both `regex incl` and `regex excl` are being applied to files and directories when generating the entry list, but are ignored when picking subdirectories to traverse (if `regex incl` is `(?i)^.+\.zip$` and `max depth` is not set, then all subdirectories will be traversed, but only the files of type zip will be counted).
- The execution time is limited by the timeout value in the agent configuration. Since large directory traversal may take longer than that, no data will be returned and the item will turn unsupported. Partial list will not be returned.
- When filtering by size, only regular files have meaningful sizes. Under Linux and BSD, directories also have non-zero sizes (a few Kb typically). Devices have zero sizes, e.g., the size of `/dev/sda1` does not reflect the respective partition size. Therefore, when using `min size` and `max size`, it is recommended to specify `types incl` as "*file*", to avoid surprises.

Examples:

```
vfs.dir.get[/dev] #retrieves the device list in /dev (Linux)
vfs.dir.get["C:\Users\ADMINI~1\AppData\Local\Temp"] #retrieves the file list in a temporary directory
```

vfs.dir.size[dir,<regex incl>,<regex excl>,<mode>,<max depth>,<regex excl dir>]

 The directory size (in bytes).
 Return value: *Integer*.
 Supported platforms: Linux, Windows. The item may work on other UNIX-like platforms.

Parameters:

- **dir** - the absolute path to directory;
- **regex incl** - a regular expression describing the name pattern of the entity (file, directory, symbolic link) to include; include all if empty (default value);
- **regex excl** - a regular expression describing the name pattern of the entity (file, directory, symbolic link) to exclude; don't exclude any if empty (default value);
- **mode** - possible values:
  - *apparent* (default) - gets apparent file sizes rather than disk usage (acts as `du -sb dir`);
  - *disk* - gets disk usage (acts as `du -s -B1 dir`). Unlike the du command, the vfs.dir.size item takes hidden files in account when calculating the directory size (acts as `du -sb .[^.]* *` within dir).
- **max depth** - the maximum depth of subdirectories to traverse:
  - **-1** (default) - unlimited;
  - **0** - no descending into subdirectories.

- **regex excl dir** - a regular expression describing the name pattern of the directory to exclude. All content of the directory will be excluded (in contrast to `regex excl`)

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- Only directories with at least the read permission for *zabbix* user are calculated. For directories with read permission only, the size of the directory itself is calculated. Directories with read & execute permissions are calculated including contents.
- With large directories or slow drives this item may time out due to the Timeout setting in agent and server/proxy configuration files. Increase the timeout values as necessary.
- The file size limit depends on large file support.

Examples:

```
vfs.dir.size[/tmp,log] #calculates the size of all files in /tmp containing 'log' in their names
vfs.dir.size[/tmp,log,^.+\.old$] #calculates the size of all files in /tmp containing 'log' in their names
```

vfs.file.cksum[file,<mode>]

 The file checksum, calculated by the UNIX cksum algorithm.
 Return value: *Integer* - with `mode` as *crc32*, *String* - with `mode` as *md5*, *sha256*.
 See supported platforms.

Parameters:

- **file** - the full path to file;
- **mode** - *crc32* (default), *md5*, or *sha256*.

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- The file size limit depends on large file support.

Example:

```
vfs.file.cksum[/etc/passwd]
```

Examples of returned values (crc32/md5/sha256 respectively):

```
675436101
9845acf68b73991eb7fd7ee0ded23c44
ae67546e4aac995e5c921042d0cf0f1f7147703aa42bfbfb65404b30f238f2dc
```

vfs.file.contents[file,<encoding>]

 Retrieving the contents of a file[7].
 Return value: *Text*.
 See supported platforms.

Parameters:

- **file** - the full path to file;
- **encoding** - the code page identifier.

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- The return value is limited to 16MB (including trailing whitespace that is truncated); database limits also apply.
- An empty string is returned if the file is empty or contains LF/CR characters only.
- The byte order mark (BOM) is excluded from the output.

Example:

```
vfs.file.contents[/etc/passwd]
```

vfs.file.exists[file,<types incl>,<types excl>]

 Checks if the file exists.
 Return value: 0 - not found; 1 - file of the specified type exists.
 See supported platforms.

Parameters:

- **file** - the full path to file;
- **types incl** - the list of file types to include, possible values:
  - *file* (regular file, default (if `types_excl` is not set));
  - *dir* (directory);
  - *sym* (symbolic link);
  - *sock* (socket);
  - *bdev* (block device);

  - *cdev* (character device);
  - *fifo* (FIFO);
  - *dev* (synonymous with "bdev,cdev");
  - *all* (all mentioned types, default if `types_excl` is set).
- **types excl** - the list of file types to exclude, see `types_incl` for possible values (by default no types are excluded)

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- Multiple types must be separated with a comma and the entire set enclosed in quotes "".
- If the same type is in both <types_incl> and <types_excl>, files of this type are excluded.
- The file size limit depends on large file support.

Examples:

```
vfs.file.exists[/tmp/application.pid]
vfs.file.exists[/tmp/application.pid,"file,dir,sym"]
vfs.file.exists[/tmp/application_dir,dir]
```

vfs.file.get[file]

 Returns information about a file.
 Return value: *JSON string*.
 See supported platforms.

Parameter:

- **file** - the full path to file

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- Supported file types on UNIX-like systems: regular file, directory, symbolic link, socket, block device, character device, FIFO.
- The file size limit depends on large file support.

Example:

```
vfs.file.get[/etc/passwd] #return a JSON with information about the /etc/passwd file (type, user, permissi
```

vfs.file.md5sum[file]

 The MD5 checksum of file.
 Return value: Character string (MD5 hash of the file).
 See supported platforms.

Parameter:

- **file** - the full path to file

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- The file size limit depends on large file support.

Example:

```
vfs.file.md5sum[/usr/local/etc/zabbix_agentd.conf]
```

Example of returned value:

```
b5052decb577e0fffd622d6ddc017e82
```

vfs.file.owner[file,<ownertype>,<resulttype>]

 Retrieves the owner of a file.
 Return value: *String*.
 See supported platforms.

Parameters:

- **file** - the full path to file;
- **ownertype** - *user* (default) or *group* (Unix only);
- **resulttype** - *name* (default) or *id*; for id - return uid/gid on Unix, SID on Windows.

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- The file size limit depends on large file support.

Examples:

```
vfs.file.owner[/tmp/zabbix_server.log] #return the file owner of /tmp/zabbix_server.log
vfs.file.owner[/tmp/zabbix_server.log,,id] #return the file owner ID of /tmp/zabbix_server.log
```

vfs.file.permissions[file]

 Return a 4-digit string containing the octal number with UNIX permissions.
 Return value: *String*.
 Supported platforms: Linux. The item may work on other UNIX-like platforms.

Parameters:

- **file** - the full path to file

The file size limit depends on large file support.

Example:

```
vfs.file.permissions[/etc/passwd] #return permissions of /etc/passwd, for example, '0644'
```

vfs.file.regexp[file,regexp,<encoding>,<start line>,<end line>,<output>]

 Retrieve a string in the file[7].
 Return value: The line containing the matched string, or as specified by the optional `output` parameter.
 See supported platforms.

Parameters:

- **file** - the full path to file;
- **regexp** - a regular expression describing the required pattern;
- **encoding** - the code page identifier;
- **start line** - the number of the first line to search (first line of file by default);
- **end line** - the number of the last line to search (last line of file by default);
- **output** - an optional output formatting template. The **\0** escape sequence is replaced with the matched part of text (from the first character where match begins until the character where match ends) while an **\N** (where N=1...9) escape sequence is replaced with Nth matched group (or an empty string if the N exceeds the number of captured groups).

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- The file size limit depends on large file support.
- Only the first matching line is returned.
- An empty string is returned if no line matched the expression.
- The byte order mark (BOM) is excluded from the output.
- Content extraction using the `output` parameter takes place on the agent.

Examples:

```
vfs.file.regexp[/etc/passwd,zabbix]
vfs.file.regexp[/path/to/some/file,"([0-9]+)$",,3,5,\1]
vfs.file.regexp[/etc/passwd,"^zabbix:.:([0-9]+)",,,,\1] → getting the ID of user *zabbix*
```

vfs.file.regmatch[file,regexp,<encoding>,<start line>,<end line>]

 Find a string in the file[7].
 Return values: 0 - match not found; 1 - found.
 See supported platforms.

Parameters:

- **file** - the full path to file;
- **regexp** - a regular expression describing the required pattern;
- **encoding** - the code page identifier;
- **start line** - the number of the first line to search (first line of file by default);
- **end line** - the number of the last line to search (last line of file by default).

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- The file size limit depends on large file support.
- The byte order mark (BOM) is ignored.

Example:

```
vfs.file.regmatch[/var/log/app.log,error]
```

vfs.file.size[file,<mode>]

 The file size (in bytes).
 Return value: *Integer*.
 See supported platforms.

Parameters:

- **file** - the full path to file;

- **mode** - possible values:
  - *bytes* (default);
  - *lines* (empty lines are counted, too).

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- The file must have read permissions for user *zabbix*.
- The file size limit depends on large file support.

Example:

```
vfs.file.size[/var/log/syslog]
```

vfs.file.time[file,<mode>]

 The file time information.
 Return value: *Integer* (Unix timestamp).
 See supported platforms.

Parameters:

- **file** - the full path to file;
- **mode** - possible values:
  - *modify* (default) - the last time of modifying file content (default);
  - *access* - the last time of reading file;
  - *change* - the last time of changing file properties.

Comments:

- On Windows, this item follows UNC paths, which can be a security risk.
- The file size limit depends on large file support.

Example:

```
vfs.file.time[/etc/passwd,modify]
```

vfs.fs.discovery

 The list of mounted filesystems with their type and mount options. Used for low-level discovery.
 Return value: *JSON string*.
 Supported platforms: Linux, FreeBSD, Solaris, HP-UX, AIX, MacOS X, OpenBSD, NetBSD, Windows.

vfs.fs.get[<mode>,<mountpoint>]

 The list of mounted filesystems with their type and mount options and, in the *full* mode, available disk space and inode statistics. Can be used for low-level discovery.
 Return value: *JSON string*.
 Supported platforms: Linux, FreeBSD, Solaris, HP-UX, AIX, MacOS X, OpenBSD, NetBSD, Windows.

Parameters:

- **mode** - possible values:
  - *full* (default) - filesystem properties with disk space and inode statistics;
  - *short* - filesystem properties only, without disk space and inode statistics;
- **mountpoint** - the mount point to return data for (for example, `/var` or `C:`). If not specified, all mounted filesystems are returned.

Comments:

- In the *short* mode, only the filesystem name, type and mount options are returned (on Windows: name, type, label and drive type). The filesystems are not queried for statistics; on Unix-like systems only the list of mounted filesystems is read (for example, `/proc/mounts` on Linux), so the check does not trigger automounting and is not affected by unresponsive network filesystems. This makes the *short* mode suitable for low-level discovery.
- The **mountpoint** value must match the `fsname` property exactly. On Windows, the comparison is case-insensitive; specify the mount point without the trailing backslash (for example, `C:`).
- If **mountpoint** is specified, data is returned for that mount point only. If no mounted filesystem matches it, an empty array (`[]`) is returned.
- The output may contain several entries with the same `fsname` if more than one filesystem is mounted on the same directory.
- In the *full* mode, available disk space and inode statistics are returned only if the Zabbix agent user has traverse permission on every directory in the mountpoint path; otherwise, the filesystem is omitted from the output. This does not apply to the *short* mode, which may therefore list filesystems that are missing from the *full* mode output.
- The `fslabel` and `fsdrivetype` properties are returned on Windows only; the `inodes` and `options` properties are not returned on Windows.
- File systems with the inode count equal to zero, which can be the case for file systems with dynamic inodes (e.g., btrfs), are also reported.

- See also: Discovery of mounted filesystems.

Examples:

```
vfs.fs.get #all mounted filesystems with statistics; same as vfs.fs.get[full]
vfs.fs.get[short] #all mounted filesystems without statistics
vfs.fs.get[full,/var] #statistics for the /var mount point only
vfs.fs.get[short,C:] #properties of the C: volume only (Windows)
```

vfs.fs.inode[fs,<mode>]

 The number or percentage of inodes.
 Return value: *Integer* - for number; *Float* - for percentage.
 See supported platforms.

Parameters:

- **fs** - the filesystem;
- **mode** - possible values: *total* (default), *free*, *used*, *pfree* (free, percentage), or *pused* (used, percentage).

If the inode count equals zero, which can be the case for file systems with dynamic inodes (e.g., btrfs), the pfree/pused values will be reported as ”100” and ”0” respectively.

Example:

```
vfs.fs.inode[/,pfree]
```

vfs.fs.size[fs,<mode>]

 The disk space in bytes or in percentage from total.
 Return value: *Integer* - for bytes; *Float* - for percentage.
 See supported platforms.

Parameters:

- **fs** - the filesystem;
- **mode** - possible values: *total* (default), *free*, *used*, *pfree* (free, percentage), or *pused* (used, percentage).

Comments:

- If the filesystem is not mounted, returns the size of a local filesystem where the mount point is located.
- The reserved space of a file system is taken into account and not included when using the *free* mode.

Example:

```
vfs.fs.size[/tmp,free]
```

vm.memory.size[<mode>]

 The memory size in bytes or in percentage from total.
 Return value: *Integer* - for bytes; *Float* - for percentage.
 See supported platforms.

Parameter:

- **mode** - possible values: *total* (default), *active*, *anon*, *buffers*, *cached*, *exec*, *file*, *free*, *inactive*, *pinned*, *shared*, *slab*, *wired*, *used*, *pused* (used, percentage), *available*, or *pavailable* (available, percentage).

Comments:

- This item accepts three categories of parameters:
  - *total* - total amount of memory;
  - platform-specific memory types: *active*, *anon*, *buffers*, *cached*, *exec*, *file*, *free*, *inactive*, *pinned*, *shared*, *slab*, *wired*;
  - user-level estimates on how much memory is used and available: *used*, *pused*, *available*, *pavailable*.
- The *active* mode parameter is supported only on FreeBSD, HP-UX, MacOS X, OpenBSD, NetBSD.
- The *anon*, *exec*, *file* mode parameters are supported only on NetBSD.
- The *buffers* mode parameter is supported only on Linux, FreeBSD, OpenBSD, NetBSD.
- The *cached* mode parameter is supported only on Linux, FreeBSD, AIX, OpenBSD, NetBSD.
- The *inactive*, *wired* mode parameters are supported only on FreeBSD, MacOS X, OpenBSD, NetBSD.
- The *pinned* mode parameter is supported only on AIX.
- The *shared* mode parameter is supported only on Linux 2.4, FreeBSD, OpenBSD, NetBSD.
- See also additional details for this item.

Example:

```
vm.memory.size[pavailable]
```

web.page.get[host,<path>,<port>,<redirect_limit>]

 Get the content of a web page.
 Return value: Web page source as text (including headers).
 See supported platforms.

Parameters:

- **host** - the hostname or URL (as `scheme://host:port/path`, where only *host* is mandatory). Allowed URL schemes: *http*, *https*[4]. A missing scheme will be treated as *http*. If a URL is specified `path` and `port` must be empty. Specifying user name/password when connecting to servers that require authentication, for example: `http://user:password@www.example.com` is only possible with cURL support [4]. Punycode is supported in hostnames.
- **path** - the path to an HTML document (default is /);
- **port** - the port number (default is 80 for HTTP)
- **redirect_limit** - the maximum number of redirects to follow [4], from 0 to 1000 (default is 0 - redirects are not followed).

Comments:

- This item turns unsupported if the resource specified in `host` does not exist or is unavailable, or if `redirect_limit` is greater than 0 and the number of redirects exceeds it.
- `host` can be a hostname, domain name, IPv4 or IPv6 address. But for IPv6 address Zabbix agent must be compiled with IPv6 support enabled.
- If redirects are followed, the returned value contains the headers of each redirect response, followed by the headers and body of the final response.

Examples:

```
web.page.get[www.example.com,index.php,80]
web.page.get[https://www.example.com]
web.page.get[https://blog.example.com/?s=zabbix]
web.page.get[localhost:80]
web.page.get["[::1]/server-status"]
web.page.get[http://www.example.com,,,5]
```

web.page.perf[host,<path>,<port>,<redirect_limit>]

 The loading time of a full web page (in seconds).
 Return value: *Float*.
 See supported platforms.

Parameters:

- **host** - the hostname or URL (as `scheme://host:port/path`, where only *host* is mandatory). Allowed URL schemes: *http*, *https*[4]. A missing scheme will be treated as *http*. If a URL is specified `path` and `port` must be empty. Specifying user name/password when connecting to servers that require authentication, for example: `http://user:password@www.example.com` is only possible with cURL support [4]. Punycode is supported in hostnames.
- **path** - the path to an HTML document (default is /);
- **port** - the port number (default is 80 for HTTP)
- **redirect_limit** - the maximum number of redirects to follow [4], from 0 to 1000 (default is 0 - redirects are not followed).

Comments:

- This item turns unsupported if the resource specified in `host` does not exist or is unavailable, or if `redirect_limit` is greater than 0 and the number of redirects exceeds it.
- `host` can be a hostname, domain name, IPv4 or IPv6 address. But for IPv6 address Zabbix agent must be compiled with IPv6 support enabled.
- If redirects are followed, the loading time includes the time spent on all redirects.

Examples:

```
web.page.perf[www.example.com,index.php,80]
web.page.perf[https://www.example.com]
web.page.perf[http://www.example.com,,,5]
```

web.page.regexp[host,<path>,<port>,regexp,<length>,<output>,<redirect_limit>]

 Find a string on the web page.
 Return value: The matched string, or as specified by the optional `output` parameter.
 See supported platforms.

Parameters:

- **host** - the hostname or URL (as `scheme://host:port/path`, where only *host* is mandatory). Allowed URL schemes: *http*, *https*[4]. A missing scheme will be treated as *http*. If a URL is specified `path` and `port` must be empty. Specifying user name/password when connecting to servers that require authentication, for example: `http://user:password@www.example.com` is only possible with cURL support [4]. Punycode is supported in hostnames.

- **path** - the path to an HTML document (default is /).
- **port** - the port number (default is 80 for HTTP).
- **regexp** - a regular expression describing the required pattern.
- **length** - the maximum number of characters to return;
- **output** - an optional output formatting template. The **\0** escape sequence is replaced with the matched part of text (from the first character where match begins until the character where match ends) while an **\N** (where N=1...9) escape sequence is replaced with Nth matched group (or an empty string if the N exceeds the number of captured groups).
- **redirect_limit** - the maximum number of redirects to follow [4], from 0 to 1000 (default is 0 - redirects are not followed).

Comments:

- This item turns unsupported if the resource specified in `host` does not exist or is unavailable, or if `redirect_limit` is greater than 0 and the number of redirects exceeds it.
- `host` can be a hostname, domain name, IPv4 or IPv6 address. But for IPv6 address Zabbix agent must be compiled with IPv6 support enabled.
- Content extraction using the `output` parameter takes place on the agent.
- If redirects are followed, the headers of each redirect response are searched as well.

Examples:

```
web.page.regexp[www.example.com,index.php,80,OK,2]
web.page.regexp[https://www.example.com,,,OK,2]
web.page.regexp[http://www.example.com,,,OK,2,,5]
```

agent.hostmetadata

 The agent host metadata.
 Return value: *String*.
 See supported platforms.

Returns the value of `HostMetadata` or `HostMetadataItem` parameters, or empty string if none are defined.

agent.hostname

 The agent host name.
 Return value: *String*.
 See supported platforms.

Returns:

- As passive check - the name of the first host listed in the Hostname parameter of the agent configuration file;
- As active check - the name of the current hostname.

agent.ping

 The agent availability check.
 Return value: Nothing - unavailable; 1 - available.
 See supported platforms.

Use the `nodata()` trigger function to check for host unavailability.

agent.variant

 The variant of Zabbix agent (Zabbix agent or Zabbix agent 2).
 Return value: 1 - Zabbix agent; 2 - Zabbix agent 2.
 See supported platforms.

agent.version

 The version of Zabbix agent.
 Return value: *String*.
 See supported platforms.

Example of returned value:

```
6.0.3
```

zabbix.stats[<ip>,<port>]

 Returns a set of Zabbix server or proxy internal metrics. If `ip` and `port` are specified, metrics are retrieved from the remote instance; otherwise, from the local instance.
 Return value: *JSON string*.
 See supported platforms.

Parameters:

- **ip** - the IP/DNS/network mask list of servers/proxies to be remotely queried (default is `127.0.0.1`);
- **port** - the port of server/proxy to be remotely queried (default is 10051)

Comments:

- A selected set of internal metrics is returned by this item. For details, see Remote monitoring of Zabbix stats.
- Note that the stats request will only be accepted from the addresses listed in the `StatsAllowedIP` server/proxy parameter on the target instance.

zabbix.stats[<ip>,<port>,queue,<from>,<to>]

 Returns the number of monitored items in the queue which are delayed on Zabbix server or proxy. If `ip` and `port` are specified, metrics are retrieved from the remote instance; otherwise, from the local instance.
 Return value: *JSON string*.
 See supported platforms.

Parameters:

- **ip** - the IP/DNS/network mask list of servers/proxies to be remotely queried (default is `127.0.0.1`);
- **port** - the port of server/proxy to be remotely queried (default is 10051)
- **queue** - constant (to be used as is)
- **from** - delayed by at least (default is 6 seconds)
- **to** - delayed by at most (default is infinity)

Note that the stats request will only be accepted from the addresses listed in the 'StatsAllowedIP' server/proxy parameter on the target instance.

Footnotes

[1] A Linux-specific note. Zabbix agent must have read-only access to filesystem `/proc`. Kernel patches from www.grsecurity.org limit access rights of non-privileged users.

[2] `vfs.dev.read[]`, `vfs.dev.write[]`: Zabbix agent will terminate "stale" device connections if the item values are not accessed for more than 3 hours. This may happen if a system has devices with dynamically changing paths or if a device gets manually removed. Note also that these items, if using an update interval of 3 hours or more, will always return '0'.

[3] `vfs.dev.read[]`, `vfs.dev.write[]`: If default `all` is used for the first parameter then the key will return summary statistics, including all block devices like sda, sdb, and their partitions (sda1, sda2, sdb3...) and multiple devices (MD raid) based on those block devices/partitions and logical volumes (LVM) based on those block devices/partitions. In such cases returned values should be considered only as relative value (dynamic in time) but not as absolute values.

[4] HTTPS and following redirects (`redirect_limit` greater than 0) are supported only if Zabbix agent is compiled with cURL support; otherwise, the item turns unsupported. Zabbix agent 2 does not require cURL for either.

[5] The `bytes` and `errors` values are not supported for loopback interfaces on Solaris systems up to and including Solaris 10 6/06 as byte, error and utilization statistics are not stored and/or reported by the kernel. However, if you're monitoring a Solaris system via net-snmp, values may be returned as net-snmp carries legacy code from the cmu-snmp dated as old as 1997 that, upon failing to read byte values from the interface statistics returns the packet counter (which does exist on loopback interfaces) multiplied by an arbitrary value of 308. This makes the assumption that the average length of a packet is 308 octets, which is a very rough estimation as the MTU limit on Solaris systems for loopback interfaces is 8892 bytes. These values should not be assumed to be correct or even closely accurate. They are guestimates. Zabbix agent does not do any guess work, but net-snmp will return a value for these fields.

[6] The command line on Solaris, obtained from `/proc/pid/psinfo`, is limited to 80 bytes and contains the command line as it was when the process was started.

[7] `vfs.file.contents[]`, `vfs.file.regexp[]`, `vfs.file.regmatch[]` items can be used for retrieving file contents. If you want to restrict access to specific files with sensitive information, run Zabbix agent under a user that has no access permissions to viewing these files.

Usage with command-line utilities

Note that when testing or using item keys with zabbix_agentd or zabbix_get from the command line you should consider shell syntax too.

For example, if a certain parameter of the key has to be enclosed in double quotes you have to explicitly escape double quotes; otherwise, they will be trimmed by the shell as special characters and will not be passed to the Zabbix utility.

Examples:

```
zabbix_agentd -t 'vfs.dir.count[/var/log,,,"file,dir",,0]'

zabbix_agentd -t vfs.dir.count[/var/log,,,\"file,dir\",,0]
```

Encoding settings

To make sure that the acquired data is not corrupted you may specify the correct encoding for processing the check (e.g., `vfs.file.contents`) in the `encoding` parameter. The list of supported encodings (code page identifiers) may be found in documentation for libiconv (GNU Project) or in Microsoft Windows SDK documentation for "Code Page Identifiers". Note that Microsoft sometimes marks some code pages as "available only to managed applications", meaning they are exposed only via the .NET runtime and not available through the native Win32 APIs. Zabbix agent implements its own encoding-conversion logic, therefore those code pages are supported by Zabbix agent even when the native Windows functions do not provide them.

If no encoding is specified in the `encoding` parameter the following resolution strategies are applied:

- If encoding is not specified (or is an empty string) it is assumed to be UTF-8, the data is processed "as-is";
- BOM analysis - applicable for items `vfs.file.contents`, `vfs.file.regexp`, `vfs.file.regmatch`. An attempt is made to determine the correct encoding by using the byte order mark (BOM) at the beginning of the file. If BOM is not present - standard resolution (see above) is applied instead.

Troubleshooting agent items

In case of passive checks, to prevent that item does not get any value because the server request to the agent times out first, the following should be noted:

- Where agent version is older than server version, the `Timeout` value in the item configuration (or global timeout) may need to be higher than the `Timeout` value in the agent configuration file.
- Where agent version is newer than server version, the `Timeout` value in the server configuration file may need to be higher than the `Timeout` value in the agent configuration file.

Zabbix agent 2

Zabbix agent 2 supports all item keys supported for Zabbix agent on Unix and Windows. This page provides details on the additional item keys, which you can use with Zabbix agent 2 only, grouped by the plugin they belong to.

The item keys are listed without parameters and additional information. Click the item key to see the full details.

Item key	Description	Plugin
ceph.df.details	The cluster's data usage and distribution among pools.	Ceph (loadable)
ceph.osd.stats	Aggregated and per OSD statistics.	
ceph.osd.discovery	The list of discovered OSDs.	
ceph.osd.dump	The usage thresholds and statuses of OSDs.	
ceph.ping	Tests whether a connection to Ceph can be established.	
ceph.pool.discovery	The list of discovered pools.	
ceph.status	The overall cluster's status.	
docker.container_info	Low-level information about a container.	Docker
docker.container_stats	The container resource usage statistics.	
docker.containers	Returns the list of containers.	
docker.containers.discovery	Returns the list of containers. Used for low-level discovery.	
docker.data.usage	Information about the current data usage.	
docker.images	Returns the list of images.	
docker.images.discovery	Returns the list of images. Used for low-level discovery.	
docker.info	The system information.	
docker.ping	Test if the Docker daemon is alive or not.	
ember.get	Returns the result of the required device.	Ember+ (loadable)
memcached.ping	Test if a connection is alive or not.	Memcached
memcached.stats	Gets the output of the STATS command.	
mongodb.cfg.discovery	Returns a list of discovered configuration servers.	MongoDB (loadable)
mongodb.collection.stats	Returns a variety of storage statistics for a given collection.	
mongodb.collections.discovery	Returns a list of discovered collections.	
mongodb.collections.usage	Returns the usage statistics for collections.	
mongodb.connpool.stats	Returns information regarding the open outgoing connections from the current database instance to other members of the sharded cluster or replica set.	
mongodb.db.stats	Returns the statistics reflecting a given database system state.	
mongodb.db.discovery	Returns a list of discovered databases.	
mongodb.jumbo_chunks.count	Returns the count of jumbo chunks.	
mongodb.oplog.stats	Returns the status of the replica set, using data polled from the oplog.	
mongodb.ping	Test if a connection is alive or not.	
mongodb.rs.config	Returns the current configuration of the replica set.	
mongodb.rs.status	Returns the replica set status from the point of view of the member where the method is run.	
mongodb.server.status	Returns the database state.	
mongodb.sh.discovery	Returns the list of discovered shards present in the cluster.	

Item key	Description	Plugin
mongodb.version	Returns the database server version.	
mqtt.get	Subscribes to a specific topic or topics (with wildcards) of the provided broker and waits for publications.	MQTT
mssql.availability.group.get	Returns availability groups.	MSSQL (loadable)
mssql.custom.query	Returns the result of a custom query.	
mssql.db.get	Returns all available MSSQL databases.	
mssql.job.status.get	Returns the status of jobs.	
mssql.last.backup.get	Returns the last backup time for all databases.	
mssql.local.db.get	Returns databases that are participating in an Always On availability group and replica (primary or secondary) and are located on the server that the connection was established to.	
mssql.mirroring.get	Returns mirroring info.	
mssql.nonlocal.db.get	Returns databases that are participating in an Always On availability group and replica (primary or secondary) located on other servers (the database is not local to the SQL Server instance that the connection was established to).	
mssql.perfcounter.get	Returns the performance counters.	
mssql.ping	Test if a connection is alive or not.	
mssql.quorum.get	Returns the quorum info.	
mssql.quorum.member.get	Returns the quorum members.	
mssql.replica.get	Returns the replicas.	
mssql.version	Returns the MSSQL version.	
mysql.custom.query	Returns the result of a custom query.	MySQL
mysql.db.discovery	Returns the list of MySQL databases.	
mysql.db.size	The database size in bytes.	
mysql.get_status_variables	Values of the global status variables.	
mysql.ping	Test if a connection is alive or not.	
mysql.replication.discovery	Returns the list of MySQL replications.	
mysql.replication.get_slave_status	The replication status.	
mysql.version	The MySQL version.	
net.dns.get	Performs a DNS query and returns detailed DNS record information.	Network
nvml.device.count	The number of GPU devices.	NVIDIA GPU (loadable)
nvml.device.decoder.utilization	GPU device decoder utilization.	
nvml.device.ecc.mode	GPU device ECC mode.	
nvml.device.encoder.stats.get	GPU device encoder statistics.	
nvml.device.encoder.utilization	GPU device encoder utilization.	
nvml.device.energy.consumption	GPU device energy consumption.	
nvml.device.errors.memory	GPU device ECC memory error information.	
nvml.device.errors.register	GPU device ECC register error information.	
nvml.device.fan.speed.avg	GPU device fan speed average.	
nvml.device.get	Discovered GPUs with UUID and name.	
nvml.device.graphics.frequency	GPU device graphics clock speed.	
nvml.device.memory.bar1.get	GPU device BAR1 memory information.	
nvml.device.memory.fb.get	GPU device framebuffer memory information.	
nvml.device.memory.frequency	GPU device memory clock speed.	
nvml.device.pci.utilization	GPU device PCI utilization.	
nvml.device.performance.state	GPU device performance state.	
nvml.device.power.limit	GPU device power limit.	
nvml.device.power.usage	GPU device power usage.	
nvml.device.serial	GPU device serial number.	
nvml.device.sm.frequency	GPU device streaming multiprocessor clock speed.	
nvml.device.temperature	GPU device temperature.	
nvml.device.utilization	GPU device utilization statistics.	
nvml.device.video.frequency	GPU device video clock speed.	
nvml.system.driver.version	The NVIDIA driver version.	
nvml.version	The NVML library version.	
oracle.diskgroups.stats	Returns the Automatic Storage Management (ASM) disk groups statistics.	Oracle
oracle.diskgroups.discovery	Returns the list of ASM disk groups.	
oracle.archive.info	The archive logs statistics.	

Item key	Description	Plugin
oracle.cdb.info	The Container Databases (CDBs) information.	
oracle.custom.query	The result of a custom query.	
oracle.datafiles.stats	Returns the data files statistics.	
oracle.db.discovery	Returns the list of databases.	
oracle.fra.stats	Returns the Fast Recovery Area (FRA) statistics.	
oracle.instance.info	The instance statistics.	
oracle.pdb.info	The Pluggable Databases (PDBs) information.	
oracle.pdb.discovery	Returns the list of PDBs.	
oracle.pga.stats	Returns the Program Global Area (PGA) statistics.	
oracle.ping	Test whether a connection to Oracle can be established.	
oracle.proc.stats	Returns the processes statistics.	
oracle.redolog.info	The log file information from the control file.	
oracle.sga.stats	Returns the System Global Area (SGA) statistics.	
oracle.sessions.stats	Returns the sessions statistics.	
oracle.sys.metrics	Returns a set of system metric values.	
oracle.sys.params	Returns a set of system parameter values.	
oracle.ts.stats	Returns the tablespaces statistics.	
oracle.ts.discovery	Returns a list of tablespaces.	
oracle.user.info	Returns Oracle user information.	
oracle.version	Returns the database server version.	
pgsql.autovacuum.count	The number of autovacuum workers.	PostgreSQL (loadable)
pgsql.archive	The information about archived files.	
pgsql.bgwriter	The combined number of checkpoints for the database cluster, broken down by checkpoint type.	
pgsql.cache.hit	The PostgreSQL buffer cache hit rate.	
pgsql.connections	Returns connections by type.	
pgsql.custom.query	Returns the result of a custom query.	
pgsql.db.age	The age of the oldest FrozenXID of the database.	
pgsql.db.bloating_tables	The number of bloating tables per database.	
pgsql.db.discovery	The list of PostgreSQL databases.	
pgsql.db.size	The database size in bytes.	
pgsql.dbstat	Collects the statistics per database.	
pgsql.dbstat.sum	The summarized data for all databases in a cluster.	
pgsql.locks	The information about granted locks per database.	
pgsql.oldest.xid	The age of the oldest XID.	
pgsql.ping	Test if a connection is alive or not.	
pgsql.queries	Query metrics by execution time.	
pgsql.replication.count	The number of standby servers.	
pgsql.replication.process	The flush lag, write lag and replay lag per each sender process.	
pgsql.replication.process.discovery	The replication process name discovery.	
pgsql.replication.recovery_role	The recovery status.	
pgsql.replication.status	The status of replication.	
pgsql.replication_lag.b	The replication lag in bytes.	
pgsql.replication_lag.sec	The replication lag in seconds.	
pgsql.uptime	The PostgreSQL uptime in milliseconds.	
pgsql.version	Returns PostgreSQL version.	
pgsql.wal.stat	The WAL statistics.	
redis.config	Gets the configuration parameters of a Redis instance that match the pattern.	Redis
redis.info	Gets the output of the INFO command.	
redis.ping	Test if a connection is alive or not.	
redis.slowlog.count	The number of slow log entries since Redis was started.	
smart.attribute.discovery	Returns a list of S.M.A.R.T. device attributes.	S.M.A.R.T.
smart.disk.discovery	Returns a list of S.M.A.R.T. devices.	
smart.disk.get	Returns all available properties of S.M.A.R.T. devices.	
systemd.unit.get	Returns all properties of a systemd unit.	Systemd
systemd.unit.info	Systemd unit information.	
systemd.unit.discovery	The list of systemd units and their details.	
web.certificate.get	Validates the certificates and returns certificate details.	Web certificates

Item key details

Parameters without angle brackets are mandatory. Parameters marked with angle brackets **< >** are optional.

ceph.df.details[connString,<user>,<apikey>]

 The cluster's data usage and distribution among pools.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, apikey** - the user name and API key for connecting to Ceph.

ceph.osd.stats[connString,<user>,<apikey>]

 Aggregated and per OSD statistics.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, apikey** - the user name and API key for connecting to Ceph.

ceph.osd.discovery[connString,<user>,<apikey>]

 The list of discovered OSDs. Used for low-level discovery.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, apikey** - the user name and API key for connecting to Ceph.

ceph.osd.dump[connString,<user>,<apikey>]

 The usage thresholds and statuses of OSDs.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, apikey** - the user name and API key for connecting to Ceph.

ceph.ping[connString,<user>,<apikey>]

 Tests whether a connection to Ceph can be established.
 Return value: *0* - connection is broken (if there is any error presented including AUTH and configuration issues); *1* - connection is successful.

Parameters:

- **connString** - the URI or session name;
- **user, apikey** - the user name and API key for connecting to Ceph.

ceph.pool.discovery[connString,<user>,<apikey>]

 The list of discovered pools. Used for low-level discovery.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, apikey** - the user name and API key for connecting to Ceph.

ceph.status[connString,<user>,<apikey>]

 The overall cluster's status.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, apikey** - the user name and API key for connecting to Ceph.

docker.container_info[<ID>,<info>]

 Low-level information about a container.
 Return value: The output of the ContainerInspect API call serialized as JSON.

Parameters:

- **ID** - the ID or name of the container;
- **info** - the amount of information returned. Supported values: *short* (default) or *full*.

The Agent 2 user (’zabbix’) must be added to the ’docker’ group for sufficient privileges. Otherwise, the check will fail.

docker.container_stats[<ID>]

 The container resource usage statistics.
 Return value: The output of the ContainerStats API call and CPU usage percentage serialized as JSON.

Parameter:

- **ID** - the ID or name of the container.

The Agent 2 user (’zabbix’) must be added to the ’docker’ group for sufficient privileges. Otherwise, the check will fail.

docker.containers

 The list of containers.
 Return value: The output of the ContainerList API call serialized as JSON.

The Agent 2 user (’zabbix’) must be added to the ’docker’ group for sufficient privileges. Otherwise, the check will fail.

docker.containers.discovery[<options>]

 Returns the list of containers. Used for low-level discovery.
 Return value: *JSON object*.

Parameter:

- **options** - specify whether all or only running containers should be discovered. Supported values: *true* - return all containers; *false* - return only running containers (default).

The Agent 2 user (’zabbix’) must be added to the ’docker’ group for sufficient privileges. Otherwise, the check will fail.

docker.data.usage

 Information about the current data usage.
 Return value: The output of the SystemDataUsage API call serialized as JSON.

The Agent 2 user (’zabbix’) must be added to the ’docker’ group for sufficient privileges. Otherwise, the check will fail.

docker.images

 Returns the list of images.
 Return value: The output of the ImageList API call serialized as JSON.

The Agent 2 user (’zabbix’) must be added to the ’docker’ group for sufficient privileges. Otherwise, the check will fail.

docker.images.discovery

 Returns the list of images. Used for low-level discovery.
 Return value: *JSON object*.

The Agent 2 user (’zabbix’) must be added to the ’docker’ group for sufficient privileges. Otherwise, the check will fail.

docker.info

 The system information.
 Return value: The output of the SystemInfo API call serialized as JSON.

The Agent 2 user (’zabbix’) must be added to the ’docker’ group for sufficient privileges. Otherwise, the check will fail.

docker.ping

 Test if the Docker daemon is alive or not.
 Return value: *1* - the connection is alive; *0* - the connection is broken.

The Agent 2 user (’zabbix’) must be added to the ’docker’ group for sufficient privileges. Otherwise, the check will fail.

ember.get[<uri>,<path>]

 Returns the result of the required device.
 Return value: *JSON object*.

Parameters:

- **uri** - Ember+ device URI. Default: 127.0.0.1:9998;
- **path** - OID path to device. Empty by default, returns root collection data.

memcached.ping[connString,<user>,<password>]

 Test if a connection is alive or not.
 Return value: *1* - the connection is alive; *0* - the connection is broken (if there is any error presented including AUTH and configuration issues).

Parameters:

- **connString** - the URI or session name;
- **user, password** - the Memcached login credentials.

memcached.stats[connString,<user>,<password>,<type>]

 Gets the output of the STATS command.
 Return value: *JSON* - the output is serialized as JSON.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the Memcached login credentials;
- **type** - stat type to be returned: *items*, *sizes*, *slabs* or *settings* (empty by default, returns general statistics).

mongodb.collection.stats[connString,<user>,<password>,<database>,collection]

 Returns a variety of storage statistics for a given collection.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials;
- **database** - the database name (default: admin);
- **collection** - the collection name.

mongodb.cfg.discovery[connString,<user>,<password>]

 Returns a list of discovered configuration servers.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mongodb.collections.discovery[connString,<user>,<password>]

 Returns a list of discovered collections. Used for low-level discovery.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mongodb.collections.usage[connString,<user>,<password>]

 Returns the usage statistics for collections.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mongodb.connpool.stats[connString,<user>,<password>]

 Returns information regarding the open outgoing connections from the current database instance to other members of the sharded cluster or replica set.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mongodb.db.stats[connString,<user>,<password>,<database>]

 Returns the statistics reflecting a given database system state.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials;
- **database** - the database name (default: admin).

mongodb.db.discovery[connString,<user>,<password>]

 Returns a list of discovered databases. Used for low-level discovery.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mongodb.jumbo_chunks.count[connString,<user>,<password>]

 Returns the count of jumbo chunks.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mongodb.oplog.stats[connString,<user>,<password>]

 Returns the status of the replica set, using data polled from the oplog.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mongodb.ping[connString,<user>,<password>]

 Test if a connection is alive or not.
 Return value: *1* - the connection is alive; *0* - the connection is broken (if there is any error presented including AUTH and configuration issues).

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mongodb.rs.config[connString,<user>,<password>]

 Returns the current configuration of the replica set.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mongodb.rs.status[connString,<user>,<password>]

 Returns the replica set status from the point of view of the member where the method is run.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mongodb.server.status[connString,<user>,<password>]

 Returns the database state.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mongodb.sh.discovery[connString,<user>,<password>]

 Returns the list of discovered shards present in the cluster.
 Return value: *JSON object*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mongodb.version[connString,<user>,<password>]

 Returns the database server version.
 Return value: *String*.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MongoDB login credentials.

mqtt.get[<broker url>,topic,<user>,<password>]

 Subscribes to a specific topic or topics (with wildcards) of the provided broker and waits for publications.
 Return value: Depending on topic content. If wildcards are used, returns topic content as JSON.

Parameters:

- **broker url** - the MQTT broker URL in the format `protocol://host:port` without query parameters (supported protocols: `tcp`, `ssl`, `ws`). If no value is specified, the agent will use `tcp://localhost:1883`. If a protocol or port are omitted, default protocol (`tcp`) or port (1883) will be used;
- **topic** - the MQTT topic (mandatory). Wildcards (+,#) are supported;
- **user, password** - the authentication credentials (if required).

Comments:

- The item must be configured as an active check ('Zabbix agent (active)' item type);
- TLS encryption certificates can be used by saving them into a default location (e.g., `/etc/ssl/certs/` directory for Ubuntu). For TLS, use the `tls://` scheme.

mssql.availability.group.get[URI,<user>,<password>]

 Returns availability groups.
 Return value: *JSON object*.

Parameters:

- **URI** - MSSQL server URI (the only supported schema is `sqlserver://`). Embedded credentials will be ignored. It is possible to specify an instance name as part of the URI, i.e.: `sqlserver://localhost/InstanceName` (no port). If port is specified, the instance name is ignored;
- **user, password** - username, password to send to protected MSSQL server.

For more information, see the MSSQL plugin readme.

mssql.custom.query[URI,<user>,<password>,queryName,<args...>]

 Returns the result of a custom query.
 Return value: *JSON object*.

Parameters:

- **URI** - MSSQL server URI (the only supported schema is `sqlserver://`). Embedded credentials will be ignored. It is possible to specify an instance name as part of the URI, i.e.: `sqlserver://localhost/InstanceName` (no port). If port is specified, the instance name is ignored;
- **user, password** - username, password to send to protected MSSQL server;
- **queryName** - name of a custom query configured in `Plugins.MSSQL.CustomQueriesDir` without the .sql extension;
- **args** - one or several comma-separated arguments to pass to a query.

For more information, see the MSSQL plugin readme.

mssql.db.get

 Returns all available MSSQL databases.
 Return value: *JSON object*.

For more information, see the MSSQL plugin readme.

mssql.job.status.get

 Returns the status of jobs.
 Return value: *JSON object*.

For more information, see the MSSQL plugin readme.

mssql.last.backup.get

 Returns the last backup time for all databases.
 Return value: *JSON object*.

For more information, see the MSSQL plugin readme.

mssql.local.db.get

 Returns databases that are participating in an Always On availability group and replica (primary or secondary) and are located on the server that the connection was established to.
 Return value: *JSON object*.

For more information, see the MSSQL plugin readme.

mssql.mirroring.get

 Returns mirroring info.
 Return value: *JSON object*.

For more information, see the MSSQL plugin readme.

mssql.nonlocal.db.get

 Returns databases that are participating in an Always On availability group and replica (primary or secondary) located on other servers (the database is not local to the SQL Server instance that the connection was established to).
 Return value: *JSON object*.

For more information, see the MSSQL plugin readme.

mssql.perfcounter.get

 Returns the performance counters.
 Return value: *JSON object*.

For more information, see the MSSQL plugin readme.

mssql.ping

 Ping the database. Test if connection is correctly configured.
 Return value: *1* - alive, *0* - not alive.

For more information, see the MSSQL plugin readme.

mssql.quorum.get

 Returns the quorum info.
 Return value: *JSON object*.

For more information, see the MSSQL plugin readme.

mssql.quorum.member.get

 Returns the quorum members.
 Return value: *JSON object*.

For more information, see the MSSQL plugin readme.

mssql.replica.get

 Returns the replicas.
 Return value: *JSON object*.

For more information, see the MSSQL plugin readme.

mssql.version

 Returns the MSSQL version.
 Return value: *String*.

For more information, see the MSSQL plugin readme.

mysql.custom.query[connString,<user>,<password>,queryName,<args...>]

 Returns the result of a custom query.
 Return value: *JSON object*.

Parameters:

- **connString** - URI or session name;
- **user, password** - MySQL login credentials;
- **queryName** - name of a custom query, must match SQL file name without an extension;
- **args** - one or several comma-separated arguments to pass to a query.

For more information, see the MySQL plugin readme, the *Custom queries* section.

mysql.db.discovery[connString,<user>,<password>]

 Returns the list of MySQL databases. Used for low-level discovery.
 Return value: The result of the "show databases" SQL query in LLD JSON format.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MySQL login credentials.

mysql.db.size[connString,<user>,<password>,<database name>]

 The database size in bytes.
 Return value: Result of the "select coalesce(sum(data_length + index_length),0) as size from information_schema.tables where table_schema=?" SQL query for specific database in bytes.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MySQL login credentials;
- **database name** - the database name.

mysql.get_status_variables[connString,<user>,<password>]

 Values of the global status variables.
 Return value: Result of the "show global status" SQL query in JSON format.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MySQL login credentials.

mysql.ping[connString,<user>,<password>]

 Test if a connection is alive or not.
 Return value: *1* - the connection is alive; *0* - the connection is broken (if there is any error presented including AUTH and configuration issues).

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MySQL login credentials.

mysql.replication.discovery[connString,<user>,<password>]

 Returns the list of MySQL replications. Used for low-level discovery.
 Return value: the result of the `SHOW SLAVE STATUS` or `SHOW REPLICA STATUS` SQL query in LLD JSON format.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MySQL login credentials.

mysql.replication.get_slave_status[connString,<user>,<password>,<master host>]

 The replication status.
 Return value: result of the `SHOW SLAVE STATUS` or `SHOW REPLICA STATUS` SQL query in JSON format.

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MySQL login credentials;
- **master host** - the replication master host name. If none found, an error is returned. If this parameter is not specified, all hosts are returned.

mysql.version[connString,<user>,<password>]

 The MySQL version.
 Return value: *String* (with the MySQL instance version).

Parameters:

- **connString** - the URI or session name;
- **user, password** - the MySQL login credentials.

net.dns.get[<ip>,name,<type>,<timeout>,<count>,<protocol>,"<flags>"]

Performs a DNS query and returns detailed DNS record information.
 This item is an extended version of the `net.dns.record` Zabbix agent item with more record types and customizable flags supported.
 Return values: *JSON object*

Parameters:

- **ip** - the IP address of DNS server (leave empty for the default DNS server);
- **name** - the DNS name to query;
- **type** - the record type to be queried (default is *SOA*);
- **timeout** - the timeout for the request in seconds (default is 1 second);
- **count** - the number of tries for the request (default is 2);
- **protocol** - the protocol used to perform DNS queries: *udp* (default) or *tcp*;
- **flags** - one or more comma-separated arguments to pass to a query.

Comments:

- The possible values for `type` are: *A*, *NS*, *MD*, *MF*, *CNAME*, *SOA*, *MB*, *MG*, *MR*, *NULL*, *PTR*, *HINFO*, *MINFO*, *MX*, *TXT*, *RP*, *AFSDB*, *X25*, *ISDN*, *RT*, *NSAPPTR*, *SIG*, *KEY*, *PX*, *GPOS*, *AAAA*, *LOC*, *NXT*, *EID*, *NIMLOC*, *SRV*, *ATMA*, *NAPTR*, *KX*, *CERT*, *DNAME*, *OPT*, *APL*, *DS*, *SSHFP*, *IPSECKEY*, *RRSIG*, *NSEC*, *DNSKEY*, *DHCID*, *NSEC3*, *NSEC3PARAM*, *TLSA*, *SMIMEA*, *HIP*, *NINFO*, *RKEY*, *TALINK*, *CDS*, *CDNSKEY*, *OPENPGPKEY*, *CSYNC*, *ZONEMD*, *SVCB*, *HTTPS*, *SPF*, *UINFO*, *UID*, *GID*, *UNSPEC*, *NID*, *L32*, *L64*, *LP*, *EUI48*, *EUI64*, *URI*, *CAA*, *AVC*, *AMTRELAY*. Note that values must be in uppercase only; lowercase or mixed case values are not supported.
- For reverse DNS lookups (when `type` is set to *PTR*), you can provide the DNS name in both reversed and non-reversed format (see examples below). Note that when PTR record is requested, the DNS name is actually an IP address.
- The possible values for `flags` are:
  - *cdflag* or *nocdflag* (default);
  - *rdflag* (default) or *nordflag*;
  - *dnssec* or *nodnssec* (default);
  - *nsid* or *nonsid* (default);
  - *edns0* (default) or *noedns0*;
  - *aaflag* or *noaaflag* (default);

    - *adflag* or *noadflag* (default). The `flags` *dnssec* and *nsid* cannot be used together with *noedns0*, as both require *edns0*. Note that values must be in lowercase only; uppercase case or mixed case values are not supported.
- Internationalized domain names are not supported, please use IDNA encoded names instead.
- The output is an object containing DNS record information based on the parameters provided (see more details).

Examples:

```
net.dns.get[192.0.2.1,example.com,DNSKEY,3,3,tcp,"cdflag,rdflag,nsid"]

net.dns.get[,198.51.100.1,PTR,,,,"cdflag,rdflag,nsid"]
net.dns.get[,1.100.51.198.in-addr.arpa,PTR,,,,"cdflag,rdflag,nsid"]

net.dns.get[,2a00:1450:400f:800::200e,PTR,,,,"cdflag,rdflag,nsid"]
net.dns.get[,e.0.0.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.8.0.f.0.0.4.0.5.4.1.0.0.a.2.ip6.arpa,PTR,,,,"cdflag,rdfla
```

nvml.device.count

 The number of GPU devices.
 Return value: *Integer*.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.decoder.utilization[<deviceUUID>]

 GPU device decoder utilization as a percentage.
 Return value: *Integer*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.ecc.mode[<deviceUUID>]

 GPU device ECC mode information (current, pending).
 Return value: *JSON*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.encoder.stats.get[<deviceUUID>]

 GPU device encoder statistics.
 Return value: *JSON*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.encoder.utilization[<deviceUUID>]

 GPU device encoder utilization as a percentage.
 Return value: *Integer*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.energy.consumption[<deviceUUID>]

 GPU device total energy consumption in millijoules (mJ) since the driver was last reloaded.
 Return value: *Integer*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.errors.memory[<deviceUUID>]

 GPU device ECC memory error information (corrected, uncorrected).
 Return value: *JSON*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.errors.register[<deviceUUID>]

 GPU device ECC register error information (corrected, uncorrected).
 Return value: *JSON*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.fan.speed.avg[<deviceUUID>]

 GPU device average fan speed as a percentage of maximum speed.
 Return value: *Integer*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.get

 Discovered GPUs with UUID and name.
 Return value: *JSON*.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.graphics.frequency[<deviceUUID>]

 GPU device graphics clock speed in MHz.
 Return value: *Integer*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.memory.fb.get[<deviceUUID>]

 GPU device framebuffer memory statistics (total, reserved, free, used).
 Return value: *JSON*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.memory.bar1.get[<deviceUUID>]

 GPU device BAR1 memory statistics (total, free, used).
 Return value: *JSON*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.memory.frequency[<deviceUUID>]

 GPU device memory clock speed in MHz.
 Return value: *Integer*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.pci.utilization[<deviceUUID>]

 GPU device PCI utilization (transmit/receive throughput in KBps).
 Return value: *JSON*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.performance.state[<deviceUUID>]

 GPU device performance state.
 Return value: *Integer*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.power.limit[<deviceUUID>]

 GPU device power limit in milliwatts.
 Return value: *Integer*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.power.usage[<deviceUUID>]

 GPU device current power usage in milliwatts.
 Return value: *Integer*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.serial[<deviceUUID>]

 GPU device serial number.
 Return value: *String*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.sm.frequency[<deviceUUID>]

 GPU device streaming multiprocessor clock speed in MHz.
 Return value: *Integer*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.temperature[<deviceUUID>]

 GPU device temperature in Celsius.
 Return value: *Integer*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.utilization[<deviceUUID>]

 GPU device utilization statistics (GPU/memory utilization as a percentage).
 Return value: *JSON*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.device.video.frequency[<deviceUUID>]

 GPU device video clock speed in MHz.
 Return value: *Integer*.

Parameter:

- **deviceUUID** - GPU device UUID.

For more information, see the NVIDIA GPU plugin readme.

nvml.system.driver.version

 The NVIDIA driver version.
 Return value: *String*.

For more information, see the NVIDIA GPU plugin readme.

nvml.version

 The NVML library version.
 Return value: *String*.

For more information, see the NVIDIA GPU plugin readme.

oracle.diskgroups.stats[connString,<user>,<password>,<service>,<diskgroup>]

 Returns the Automatic Storage Management (ASM) disk groups statistics.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., `1521`) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
  - session name - if such a name is found in the plugin configuration file;
  - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
  - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name;
- **diskgroup** - the name of the ASM disk group to query.

oracle.diskgroups.discovery[connString,<user>,<password>,<service>]

 Returns the list of ASM disk groups. Used for low-level discovery.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., `1521`) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
  - session name - if such a name is found in the plugin configuration file;
  - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
  - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

oracle.archive.info[connString,<user>,<password>,<service>,<destination>]

 The archive logs statistics.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., `1521`) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
  - session name - if such a name is found in the plugin configuration file;
  - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
  - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name;
- **destination** - the name of the destination to query.

oracle.cdb.info[connString,<user>,<password>,<service>,<database>]

 The Container Databases (CDBs) information.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:

    - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
    - session name - if such a name is found in the plugin configuration file;
    - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
    - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name;
- **database** - the name of the database to query.

oracle.custom.query[connString,<user>,<password>,<service>,queryName,<args...>]

 The result of a custom query.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
    - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
    - session name - if such a name is found in the plugin configuration file;
    - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
    - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name;
- **queryName** - the name of a custom query, must match SQL file name without an extension;
- **args** - one or several comma-separated arguments to pass to the query.

Comments: - Returned data is automatically converted into JSON. - Avoid returning JSON directly from queries, as it will become corrupted when the plugin attempts to convert it into JSON again.

oracle.datafiles.stats[connString,<user>,<password>,<service>]

 Returns the data files statistics.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
    - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
    - session name - if such a name is found in the plugin configuration file;
    - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
    - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

oracle.db.discovery[connString,<user>,<password>,<service>]

 Returns the list of databases. Used for low-level discovery.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
    - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;

    - session name - if such a name is found in the plugin configuration file;
    - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
    - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

oracle.fra.stats[connString,<user>,<password>,<service>]

 Returns the Fast Recovery Area (FRA) statistics.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
    - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
    - session name - if such a name is found in the plugin configuration file;
    - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
    - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

oracle.instance.info[connString,<user>,<password>,<service>]

 The instance statistics.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
    - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
    - session name - if such a name is found in the plugin configuration file;
    - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
    - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

oracle.pdb.info[connString,<user>,<password>,<service>,<database>]

 The Pluggable Databases (PDBs) information.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
    - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
    - session name - if such a name is found in the plugin configuration file;
    - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
    - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name;

- **database** - the name of the database to query.

oracle.pdb.discovery[connString,<user>,<password>,<service>]

 Returns the list of PDBs. Used for low-level discovery.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
  - session name - if such a name is found in the plugin configuration file;
  - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
  - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

oracle.pga.stats[connString,<user>,<password>,<service>]

 Returns the Program Global Area (PGA) statistics.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
  - session name - if such a name is found in the plugin configuration file;
  - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
  - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

oracle.ping[connString,<user>,<password>,<service>]

 Test whether a connection to Oracle can be established.
 Return value: *1* - the connection is successful; *0* - the connection is broken (if there is any error presented including AUTH and configuration issues).

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
  - session name - if such a name is found in the plugin configuration file;
  - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
  - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

Examples:

```
oracle.ping[tcp://127.0.0.1:1521,ZABBIX_MON,zabbix,xe]
oracle.ping[localhost,ZABBIX_MON,zabbix,xe]
oracle.ping[zbx_tns_example,ZABBIX_MON,zabbix,xe]
oracle.ping["(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=localhost)(PORT=1521))(CONNECT_DATA=(SERVICE_NAME=x
```

oracle.proc.stats[connString,<user>,<password>,<service>]

 Returns the processes statistics.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
  - session name - if such a name is found in the plugin configuration file;
  - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
  - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

oracle.redolog.info[connString,<user>,<password>,<service>]

 The log file information from the control file.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
  - session name - if such a name is found in the plugin configuration file;
  - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
  - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

oracle.sga.stats[connString,<user>,<password>,<service>]

 Returns the System Global Area (SGA) statistics.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
  - session name - if such a name is found in the plugin configuration file;
  - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
  - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

oracle.sessions.stats[connString,<user>,<password>,<service>,<lockMaxTime>]

 Returns the sessions statistics.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;

    - session name - if such a name is found in the plugin configuration file;
    - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
    - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name;
- **lockMaxTime** - the maximum session lock duration in seconds to count the session as a prolongedly locked. Default: 600 seconds.

oracle.sys.metrics[connString,<user>,<password>,<service>,<duration>]

 Returns a set of system metric values.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
    - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
    - session name - if such a name is found in the plugin configuration file;
    - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
    - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name;
- **duration** - the capturing interval (in seconds) of system metric values. Possible values: *60* - long duration (default), *15* - short duration.

oracle.sys.params[connString,<user>,<password>,<service>]

 Returns a set of system parameter values.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
    - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
    - session name - if such a name is found in the plugin configuration file;
    - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
    - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

oracle.ts.stats[connString,<user>,<password>,<service>,<tablespace>,<type>,<conname>]

 Returns the tablespaces statistics.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
    - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
    - session name - if such a name is found in the plugin configuration file;
    - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
    - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.

- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name;
- **tablespace** - name of the tablespace to query. Default (if left empty and `type` is set):
  - "TEMP" (if `type` is set to "TEMPORARY");
  - "USERS" (if `type` is set to "PERMANENT").
- **type** - the type of the tablespace to query. Default (if `tablespace` is set): "PERMANENT".
- **conname** - name of the container for which the information is required.

If `tablespace`, `type`, or `conname` is omitted, the item will return tablespace statistics for all matching containers (including PDBs and CDB).

oracle.ts.discovery[connString,<user>,<password>,<service>]

 Returns a list of tablespaces. Used for low-level discovery.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
  - session name - if such a name is found in the plugin configuration file;
  - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
  - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

oracle.user.info[connString,<user>,<password>,<service>,<username>]

 Returns Oracle user information.
 Return value: *JSON object*.

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
  - session name - if such a name is found in the plugin configuration file;
  - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;
  - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name;
- **username** - the username for which the information is needed. Lowercase usernames are not supported. Default: current user.

oracle.version[connString,<user>,<password>,<service>]

 Returns the database server version.
 Return value: *String*.

Parameters:

- **connString** - it can be:
  - the URI - if no session with the given name is found, and `Plugins.Oracle.ResolveTNS` is set to false. If it contains a schema (e.g., "tcp://"), port (e.g., 1521) or both, the ResolveTNS option will not be taken into account, and it anyway will be treated as URI;
  - session name - if such a name is found in the plugin configuration file;
  - TNS key - if the ResolveTNS option is set to true, and none of the above conditions apply;

  - TNS value - if it starts with the open bracket "(" (leading spaces ignored). In this case, the ResolveTNS option is not taken into account.
- **user** - the Oracle username, supports appending one of the login options `as sysdba`, `as sysoper`, `as sysasm`, `as sysbackup`, `as sysdg`, `as syskm`, or `as sysrac` in the format `user as sysdba` (a login option is case-insensitive, must not contain a trailing space).
- **password** - the Oracle password;
- **service** - the Oracle service name.

pgsql.autovacuum.count[uri,<username>,<password>,<database name>]

 The number of autovacuum workers.
 Return value: *Integer*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.archive[uri,<username>,<password>,<database name>]

 The information about archived files.
 Return value: *JSON object*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.bgwriter[uri,<username>,<password>,<database name>]

 The combined number of checkpoints for the database cluster, broken down by checkpoint type.
 Return value: *JSON object*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.cache.hit[uri,<username>,<password>,<database name>]

 The PostgreSQL buffer cache hit rate.
 Return value: *Float*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.connections[uri,<username>,<password>,<database name>]

 Returns connections by type.
 Return value: *JSON object*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.custom.query[uri,<username>,<password>,<database name>,queryName,<args...>]

 Returns the result of a custom query.
 Return value: *JSON object*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of Plugins.PostgreSQL.Default.Database from postgresql.conf is used);
- **queryName** - the name of a custom query, must match the SQL file name without an extension;
- **args** - one or several comma-separated arguments to pass to a query.

pgsql.db.age[uri,<username>,<password>,<database name>]

 The age of the oldest FrozenXID of the database (number of transactions).
 Return value: *Integer*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.db.bloating_tables[uri,<username>,<password>,<database name>]

 The number of bloating tables per database.
 Return value: *Integer*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.db.discovery[uri,<username>,<password>,<database name>]

 The list of PostgreSQL databases. Used for low-level discovery.
 Return value: *JSON object*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.db.size[uri,<username>,<password>,<database name>]

 The database size in bytes.
 Return value: *Integer*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.dbstat[uri,<username>,<password>,<database name>]

 Collects the statistics per database. Used for low-level discovery.
 Return value: *JSON object*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.dbstat.sum[uri,<username>,<password>,<database name>]

 The summarized data for all databases in a cluster.
 Return value: *JSON object*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.locks[uri,<username>,<password>,<database name>]

 The information about granted locks per database. Used for low-level discovery.
 Return value: *JSON object*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.oldest.xid[uri,<username>,<password>,<database name>]

 The age of the oldest XID (number of transactions).
 Return value: *Integer*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.ping[uri,<username>,<password>,<database name>]

 Tests whether a connection is alive or not.
 Return value: *1* - the connection is alive; *0* - the connection is broken (if there is any error presented including AUTH and configuration issues).

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.queries[uri,<username>,<password>,<database name>,time period]

 Queries metrics by execution time.
 Return value: *JSON object*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used);
- **time period** - the execution time limit for the count of slow queries (must be a positive integer).

pgsql.replication.count[uri,<username>,<password>]

 The number of standby servers.
 Return value: *Integer*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials.

pgsql.replication.process[uri,<username>,<password>]

 The flush lag, write lag and replay lag per each sender process.
 Return value: *JSON object*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials.

pgsql.replication.process.discovery[uri,<username>,<password>]

 The replication process name discovery.
 Return value: *JSON object*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials.

pgsql.replication.recovery_role[uri,<username>,<password>]

 The recovery status.
 Return value: *0* - master mode; *1* - recovery is still in progress (standby mode).

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials.

pgsql.replication.status[uri,<username>,<password>]

 The status of replication.
 Return value: *0* - streaming is down; *1* - streaming is up; *2* - master mode.

Parameters:

- **uri** - the URI or session name;

- **username, password** - the PostgreSQL credentials.

pgsql.replication_lag.b[uri,<username>,<password>]

 The replication lag in bytes.
 Return value: *Integer*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials.

pgsql.replication_lag.sec[uri,<username>,<password>]

 The replication lag in seconds.
 Return value: *Integer*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials.

pgsql.uptime[uri,<username>,<password>,<database name>]

 The PostgreSQL uptime in milliseconds.
 Return value: *Float*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.version[uri,<username>,<password>,<database name>]

 Returns PostgreSQL version.
 Return value: *String*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

pgsql.wal.stat[uri,<username>,<password>,<database name>]

 The WAL statistics.
 Return value: *JSON object*.

Parameters:

- **uri** - the URI or session name;
- **username, password** - the PostgreSQL credentials;
- **database name** - the database name (if omitted, the value of `Plugins.PostgreSQL.Default.Database` from `postgresql.conf` is used).

redis.config[connString,<password>,<pattern>]

 Gets the configuration parameters of a Redis instance that match the pattern.
 Return value: *JSON* - if a glob-style pattern was used; single value - if a pattern did not contain any wildcard character.

Parameters:

- **connString** - the URI or session name;
- **password** - the Redis password;
- **pattern** - a glob-style pattern (* by default).

Comment:

- Security warning: `redis.config` runs the Redis `CONFIG GET` command and may return sensitive configuration parameters (for example, `requirepass`, `masterauth`, TLS-related settings, and other secrets). Using the default pattern (`*`) or wide patterns may expose these secrets to anyone with permission to view item values in Zabbix.

redis.info[connString,<password>,<section>]

 Gets the output of the INFO command.
 Return value: *JSON* - the output is serialized as JSON.

Parameters:

- **connString** - the URI or session name;
- **password** - the Redis password;

- **section** - the section of information (*default* by default).

redis.ping[connString,<password>]

 Test if a connection is alive or not.
 Return value: *1* - the connection is alive; *0* - the connection is broken (if there is any error presented including AUTH and configuration issues).

Parameters:

- **connString** - the URI or session name;
- **password** - the Redis password.

redis.slowlog.count[connString,<password>]

 The number of slow log entries since Redis was started.
 Return value: *Integer*.

Parameters:

- **connString** - the URI or session name;
- **password** - the Redis password.

smart.attribute.discovery

 Returns a list of S.M.A.R.T. device attributes.
 Return value: *JSON object*.

Comments:

- The following macros and their values are returned: `{#NAME}`, `{#DISKTYPE}`, `{#ID}`, `{#ATTRNAME}`, `{#THRESH}`;
- HDD, SSD and NVME drive types are supported. Drives can be alone or combined in a RAID. `{#NAME}` will have an add-on in case of RAID, e.g: `{"{#NAME}": "/dev/sda cciss,2"}`.

smart.disk.discovery[<type>]

 Returns a list of S.M.A.R.T. devices.
 Return value: *JSON object*.

Parameter:

- **type** - specifies a value to scan for the disks. Possible values: *id*, *name* (default). *id* is not supported on Windows, and will return an error if selected

Comments:

- The following macros and their values are returned: `{#NAME}`, `{#DISKTYPE}`, `{#MODEL}`, `{#SN}`, `{#PATH}`, `{#ATTRIBUTES}`, `{#RAIDTYPE}`;
- HDD, SSD and NVME drive types are supported. If a drive does not belong to a RAID, the `{#RAIDTYPE}` will be empty. `{#NAME}` will have an add-on in case of RAID, e.g: `{"{#NAME}": "/dev/sda cciss,2"}`.

smart.disk.get[<path>,<raid type>]

 Returns all available properties of S.M.A.R.T. devices.
 Return value: *JSON object*.

Parameters:

- **path** - the disk path, the `{#PATH}` macro may be used as a value;
- **raid_type** - the RAID type, the `{#RAID}` macro may be used as a value

Comments:

- HDD, SSD and NVME drive types are supported. Drives can be alone or combined in a RAID;
- The data includes smartctl version and call arguments, and additional fields:
disk_name - holds the name with the required add-ons for RAID discovery, e.g: `{"disk_name": "/dev/sda cciss,2"}`
disk_type - holds the disk type HDD, SSD, or NVME, e.g: `{"disk_type": "ssd"}`;
- If no parameters are specified, the item will return information about all disks.

systemd.unit.get[unit name,<interface>]

 Returns all properties of a systemd unit.
 Return value: *JSON object*.

Parameters:

- **unit name** - the unit name (you may want to use the `{#UNIT.NAME}` macro in item prototype to discover the name);
- **interface** - the unit interface type, possible values: *Unit* (default), *Service*, *Socket*, *Device*, *Mount*, *Automount*, *Swap*, *Target*, *Path*.

Comments:

- This item is supported on Linux platform only;
- LoadState, ActiveState and UnitFileState for Unit interface are returned as text and integer: `"ActiveState":{"state":1,"text":"a`

systemd.unit.info[unit name,<property>,<interface>]

 A systemd unit information.
 Return value: *String*.

Parameters:

- **unit name** - the unit name (you may want to use the `{#UNIT.NAME}` macro in item prototype to discover the name);
- **property** - unit property (e.g., ActiveState (default), LoadState, Description);
- **interface** - the unit interface type (e.g., Unit (default), Socket, Service).

Comments:

- This item is supported on Linux platform only;
- This item allows you to retrieve a specific property from specific type of interface as described in dbus API.

Examples:

```
systemd.unit.info["{#UNIT.NAME}"] #collect active state (active, reloading, inactive, failed, activating,
systemd.unit.info["{#UNIT.NAME}",LoadState] #collect load state info on discovered systemd units
systemd.unit.info[mysqld.service,Id] #retrieve the service technical name (mysqld.service)
systemd.unit.info[mysqld.service,Description] #retrieve the service description (MySQL Server)
systemd.unit.info[mysqld.service,ActiveEnterTimestamp] #retrieve the last time the service entered the act
systemd.unit.info[dbus.socket,NConnections,Socket] #collect the number of connections from this socket uni
```

systemd.unit.discovery[<type>]

 List of systemd units and their details. Used for low-level discovery.
 Return value: *JSON object*.

Parameter:

- **type** - possible values: *all*, *automount*, *device*, *mount*, *path*, *service* (default), *socket*, *swap*, *target*.

This item is supported on Linux platform only.

web.certificate.get[hostname,<port>,<address>]

 Validates the certificates and returns certificate details.
 Return value: *JSON object*.

Parameter:

- **hostname** - can be either IP or DNS.
May contain the URL scheme (*https* only), path (it will be ignored), and port.
If a port is provided in both the first and the second parameters, their values must match.
If address (the 3rd parameter) is specified, the hostname is only used for SNI and hostname verification;
- **port** - the port number (default is 443 for HTTPS);
- **address** - can be either IP or DNS. If specified, it will be used for connection, and hostname (the 1st parameter) will be used for SNI, and host verification. In case, the 1st parameter is an IP and the 3rd parameter is DNS, the 1st parameter will be used for connection, and the 3rd parameter will be used for SNI and host verification.

Comments:

- This item turns unsupported if the destination specified in the host configuration does not exist, is unavailable, or if the TLS handshake fails with any error except an invalid certificate;
- Currently, AIA (Authority Information Access) X.509 extension, CRLs and OCSP (including OCSP stapling), and Certificate Transparency are not supported;
- JSON response fields:
  - *x509*: contains the details of the X.509 certificate.
    - *version*: the X.509 version (e.g., 3).
    - *serial_number*: the serial number of the certificate.
    - *signature_algorithm*: the algorithm used to sign the certificate (e.g., `SHA256-RSA`).
    - *issuer*: the issuer of the certificate.
    - *not_before*: the start date of the certificate's validity.
    - *not_after*: the expiration date of the certificate.
    - *subject*: the subject of the certificate.
    - *public_key_algorithm*: the algorithm used for the public key (e.g., `RSA`).
    - *alternative_names*: subject alternative names (if present), otherwise, `null`.
  - result: contains the validation result.
    - *value*: the validation status (see possible values below).
    - *message*: detailed validation message (e.g., `"certificate verified successfully"`).
  - *sha1_fingerprint*: the SHA-1 fingerprint of the certificate.
  - *sha256_fingerprint*: the SHA-256 fingerprint of the certificate.
- The `$.result.value` field indicates the certificate validation result. Possible values include:
  - *valid* - the certificate is valid and trusted.

- *valid-but-self-signed* - the certificate is valid but self-signed, meaning its subject matches its issuer.
- *invalid* - the certificate is invalid due to an issue such as expiration, incorrect hostname, or an unknown signing authority.

Example:

```
web.certificate.get[example.com,443]
```

JSON Response:


```
{
  "x509": {
    "version": 3,
    "serial_number": "0ad893bafa68b0b7fb7a404f06ecaf9a",
    "signature_algorithm": "ECDSA-SHA384",
    "issuer": "CN=DigiCert Global G3 TLS ECC SHA384 2020 CA1,O=DigiCert Inc,C=US",
    "not_before": {
      "value": "Jan 15 00:00:00 2025 GMT",
      "timestamp": 1736899200
    },
    "not_after": {
      "value": "Jan 15 23:59:59 2026 GMT",
      "timestamp": 1768521599
    },
    "subject": "CN=*.example.com,O=Internet Corporation for Assigned Names and Numbers,L=Los Angeles,ST=Ca
    "public_key_algorithm": "ECDSA",
    "alternative_names": [
      "*.example.com",
      "example.com"
    ]
  },
  "result": {
    "value": "valid",
    "message": "certificate verified successfully"
  },
  "sha1_fingerprint": "310db7af4b2bc9040c8344701aca08d0c69381e3",
  "sha256_fingerprint": "455943cf819425761d1f950263ebf54755d8d684c25535943976f488bc79d23b"
}
```


Windows Zabbix agent

Overview

The Windows Zabbix agent items are presented in two lists:

- Shared items - the item keys that are shared with the UNIX Zabbix agent;
- Windows-specific items - the item keys that are supported **only** on Windows.

Note that all item keys supported by Zabbix agent on Windows are also supported by the new generation Zabbix agent 2. See the additional item keys that you can use with the agent 2 only.

See also: Minimum permissions for Windows items

Shared items

The table below lists Zabbix agent items that are supported on Windows and are shared with the UNIX Zabbix agent:

- The item key is a link to full details of the UNIX Zabbix agent item
- Windows-relevant item comments are included

Item key	Description	Item group
log	The monitoring of a log file. This item is not supported for Windows Event Log. The `persistent_dir` parameter is not supported on Windows.	Log monitoring
log.count	The count of matched lines in a monitored log file. This item is not supported for Windows Event Log. The `persistent_dir` parameter is not supported on Windows.	

Item key	Description	Item group
logrt	The monitoring of a log file that is rotated. This item is not supported for Windows Event Log. The `persistent_dir` parameter is not supported on Windows.	
logrt.count	The count of matched lines in a monitored log file that is rotated. This item is not supported for Windows Event Log. The `persistent_dir` parameter is not supported on Windows.	
modbus.get	Reads Modbus data.	Modbus
net.dns	Checks if the DNS service is up. The `ip`, `timeout` and `count` parameters are ignored on Windows unless using Zabbix agent 2.	Network
net.dns.perf	Checks the performance of a DNS service. The `ip`, `timeout` and `count` parameters are ignored on Windows unless using Zabbix agent 2.	
net.dns.record	Performs a DNS query. The `ip`, `timeout` and `count` parameters are ignored on Windows unless using Zabbix agent 2.	
net.if.discovery	The list of network interfaces. Some Windows versions (for example, Server 2008) might require the latest updates installed to support non-ASCII characters in interface names.	
net.if.in	The incoming traffic statistics on a network interface. On Windows, the item gets values from 64-bit counters if available. 64-bit interface statistic counters were introduced in Windows Vista and Windows Server 2008. If 64-bit counters are not available, the agent uses 32-bit counters. Multi-byte interface names on Windows are supported. You may obtain network interface descriptions on Windows with net.if.discovery or net.if.list items.	
net.if.out	The outgoing traffic statistics on a network interface. On Windows, the item gets values from 64-bit counters if available. 64-bit interface statistic counters were introduced in Windows Vista and Windows Server 2008. If 64-bit counters are not available, the agent uses 32-bit counters. Multi-byte interface names on Windows are supported. You may obtain network interface descriptions on Windows with net.if.discovery or net.if.list items.	
net.if.total	The sum of incoming and outgoing traffic statistics on a network interface. On Windows, the item gets values from 64-bit counters if available. 64-bit interface statistic counters were introduced in Windows Vista and Windows Server 2008. If 64-bit counters are not available, the agent uses 32-bit counters. You may obtain network interface descriptions on Windows with net.if.discovery or net.if.list items.	
net.tcp.listen	Checks if this TCP port is in LISTEN state.	
net.tcp.port	Checks if it is possible to make a TCP connection to the specified port.	
net.tcp.service	Checks if a service is running and accepting TCP connections. Checking of LDAP and HTTPS on Windows is only supported by Zabbix agent 2.	
net.tcp.service.perf	Checks the performance of a TCP service. Checking of LDAP and HTTPS on Windows is only supported by Zabbix agent 2.	
net.tcp.socket.count	Returns the number of TCP sockets that match parameters. This item is supported on Linux by Zabbix agent, but on Windows it is supported only by Zabbix agent 2 on 64-bit Windows.	
net.udp.service	Checks if a service is running and responding to UDP requests.	
net.udp.service.perf	Checks the performance of a UDP service.	
net.udp.socket.count	Returns the number of UDP sockets that match parameters. This item is supported on Linux by Zabbix agent, but on Windows it is supported only by Zabbix agent 2 on 64-bit Windows.	
proc.get	The list of OS processes and their parameters. The `cmdline` parameter is not supported on Windows.	Processes
proc.num	The number of processes. On Windows, only the `name` and `user` parameters are supported.	
system.cpu.discovery	The list of detected CPUs/CPU cores.	System
system.cpu.load	The CPU load. When a collector process is started on Zabbix agent, the following performance counters are initialized and later used for this item: `\System\Processor Queue Length`	

Item key	Description	Item group
system.cpu.num	The number of CPUs.	
system.cpu.util	The CPU utilization percentage. The value is acquired using the *Processor Time* performance counter. Note that since Windows 8 its Task Manager shows CPU utilization based on the *Processor Utility* performance counter, while in previous versions it was the *Processor Time* counter (see more details). *system* is the only `type` parameter supported on Windows.	
system.hostname	The system host name. The value is acquired by either `GetComputerName()` (for **netbios**), `GetComputerNameExA()` (for **fqdn**), or `gethostname()` (for **host**) functions on Windows. See also a more detailed description.	
system.localtime	The system time.	
system.run	Run the specified command on the host.	
system.sw.arch	The software architecture information.	
system.sw.os	The operating system information.	
system.sw.os.get	Detailed information about the operating system (version, type, distribution name, minor and major version, etc).	
system.swap.size	The swap space size in bytes or in percentage from total. The `pused` type parameter is supported on Linux by Zabbix agent, but on Windows it is supported only by Zabbix agent 2. Note that this key might report incorrect swap space size/percentage on virtualized (VMware ESXi, VirtualBox) Windows platforms. In this case you may use the `perf_counter[\700(_Total)\702]` key to obtain correct swap space percentage.	
system.uname	Identification of the system. On Windows the value for this item is obtained from `Win32_OperatingSystem` and `Win32_Processor` WMI classes. The OS name (including edition) might be translated to the user's display language. On some versions of Windows it contains trademark symbols and extra spaces.	
system.uptime	The system uptime in seconds.	
vfs.dir.count	The directory entry count. On Windows, directory symlinks are skipped and hard links are counted only once.	Virtual file systems
vfs.dir.get	The directory entry list. On Windows, directory symlinks are skipped and hard links are counted only once.	
vfs.dir.size	The directory size. On Windows any symlink is skipped and hard links are taken into account only once.	
vfs.file.cksum	The file checksum, calculated by the UNIX cksum algorithm.	
vfs.file.contents	Retrieving the contents of a file.	
vfs.file.exists	Checks if the file exists. On Windows the double quotes have to be backslash '\' escaped and the whole item key enclosed in double quotes when using the command line utility for calling `zabbix_get.exe` or `agent2`. Note that the item may turn unsupported on Windows if a directory is searched within a non-existing directory, e.g., `vfs.file.exists[C:\no\dir,dir]` (where 'no' does not exist).	
vfs.file.get	Returns information about a file. Supported file types on Windows: regular file, directory, symbolic link	
vfs.file.md5sum	The MD5 checksum of file.	
vfs.file.owner	Retrieves the owner of a file.	
vfs.file.regexp	Retrieve a string in the file.	
vfs.file.regmatch	Find a string in the file.	
vfs.file.size	The file size.	
vfs.file.time	The file time information. On Windows XP `vfs.file.time[file,change]` may be equal to `vfs.file.time[file,access]`.	
vfs.fs.discovery	The list of mounted filesystems with their type and mount options. The {#FSLABEL} macro is supported on Windows.	
vfs.fs.get	The list of mounted filesystems with their type, available disk space, inode statistics and mount options. The {#FSLABEL} macro is supported on Windows.	

Item key	Description	Item group
vfs.fs.size	The disk space in bytes or in percentage from total.	
vm.memory.size	The memory size in bytes or in percentage from total.	Virtual memory
web.page.get	Get the content of a web page.	Web monitoring
web.page.perf	The loading time of a full web page.	
web.page.regexp	Find a string on the web page.	
agent.hostmetadata	The agent host metadata.	Zabbix
agent.hostname	The agent host name.	
agent.ping	The agent availability check.	
agent.variant	The variant of Zabbix agent (Zabbix agent or Zabbix agent 2).	
agent.version	The version of Zabbix agent.	
zabbix.stats	Returns a set of Zabbix server or proxy internal metrics remotely.	
zabbix.stats (queue)	Returns the number of monitored items in the queue which are delayed on Zabbix server or proxy remotely.	

Windows-specific items

The table provides details on the item keys that are supported **only** by the Windows Zabbix agent.

Windows-specific items sometimes are an approximate counterpart of a similar agent item, for example, `proc_info`, supported on Windows, roughly corresponds to the `proc.mem` item, not supported on Windows.

The item key is a link to full item key details.

Item key	Description	Item group
eventlog	The Windows event log monitoring.	Log monitoring
eventlog.count	The count of lines in the Windows event log.	
net.if.list	The network interface list (includes interface type, status, IPv4 address, description).	Network
perf_counter	The value of any Windows performance counter.	Performance counters
perf_counter_en	The value of any Windows performance counter in English.	
perf_instance.discovery	The list of object instances of Windows performance counters.	
perf_instance_en.discovery	The list of object instances of Windows performance counters, discovered using the object names in English.	
proc_info	Various information about specific process(es).	Processes
registry.data	Return data for the specified value name in the Windows Registry key.	Registry
registry.get	The list of Windows Registry values or keys located at given key.	
service.discovery	The list of Windows services.	Services
service.info	Information about a service.	
services	The listing of services.	
vm.vmemory.size	The virtual memory size in bytes or in percentage from the total.	Virtual memory
wmi.get	Execute a WMI query and return the first selected object.	WMI
wmi.getall	Execute a WMI query and return the whole response.	

Item key details

Parameters without angle brackets are mandatory. Parameters marked with angle brackets **< >** are optional.

eventlog[name,<regexp>,<severity>,<source>,<eventid>,<maxlines>,<mode>]

 The event log monitoring.
 Return value: *Log*.

Parameters:

- **name** - the name of the event log channel (*Log Name* in the Event Viewer GUI);
- **regexp** - a regular expression describing the required pattern (case sensitive);
- **severity** - a regular expression describing severity (case insensitive). This parameter accepts a regular expression based on the following values: ”Information”, ”Warning”, ”Error”, ”Critical”, ”Verbose” (running on Windows Vista or newer).
- **source** - a regular expression describing the source identifier (case insensitive);

- **eventid** - a regular expression describing the event identifier(s) (case sensitive);
- **maxlines** - the maximum number of new lines per second the agent will send to Zabbix server or proxy. This parameter overrides the value of MaxLinesPerSecond in zabbix_agentd.conf.
- **mode** - possible values:
  - *all* (default);
  - *skip* - skip the processing of older data (affects only newly created items).

Comments:

- The item must be configured as an active check;
- The agent is unable to send in events from the "Forwarded events" log;
- Windows Eventing 6.0 is supported;
- Selecting a non-Log type of information for this item will lead to the loss of local timestamp, as well as log severity and source information;
- See also additional information on log monitoring.

Examples:

```
eventlog[Application]
eventlog[Microsoft-Windows-Application-Experience/Program-Compatibility-Assistant]
eventlog[Security,,"Failure Audit",,^(529|680)$]
eventlog[System,,"Warning|Error"]
eventlog[System,,,,^1$]
eventlog[Windows PowerShell,,,,,,skip]
eventlog[System,,,,@TWOSHORT] #here a custom regular expression named `TWOSHORT` is referenced (defined as
```

eventlog.count[name,<regexp>,<severity>,<source>,<eventid>,<maxproclines>,<mode>]

 The count of lines in the Windows event log.
 Return value: *Integer*.

Parameters:

- **name** - the name of the event log channel (*Log Name* in the Event Viewer GUI);
- **regexp** - a regular expression describing the required pattern (case sensitive);
- **severity** - a regular expression describing severity (case insensitive). This parameter accepts a regular expression based on the following values: "Information", "Warning", "Error", "Critical", "Verbose" (running on Windows Vista or newer).
- **source** - a regular expression describing the source identifier (case insensitive);
- **eventid** - a regular expression describing the event identifier(s) (case sensitive);
- **maxproclines** - the maximum number of new lines per second the agent will analyze (cannot exceed 10000). The default value is 10*MaxLinesPerSecond in zabbix_agentd.conf.
- **mode** - possible values:
  - *all* (default);
  - *skip* - skip the processing of older data (affects only newly created items).

Comments:

- The item must be configured as an active check;
- The agent is unable to send in events from the "Forwarded events" log;
- Windows Eventing 6.0 is supported;
- Selecting a non-Log type of information for this item will lead to the loss of local timestamp, as well as log severity and source information;
- See also additional information on log monitoring.

Examples:

```
eventlog.count[System,,"Warning|Error"]
eventlog.count[Windows PowerShell,,,,,,skip]
```

net.if.list

 The network interface list (includes interface type, status, IPv4 address, description).
 Return value: *Text*.

Comments:

- Multi-byte interface names supported;
- Disabled interfaces are not listed;
- Enabling/disabling some components may change their ordering in the Windows interface name;
- Some Windows versions (for example, Server 2008) might require the latest updates installed to support non-ASCII characters in interface names.

perf_counter[counter,<interval>]

 The value of any Windows performance counter.
 Return value: *Integer*, *float*, *string* or *text* (depending on the request).

Parameters:

- **counter** - the path to the counter;
- **interval** - the last N seconds for storing the average value. The `interval` must be between 1 and 900 seconds (included) and the default value is 1.

Comments:

- This item follows UNC paths, which can be a security risk;
- `interval` is used for counters that require more than one sample (like CPU utilization), so the check returns an average value for last "interval" seconds every time;
- Performance Monitor can be used to obtain the list of available counters.
- See also: Windows performance counters.

perf_counter_en[counter,<interval>]

 The value of any Windows performance counter in English.
 Return value: *Integer*, *float*, *string* or *text* (depending on the request).

Parameters:

- **counter** - the path to the counter in English;
- **interval** - the last N seconds for storing the average value. The `interval` must be between 1 and 900 seconds (included) and the default value is 1.

Comments:

- This item follows UNC paths, which can be a security risk;
- `interval` is used for counters that require more than one sample (like CPU utilization), so the check returns an average value for last "interval" seconds every time;
- This item is only supported on **Windows Server 2008/Vista** and above;
- You can find the list of English strings by viewing the following registry key: `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib\009`.

perf_instance.discovery[object]

 The list of object instances of Windows performance counters. Used for low-level discovery.
 Return value: *JSON object*.

Parameter:

- **object** - the object name (localized).

perf_instance_en.discovery[object]

 The list of object instances of Windows performance counters, discovered using the object names in English. Used for low-level discovery.
 Return value: *JSON object*.

Parameter:

- **object** - the object name (in English).

proc_info[process,<attribute>,<type>]

 Various information about specific process(es).
 Return value: *Float*.

Parameters:

- **process** - the process name;
- **attribute** - the requested process attribute;
- **type** - the representation type (meaningful when more than one process with the same name exists)

Comments:

- The following `attributes` are supported:
  - *vmsize* (default) - size of process virtual memory in Kbytes;
  - *wkset* - size of process working set (amount of physical memory used by process) in Kbytes;
  - *pf* - number of page faults;
  - *ktime* - process kernel time in milliseconds;
  - *utime* - process user time in milliseconds;
  - *io_read_b* - number of bytes read by process during I/O operations;
  - *io_read_op* - number of read operation performed by process;
  - *io_write_b* - number of bytes written by process during I/O operations;
  - *io_write_op* - number of write operation performed by process;

    - *io_other_b* - number of bytes transferred by process during operations other than read and write operations;
    - *io_other_op* - number of I/O operations performed by process, other than read and write operations;
    - *gdiobj* - number of GDI objects used by process;
    - *userobj* - number of USER objects used by process.
- Valid `types` are:
    - *avg* (default) - average value for all processes named <process>;
    - *min* - minimum value among all processes named <process>;
    - *max* - maximum value among all processes named <process>;
    - *sum* - sum of values for all processes named <process>.
- On a 64-bit system, a 64-bit Zabbix agent is required for this item to work correctly.

Examples:

```
proc_info[iexplore.exe,wkset,sum] #retrieve the amount of physical memory taken by all Internet Explorer p
proc_info[iexplore.exe,pf,avg] #retrieve the average number of page faults for Internet Explorer processes
```

registry.data[key,<value name>]

 Return data for the specified value name in the Windows Registry key.
 Return value: *Integer*, *string* or *text* (depending on the value type)

Parameters:

- **key** - the registry key including the root key; root abbreviations (e.g., HKLM) are allowed;
- **value name** - the registry value name in the key (empty string "" by default). The default value is returned if the value name is not supplied.

Comments:

- Supported root abbreviations:
    - HKCR - HKEY_CLASSES_ROOT
    - HKCC - HKEY_CURRENT_CONFIG
    - HKCU - HKEY_CURRENT_USER
    - HKCULS - HKEY_CURRENT_USER_LOCAL_SETTINGS
    - HKLM - HKEY_LOCAL_MACHINE
    - HKPD - HKEY_PERFORMANCE_DATA
    - HKPN - HKEY_PERFORMANCE_NLSTEXT
    - HKPT - HKEY_PERFORMANCE_TEXT
    - HKU - HKEY_USERS
- Keys with spaces must be double-quoted.

Examples:

```
registry.data["HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\Windows Error Reporting"] #return the data of
registry.data["HKLM\SOFTWARE\Microsoft\Windows\Windows Error Reporting","EnableZip"] #return the data of t
```

registry.get[key,<mode>,<name regexp>]

 The list of Windows Registry values or keys located at given key.
 Return value: *JSON object*.

Parameters:

- **key** - the registry key including the root key; root abbreviations (e.g., HKLM) are allowed (see comments for `registry.data[]` to see full list of abbreviations);
- **mode** - possible values:
values (default) or *keys*;
- **name regexp** - only discover values with names that match the regexp (default - discover all values). Allowed only with *values* as `mode`.

Keys with spaces must be double-quoted.

Examples:

```
registry.get[HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall,values,"^DisplayName|DisplayVersion$
The JSON will include details of the key, last subkey, value name, value type and value data.
registry.get[HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall,values] #return the data of the all
The JSON will include details of the key, last subkey, value name, value type and value data.
registry.get[HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall,keys] #return all subkeys of this ke
The JSON will include details of the key and last subkey.
```

service.discovery

 The list of Windows services. Used for low-level discovery.
 Return value: *JSON object*.

service.info[service,<param>]

 Information about a service.
 Return value: - *Integer* - with `param` as *state*, *startup*;

 Specifically for *state*: - 0 - running - 1 - paused - 2 - start pending - 3 - pause pending - 4 - continue pending - 5 - stop pending - 6 - stopped - 7 - unknown - 255 - no such service

 Specifically for *startup*: - 0 - automatic - 1 - automatic delayed - 2 - manual - 3 - disabled - 4 - unknown - 5 - automatic trigger start - 6 - automatic delayed trigger start - 7 - manual trigger start - *String* - with `param` as *displayname*, *path*, *user*; - *Text* - with `param` as *description*.

Parameters:

- **service** - a real service name or its display name as seen in the MMC Services snap-in;
- **param** - *state* (default), *displayname*, *path*, *user*, *startup*, or *description*.

Comments:

- Items like `service.info[service,state]` and `service.info[service]` will return the same information;
- Only with `param` as *state* this item returns a value for non-existing services (255).

Examples:

```
service.info[SNMPTRAP] - state of the SNMPTRAP service;
service.info[SNMP Trap] - state of the same service, but with the display name specified;
service.info[EventLog,startup] - the startup type of the EventLog service
```

services[<type>,<state>,<exclude>]

 The listing of services.
 Return value: *0* - if empty; *Text* - the list of services separated by a newline.

Parameters:

- **type** - *all* (default), *automatic*, *manual*, or *disabled*;
- **state** - *all* (default), *stopped*, *started*, *start_pending*, *stop_pending*, *running*, *continue_pending*, *pause_pending*, or *paused*;
- **exclude** - the services to exclude from the result. Excluded services should be listed in double quotes, separated by comma, without spaces.

Examples:

```
services[,started] #returns the list of started services;
services[automatic, stopped] #returns the list of stopped services that should be running;
services[automatic, stopped, "service1,service2,service3"] #returns the list of stopped services that shou
```

vm.vmemory.size[<type>]

 The virtual memory size in bytes or in percentage from the total.
 Return value: *Integer* - for bytes; *float* - for percentage.

Parameter:

- **type** - possible values: *available* (available virtual memory), *pavailable* (available virtual memory, in percent), *pused* (used virtual memory, in percent), *total* (total virtual memory, default), or *used* (used virtual memory)

Comments:

- The monitoring of virtual memory statistics is based on:
  - Total virtual memory on Windows (total physical + page file size);
  - The maximum amount of memory Zabbix agent can commit;
  - The current committed memory limit for the system or Zabbix agent, whichever is smaller.

Example:

```
vm.vmemory.size[pavailable] #return the available virtual memory, in percentage
```

wmi.get[<namespace>,<query>]

 Execute a WMI query and return the first selected object.
 Return value: *Integer*, *float*, *string* or *text* (depending on the request).

Parameters:

- **namespace** - the WMI namespace;
- **query** - the WMI query returning a single object.

WMI queries are performed with WQL.

Example:

```
wmi.get[root\cimv2,select status from Win32_DiskDrive where Name like '%PHYSICALDRIVE0%'] #returns the sta
```

wmi.getall[<namespace>,<query>]

 Execute a WMI query and return the whole response. Can be used for low-level discovery.
 Return value: *JSON object*

Parameters:

- **namespace** - the WMI namespace;
- **query** - the WMI query.

Comments:

- WMI queries are performed with WQL.
- JSONPath preprocessing can be used to point to more specific values in the returned JSON.

Example:

```
wmi.getall[root\cimv2,select * from Win32_DiskDrive where Name like '%PHYSICALDRIVE%'] #returns status inf
```

Monitoring Windows services

This tutorial provides step-by-step instructions for setting up the monitoring of Windows services. It is assumed that Zabbix server and agent are configured and operational.

Step 1

Get the service name.

You can get the service name by going to the MMC Services snap-in and bringing up the properties of the service. In the *General* tab you should see a field called *Service name*. The value that follows is the name you will use when setting up an item for monitoring. For example, if you wanted to monitor the "workstation" service, then your service might be: **lanmanworkstation**.

Step 2

Configure an item for monitoring the service.

The item `service.info[service,<param>]` retrieves information about a particular service. Depending on the information you need, specify the `param` option which accepts the following values: *displayname*, *state*, *path*, *user*, *startup* or *description*. The default value is *state* if `param` is not specified (`service.info[service]`).

The type of return value depends on chosen `param`: integer for *state* and *startup*; character string for *displayname*, *path* and *user*; text for *description*.

Example:

- *Key:* `service.info[lanmanworkstation]`
- *Type of information:* Numeric (unsigned)

The item `service.info[lanmanworkstation]` will retrieve information about the state of the service as a numerical value. To map a numerical value to a text representation in the frontend ("0" as "Running", "1" as "Paused", etc.), you can configure value mapping on the host on which the item is configured. To do this, either link the template *Windows services by Zabbix agent* or *Windows services by Zabbix agent active* to the host, or configure on the host a new value map that is based on the *Windows service state* value map configured on the mentioned templates.

Note that both of the mentioned templates have a discovery rule configured that will discover services automatically. If you do not want this, you can disable the discovery rule on the host level once the template has been linked to the host.

Discovery of Windows services

Low-level discovery provides a way to automatically create items, triggers, and graphs for different entities on a computer. Zabbix can automatically start monitoring Windows services on your machine, without the need to know the exact name of a service or create items for each service manually. A filter can be used to generate real items, triggers, and graphs only for services of interest.

Log file monitoring

Overview

Zabbix can be used for centralized monitoring and analysis of log files with/without log rotation support.

Notifications can be used to warn users when a log file contains certain strings or string patterns.

To monitor a log file you must have:

- Zabbix agent running on the host.
- Log monitoring item set up.

Attention:
The size limit of a monitored log file depends on large file support.

Configuration

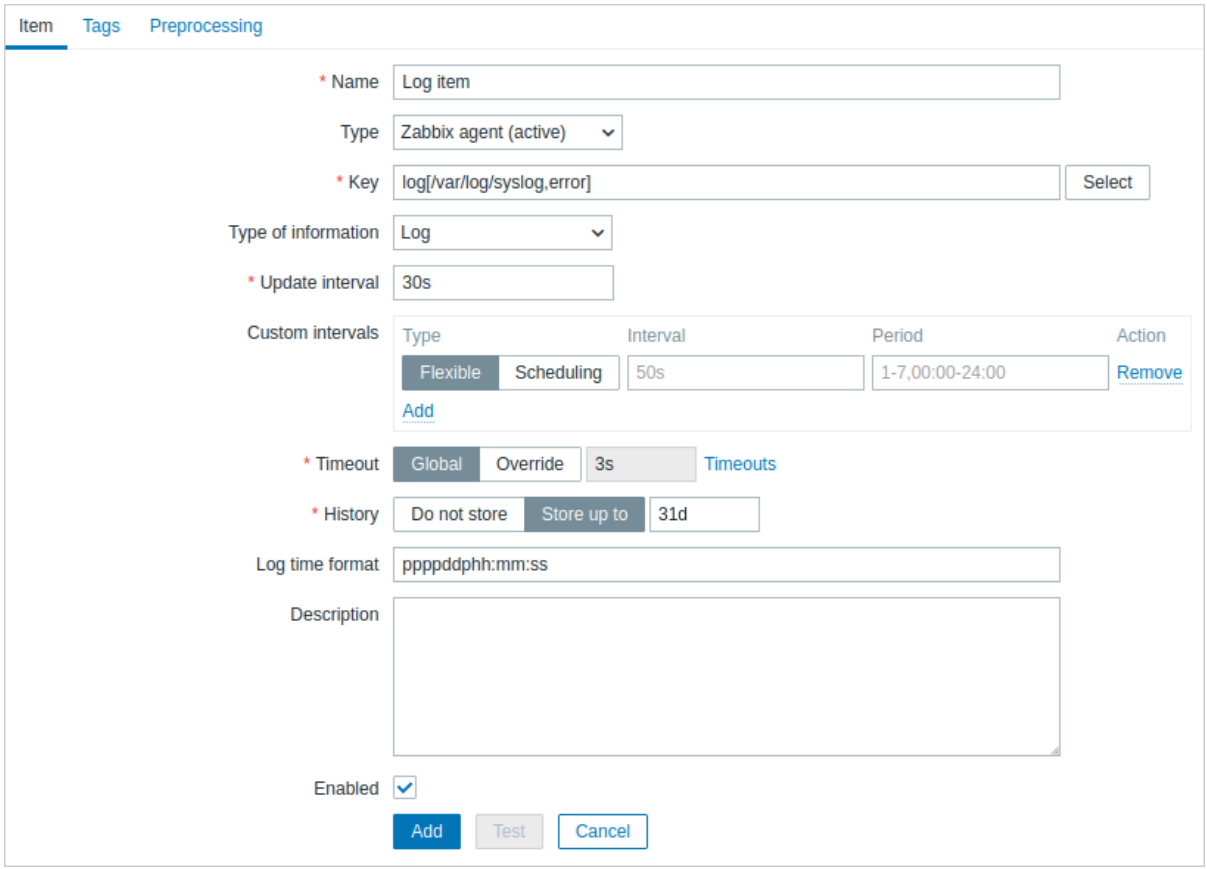
Verify agent parameters

Make sure that in the agent configuration file:

- `Hostname` parameter matches the host name in the frontend.
- Servers in the `ServerActive` parameter are specified for the processing of active checks.

Item configuration

Configure a log monitoring item.



All mandatory input fields are marked with a red asterisk.

Specifically for log monitoring items you enter:

Type	Select **Zabbix agent (active)** here.
Key	Use one of the following item keys: `log[]` or `logrt[]`: These two item keys allow you to monitor logs and filter log entries by the content regexp, if present. For example: `log[/var/log/syslog,error]`. Make sure that the file has read permissions for the 'zabbix' user; otherwise, the item status will be set to 'unsupported'. **log.count[]** or **logrt.count[]**: These two item keys allow you to return the number of matching lines only. See supported Zabbix agent item key section for details on using these item keys and their parameters.

Type of information	Prefilled automatically: For `log[]` or `logrt[]` items - *Log*; For `log.count[]` or `logrt.count[]` items - *Numeric (unsigned)*. If optionally using the `output` parameter, you may manually select the appropriate type of information other than *Log*. Note that choosing a non-Log type of information will lead to the loss of local timestamp.
Update interval (in sec)	The parameter defines how often Zabbix agent will check for any changes in the log file. Setting it to 1 second will make sure that you get new records as soon as possible.
Log time format	In this field you may optionally specify the pattern for parsing the log line timestamp. Supported placeholders: * **y**: *Year (1970-2038)* * **M**: *Month (01-12)* * **d**: *Day (01-31)* * **h**: *Hour (00-23)* * **m**: *Minute (00-59)* * **s**: *Second (00-59)* If left blank, the timestamp will be set to 0 in Unix time, representing `January 1, 1970`. For example, consider the following line from the Zabbix agent log file: `23480:20100328:154718.045 Zabbix agent started. Zabbix 1.8.2 (revision 11211).` It begins with six character positions for PID, followed by date, time, and the rest of the message. The log time format for this line would be "ppppppp:yyyyMMdd:hhmmss". Note that "p" and ":" characters are placeholders and can be any character except "yMdhms".

Important notes

- The server and agent keep the trace of a monitored log's size and last modification time (for `logrt`) in two counters. Additionally:
  - The agent also internally uses inode numbers (on UNIX/GNU/Linux), file indexes (on Microsoft Windows) and MD5 sums of the first 512 log file bytes for improving decisions when logfiles get truncated and rotated.
  - On UNIX/GNU/Linux systems it is assumed that the file systems where log files are stored report inode numbers, which can be used to track files.
  - On Microsoft Windows Zabbix agent determines the file system type the log files reside on and uses:
    - On NTFS file systems 64-bit file indexes.
    - On ReFS file systems (only from Microsoft Windows Server 2012) 128-bit file IDs.
    - On file systems where file indexes change (e.g., FAT32, exFAT) a fall-back algorithm is used to take a sensible approach in uncertain conditions when log file rotation results in multiple log files with the same last modification time.
  - The inode numbers, file indexes and MD5 sums are internally collected by Zabbix agent. They are not transmitted to Zabbix server and are lost when Zabbix agent is stopped.
  - Do not modify a log file's last-modified time (for example, with `touch`), and do not replace a monitored log file by copying a file back to its original name (this creates a new inode). In either case Zabbix may treat the file as a different file and re-read it from the start, which can produce duplicate alerts.
  - If there are several matching log files for `logrt[]` item and Zabbix agent is following the most recent of them and this most recent log file is deleted, a warning message `"there are no files matching "<regexp mask>" in "<directory>"` is logged. Zabbix agent ignores log files with modification time less than the most recent modification time seen by the agent for the `logrt[]` item being checked.
- The agent starts reading the log file from the point it stopped the previous time.
- The number of bytes already analyzed (the size counter) and last modification time (the time counter) are stored in the Zabbix database and are sent to the agent to make sure the agent starts reading the log file from this point in cases when the agent is just started or has received items which were previously disabled or not supported. However, if the agent has received a non-zero size counter from server, but the `logrt[]` or `logrt.count[]` item is unable to find matching files, the size counter is reset to 0 to analyze from the start if the files appear later.
- Whenever the log file becomes smaller than the log size counter known by the agent, the counter is reset to zero and the agent starts reading the log file from the beginning taking the time counter into account.
- If there are several matching files with the same last modification time in the directory, then the agent tries to correctly analyze all log files with the same modification time and avoid skipping data or analyzing the same data twice, although it cannot be guaranteed in all situations. The agent does not assume any particular log file rotation scheme nor determines one. When presented multiple log files with the same last modification time, the agent will process them in a lexicographically descending order. Thus, for some rotation schemes the log files will be analyzed and reported in their original order. For other rotation schemes the original log file order will not be honored, which can lead to reporting matched log file records in altered order (the problem does not happen if log files have different last modification times).

- Zabbix agent processes new records of a log file once per *Update interval* seconds.
- Zabbix agent does not send more than `maxlines` of a log file per second. The limit prevents overloading of network and CPU resources and overrides the default value provided by `MaxLinesPerSecond` parameter in the agent configuration file.
- To find the required string Zabbix will process 10 times more new lines than set in `MaxLinesPerSecond`. Thus, for example, if a `log[]` or `logrt[]` item has *Update interval* of 1 second, by default the agent will analyze no more than 200 log file records and will send no more than 20 matching records to Zabbix server in one check. By increasing `MaxLinesPerSecond` in the agent configuration file or setting `maxlines` parameter in the item key, the limit can be increased up to 10000 analyzed log file records and 1000 matching records sent to Zabbix server in one check. If the *Update interval* is set to 2 seconds the limits for one check would be set 2 times higher than with *Update interval* of 1 second.
- Additionally, `log` and `log.count` values are always limited to 50% of the agent send buffer size, even if there are no non-log values in it. So for the `maxlines` values to be sent in one connection (and not in several connections), the agent `BufferSize` parameter must be at least `maxlines` x 2. Zabbix agent can upload data during log gathering and thus free the buffer, whereas Zabbix agent 2 will stop log gathering until the data is uploaded and the buffer is freed, which is performed asynchronously.
- In the absence of log items all agent buffer size is used for non-log values. When log values come in they replace the older non-log values as needed, up to the designated 50%.
- For log file records longer than 256kB, only the first 256kB are matched against the regular expression and the rest of the record is ignored. However, if Zabbix agent is stopped while it is dealing with a long record the agent internal state is lost and the long record may be analyzed again and differently after the agent is started again.
- Special note for `\` path separators: if `file_format` is `file\.log`, then there should not be a `file` directory, since it is not possible to unambiguously define whether `.` is escaped or is the first symbol of the file name.
- Regular expressions for `logrt` are supported in filename only, directory regular expression matching is not supported.
- On UNIX platforms a `logrt[]` item becomes NOTSUPPORTED if a directory where the log files are expected to be found does not exist.
- On Microsoft Windows, if a directory does not exist the item will not become NOTSUPPORTED (for example, if directory is misspelled in item key).
- An absence of log files for `logrt[]` item does not make it NOTSUPPORTED. Errors of reading log files for `logrt[]` item are logged as warnings into Zabbix agent log file but do not make the item NOTSUPPORTED.
- Zabbix agent log file can be helpful to find out why a `log[]` or `logrt[]` item became NOTSUPPORTED. Zabbix can monitor its agent log file, except when at `DebugLevel=4` or `DebugLevel=5`.
- Searching for a question mark using a regular expression, e.g., `\?` may result in false positives if the text file contains NUL symbols, as those are replaced with `?` by Zabbix to continue processing the line until the newline character.

Extracting matching part of regular expression

Sometimes we may want to extract only the interesting value from a target file instead of returning the whole line when a regular expression match is found.

Log items have the ability to extract desired values from matched lines. This is accomplished by the additional **output** parameter in `log` and `logrt` items.

Using the 'output' parameter allows you to indicate the "capturing group" of the match that we may be interested in.

So, for example:

```
log[/path/to/the/file,"large result buffer allocation.*Entries: ([0-9]+)",,,,\1]
```

should allow returning the entry count as found in the content of:

```
Fr Feb 07 2014 11:07:36.6690 */ Thread Id 1400 (GLEWF) large result
buffer allocation - /Length: 437136/Entries: 5948/Client Ver: >=10/RPC
ID: 41726453/User: AUser/Form: CFG:ServiceLevelAgreement
```

Only the number will be returned because `\1` refers to the first and only capturing group: `([0-9]+)`.

And, with the ability to extract and return a number, the value can be used to define triggers.

Using maxdelay parameter

The `maxdelay` parameter in log items allows ignoring some older lines from log files in order to get the most recent lines analyzed within the `maxdelay` seconds.

> **Warning:**
> Specifying `maxdelay` > 0 may lead to **ignoring important log file records and missed alerts**. Use it carefully at your own risk only when necessary.

By default, items for log monitoring follow all new lines appearing in the log files. However, there are applications which in some situations start writing an enormous number of messages in their log files. For example, if a database or a DNS server is unavailable,

such applications flood log files with thousands of nearly identical error messages until normal operation is restored. By default, all those messages will be dutifully analyzed and matching lines sent to server as configured in `log` and `logrt` items.

Built-in protection against overload consists of a configurable `maxlines` parameter (protects server from too many incoming matching log lines) and a 10*`maxlines` limit (protects host CPU and I/O from overloading by agent in one check). Still, there are 2 problems with the built-in protection. First, a large number of potentially not-so-informative messages are reported to server and consume space in the database. Second, due to the limited number of lines analyzed per second the agent may lag behind the newest log records for hours. Quite likely, you might prefer to be sooner informed about the current situation in the log files instead of crawling through old records for hours.

The solution to both problems is using the `maxdelay` parameter. If `maxdelay` > 0 is specified, during each check the number of processed bytes, the number of remaining bytes and processing time is measured. From these numbers the agent calculates an estimated delay - how many seconds it would take to analyze all remaining records in a log file.

If the delay does not exceed `maxdelay` then the agent proceeds with analyzing the log file as usual.

If the delay is greater than `maxdelay` then the agent **ignores a chunk of a log file by "jumping" over it** to a new estimated position so that the remaining lines could be analyzed within `maxdelay` seconds.

Note that agent does not even read ignored lines into buffer, but calculates an approximate position to jump to in a file.

The fact of skipping log file lines is logged in the agent log file like this:

```
14287:20160602:174344.206 item:"logrt["/home/zabbix32/test[0-9].log",ERROR,,1000,,,120.0]"
logfile:"/home/zabbix32/test1.log" skipping 679858 bytes
(from byte 75653115 to byte 76332973) to meet maxdelay
```

The "to byte" number is approximate because after the "jump" the agent adjusts the position in the file to the beginning of a log line which may be further in the file or earlier.

Depending on how the speed of growing compares with the speed of analyzing the log file you may see no "jumps", rare or often "jumps", large or small "jumps", or even a small "jump" in every check. Fluctuations in the system load and network latency also affect the calculation of delay and hence, "jumping" ahead to keep up with the `maxdelay` parameter.

Setting `maxdelay` < `update interval` is not recommended (it may result in frequent small "jumps").

Notes on handling 'copytruncate' log file rotation

`logrt` with the `copytruncate` option assumes that different log files have different records (at least their timestamps are different), therefore MD5 sums of initial blocks (up to the first 512 bytes) will be different. Two files with the same MD5 sums of initial blocks means that one of them is the original, another - a copy.

`logrt` with the `copytruncate` option makes effort to correctly process log file copies without reporting duplicates. However, things like producing multiple log file copies with the same timestamp, log file rotation more often than `logrt[]` item update interval, frequent restarting of agent are not recommended. The agent tries to handle all these situations reasonably well, but good results cannot be guaranteed in all circumstances.

Notes on persistent files for log*[] items

Purpose of persistent files

When Zabbix agent is started it receives a list of active checks from Zabbix server or proxy. For `log*[]` metrics it receives the processed log size and the modification time for finding where to start log file monitoring from. Depending on the actual log file size and modification time reported by file system the agent decides either to continue log file monitoring from the processed log size or re-analyze the log file from the beginning.

A running agent maintains a larger set of attributes for tracking all monitored log files between checks. This in-memory state is lost when the agent is stopped.

The new optional parameter `persistent_dir` specifies a directory for storing this state of `log[]`, `log.count[]`, `logrt[]` or `logrt.count[]` item in a file. The state of log item is restored from the persistent file after the Zabbix agent is restarted.

The primary use-case is monitoring of log file located on a mirrored file system. Until some moment in time the log file is written to both mirrors. Then mirrors are split. On the active copy the log file is still growing, getting new records. Zabbix agent analyzes it and sends processed logs size and modification time to server. On the passive copy the log file stays the same, well behind the active copy. Later the operating system and Zabbix agent are rebooted from the passive copy. The processed log size and modification time the Zabbix agent receives from server may not be valid for situation on the passive copy. To continue log file monitoring from the place the agent left off at the moment of file system mirror split the agent restores its state from the persistent file.

Agent operation with persistent file

On startup Zabbix agent knows nothing about persistent files. Only after receiving a list of active checks from Zabbix server (proxy) the agent sees that some log items should be backed by persistent files under specified directories.

During agent operation the persistent files are opened for writing (with fopen(filename, "w")) and overwritten with the latest data. The chance of losing persistent file data if the overwriting and file system mirror split happen at the same time is very small, no special handling for it. Writing into persistent file is NOT followed by enforced synchronization to storage media (`fsync()` is not called).

Overwriting with the latest data is done after successful reporting of matching log file record or metadata (processed log size and modification time) to Zabbix server. That may happen as often as every item check if log file keeps changing.

No special actions during agent shutdown.

After receiving a list of active checks the agent marks obsolete persistent files for removal. A persistent file becomes obsolete if:

1. The corresponding log item is no longer monitored.
2. A log item is reconfigured with a different `persistent_dir` location than before.

Removing is done with delay 24 hours because log files in NOTSUPPORTED state are not included in the list of active checks but they may become SUPPORTED later and their persistent files will be useful.

If the agent is stopped before 24 hours expire, then the obsolete files will not be deleted as Zabbix agent is not getting info about their location from Zabbix server anymore.

Warning:
Reconfiguring a log item's `persistent_dir` back to the old `persistent_dir` location while the agent is stopped, without deleting the old persistent file by user - will cause restoring the agent state from the old persistent file resulting in missed messages or false alerts.

Naming and location of persistent files

Zabbix agent distinguishes active checks by their keys. For example, `logrt[/home/zabbix/test.log]` and `logrt[/home/zabbix/test` are different items. Modifying the item `logrt[/home/zabbix/test.log,,,10]` in frontend to `logrt[/home/zabbix/test.log,,,20]` will result in deleting the item `logrt[/home/zabbix/test.log,,,10]` from the agent's list of active checks and creating `logrt[/home/zabbix/test.log,,,20]` item (some attributes are carried across modification in frontend/server, not in agent).

The file name is composed of MD5 sum of item key with item key length appended to reduce possibility of collisions. For example, the state of `logrt[/home/zabbix50/test.log,,,,,,,,/home/zabbix50/agent_private]` item will be kept in persistent file c963ade4008054813bbc0a650bb8e09266.

Multiple log items can use the same value of `persistent_dir`.

`persistent_dir` is specified by taking into account specific file system layouts, mount points and mount options and storage mirroring configuration - the persistent file should be on the same mirrored filesystem as the monitored log file.

If the `persistent_dir` directory cannot be created or does not exist, or if the Zabbix agent does not have sufficient permissions to create, write, read, or delete files, the log item becomes `NOTSUPPORTED`.

If access rights to persistent storage files are removed during agent operation or other errors occur (e.g., disk full) then errors are logged into the agent log file but the log item does not become `NOTSUPPORTED`.

Load on I/O

Item's persistent file is updated after successful sending of every batch of data (containing item's data) to server. For example, default `BufferSize` is 100. If a log item has found 70 matching records then the first 50 records will be sent in one batch, persistent file will be updated, then remaining 20 records will be sent (maybe with some delay when more data is accumulated) in the 2nd batch, and the persistent file will be updated again.

Actions if communication fails between agent and server

Each matching line from `log[]` and `logrt[]` item and a result of each `log.count[]` and `logrt.count[]` item check requires a free slot in the designated 50% area in the agent send buffer. The buffer elements are regularly sent to server (or proxy) and the buffer slots are free again.

While there are free slots in the designated log area in the agent send buffer and communication fails between agent and server (or proxy) the log monitoring results are accumulated in the send buffer. This helps to mitigate short communication failures.

During longer communication failures all log slots get occupied and the following actions are taken:

- `log[]` and `logrt[]` item checks are stopped. When communication is restored and free slots in the buffer are available the checks are resumed from the previous position. No matching lines are lost, they are just reported later.

- `log.count[]` and `logrt.count[]` checks are stopped if `maxdelay = 0` (default). Behavior is similar to `log[]` and `logrt[]` items as described above. Note that this can affect `log.count[]` and `logrt.count[]` results: for example, one check counts 100 matching lines in a log file, but as there are no free slots in the buffer the check is stopped. When communication is restored the agent counts the same 100 matching lines and also 70 new matching lines. The agent now sends count = 170 as if they were found in one check.
- `log.count[]` and `logrt.count[]` checks with `maxdelay > 0`: if there was no "jump" during the check, then behavior is similar to described above. If a "jump" over log file lines took place then the position after "jump" is kept and the counted result is discarded. So, the agent tries to keep up with a growing log file even in case of communication failure.

Handling of regular expression compilation and runtime errors

If a regular expression used in `log[]`, `logrt[]`, `log.count[]` or `logrt.count[]` item cannot be compiled by PCRE or PCRE2 library then the item goes into NOTSUPPORTED state with an error message. To continue monitoring the log item, the regular expression should be fixed.

If the regular expression compiles successfully, but fails at runtime (on some or on all log records), then the log item remains supported and monitoring continues. The runtime error is logged in the Zabbix agent log file (without the log file record).

The logging rate is limited to one runtime error per check to allow Zabbix agent to monitor its own log file. For example, if 10 records are analyzed and 3 records fail with a regexp runtime error, one record is produced in the agent log.

Exception: if `MaxLinesPerSecond=1` and `update interval=1` (allowing only one record to be analyzed per check), regexp runtime errors are not logged.

zabbix_agentd logs the item key in case of a runtime error, zabbix_agent2 logs the item ID to help identify which log item has runtime errors. It is recommended to redesign the regular expression in case of runtime errors.

Simple check

Overview

Simple checks are normally used for remote agent-less checks of services.

Note that Zabbix agent is not needed for simple checks. Zabbix server/proxy is responsible for the processing of simple checks (making external connections, etc).

Examples of using simple checks:

```
net.tcp.service[ftp,,155]
net.tcp.service[http]
net.tcp.service.perf[http,,8080]
net.udp.service.perf[ntp]
```

Note:
User name and *Password* fields (limited to 255 characters) in simple check item configuration are used for VMware monitoring items; ignored otherwise.

Supported checks

The item keys are listed without optional parameters and additional information. Click the item key to see the full details.

See also VMware monitoring item keys.

Item key	Description
icmpping	The host accessibility by ICMP ping.
icmppingloss	The percentage of lost packets.
icmppingretry	The host accessibility by ICMP ping with retries.
icmppingsec	The ICMP ping response time.
net.tcp.service	Checks if a service is running and accepting TCP connections.
net.tcp.service.perf	Checks the performance of a TCP service.
net.udp.service	Checks if a service is running and responding to UDP requests.
net.udp.service.perf	Checks the performance of a UDP service.

Item key details

Parameters without angle brackets are mandatory. Parameters marked with angle brackets **< >** are optional.

icmpping[<target>,<packets>,<interval>,<size>,<timeout>,<options>]

 The host accessibility by ICMP ping.
 Return value: *0* - ICMP ping fails; *1* - ICMP ping successful.

Parameters:

- **target** - the host IP or DNS name;
- **packets** - the number of packets;
- **interval** - the time between successive packets in milliseconds;
- **size** - the packet size in bytes;
- **timeout** - the timeout in milliseconds;
- **options** - used for allowing redirect: if empty (default value), redirected responses are treated as target host down; if set to *allow_redirect*, redirected responses are treated as target host up.

See also the table of default values.

Comment:

- If at least one of the packets defined by the `packets` parameter is returned, the item returns 1. If all of the packets fail to return, the item returns 0.

Example:

```
icmpping[,4] #If at least one of the four packets is returned, the item returns 1.
```

icmppingloss[<target>,<packets>,<interval>,<size>,<timeout>,<options>]

 The percentage of lost packets.
 Return value: *Float*.

Parameters:

- **target** - the host IP or DNS name;
- **packets** - the number of packets;
- **interval** - the time between successive packets in milliseconds;
- **size** - the packet size in bytes;
- **timeout** - the timeout in milliseconds;
- **options** - used for allowing redirect: if empty (default value), redirected responses are treated as target host down; if set to *allow_redirect*, redirected responses are treated as target host up.

See also the table of default values.

icmppingretry[<target>,<retries>,<backoff>,<size>,<timeout>,<options>]

 The host accessibility by ICMP ping with retries. If the first packet succeeds, stop; if the packet fails, retry until the maximum number of attempts defined in the `retries` parameter is reached. This item is useful to reduce the number of packets sent over the network.
 Return value: *0* - ICMP ping fails; *1* - ICMP ping successful.

Parameters:

- **target** - the host IP or DNS name;
- **retries** - the number of times an attempt at pinging a target will be made, not including the first try (0 or greater; default 1);
- **backoff** - the number by which the wait time is multiplied on each successive request (1.0-5.0 range; default 1.0);
- **size** - the packet size in bytes;
- **timeout** - the timeout in milliseconds;
- **options** - used for allowing redirect: if empty (default value), redirected responses are treated as target host down; if set to *allow_redirect*, redirected responses are treated as target host up.

See also the table of default values.

icmppingsec[<target>,<packets>,<interval>,<size>,<timeout>,<mode>,<options>]

 The ICMP ping response time (in seconds).
 Return value: *Float*.

Parameters:

- **target** - the host IP or DNS name;
- **packets** - the number of packets;
- **interval** - the time between successive packets in milliseconds;
- **size** - the packet size in bytes;
- **timeout** - the timeout in milliseconds;
- **mode** - possible values: *min*, *max*, or *avg* (default);
- **options** - used for allowing redirect: if empty (default value), redirected responses are treated as target host down; if set to *allow_redirect*, redirected responses are treated as target host up.

Comments:

- Packets which are lost or timed out are not used in the calculation;
- If the host is not available (timeout reached), the item will return 0;
- If the return value is less than 0.0001 seconds, the value will be set to 0.0001 seconds;
- See also the table of default values.

net.tcp.service[service,<ip>,<port>]

 Checks if a service is running and accepting TCP connections.
 Return value: *0* - the service is down; *1* - the service is running.

Parameters:

- **service** - possible values: *ssh*, *ldap*, *smtp*, *ftp*, *http*, *pop*, *nntp*, *imap*, *tcp*, *https*, *telnet* (see details);
- **ip** - the IP address or DNS name (by default the host IP/DNS is used);
- **port** - the port number (by default the standard service port number is used).

Comments:

- Note that with *tcp* service indicating the port is mandatory;
- These checks may result in additional messages in system daemon logfiles (SMTP and SSH sessions being logged usually);
- Checking of encrypted protocols (like IMAP on port 993 or POP on port 995) is currently not supported. As a workaround, please use `net.tcp.service[tcp,<ip>,port]` for checks like these.

Example:

```
net.tcp.service[ftp,,45] #This item can be used to test the availability of FTP server on TCP port 45.
```

Attention:
If SELinux is running in enforced mode, custom TCP/UDP simple checks may be blocked by the policy. To verify and allow the new outgoing connection, review audit denials: with `grep denied /var/log/audit/audit.log`

net.tcp.service.perf[service,<ip>,<port>]

 Checks the performance of a TCP service.
 Return value: *Float*: *0.000000* - the service is down; *seconds* - the number of seconds spent while connecting to the service.

Parameters:

- **service** - possible values: *ssh*, *ldap*, *smtp*, *ftp*, *http*, *pop*, *nntp*, *imap*, *tcp*, *https*, *telnet* (see details);
- **ip** - the IP address or DNS name (by default the host IP/DNS is used);
- **port** - the port number (by default the standard service port number is used).

Comments:

- Note that with *tcp* service indicating the port is mandatory;
- Checking of encrypted protocols (like IMAP on port 993 or POP on port 995) is currently not supported. As a workaround, please use `net.tcp.service[tcp,<ip>,port]` for checks like these.

Example:

```
net.tcp.service.perf[ssh] #This item can be used to test the speed of initial response from SSH server.
```

net.udp.service[service,<ip>,<port>]

 Checks if a service is running and responding to UDP requests.
 Return value: *0* - the service is down; *1* - the service is running.

Parameters:

- **service** - possible values: *ntp* (see details);
- **ip** - the IP address or DNS name (by default the host IP/DNS is used);
- **port** - the port number (by default the standard service port number is used).

Example:

```
net.udp.service[ntp,,45] #This item can be used to test the availability of NTP service on UDP port 45.
```

net.udp.service.perf[service,<ip>,<port>]

 Checks the performance of a UDP service.
 Return value: *Float*: *0.000000* - the service is down; *seconds* - the number of seconds spent waiting for response from the service.

Parameters:

- **service** - possible values: *ntp* (see details);

- **ip** - the IP address or DNS name (by default the host IP/DNS is used);
- **port** - the port number (by default the standard service port number is used).

Example:

```
net.udp.service.perf[ntp] #This item can be used to test the response time from NTP service.
```

Attention:
For SourceIP support in LDAP simple checks (e.g., `net.tcp.service[ldap]`), OpenLDAP version 2.6.1 or above is required.

Timeout processing

Flexible item timeouts, even though supported for simple checks, do not apply to `icmpping*` and VMware items. See Flexible timeout support.

ICMP pings

Zabbix uses an external utility **fping** to process ICMP pings (**icmpping**, **icmppingloss**, **icmppingretry**, **icmppingsec**).

Installation

fping is not included with Zabbix and needs to be installed separately:

- Various Unix-based platforms have the fping package in their default repositories, but it is not pre-installed. In this case you can use the package manager to install fping.
- Zabbix provides fping packages for RHEL and its derivatives. Please note that these packages are provided without official support.
- fping can also be compiled from source.

Configuration

Specify fping location in the *FpingLocation* parameter of Zabbix server/proxy configuration file (or *Fping6Location* parameter for using IPv6 addresses).

fping should be executable by the user Zabbix server/proxy run as and this user should have sufficient rights.

See also: Known issues for processing simple checks with fping versions below 3.10.

Default values

Defaults, limits and description of values for ICMP check parameters:

Parameter	Unit	Description	Fping's flag	Defaults set by		Allowed limits by Zabbix	
				fping	**Zabbix**	**min**	**max**
packets	number	Number of request packets sent to a target	-C		3	1	10000
interval	milliseconds	Time to wait between successive packets to an individual target	-p	1000		20	unlimited
size	bytes	Packet size in bytes 56 bytes on x86, 68 bytes on x86_64	-b	56 or 68		24	65507
timeout	milliseconds	**fping v3.x** - timeout to wait after last packet sent, affected by -C flag **fping v4.x** - individual timeout for each packet	-t	**fping v3.x** - 500 **fping v4.x** and newer - inherited from -p flag, but not more than 2000		50	unlimited

Parameter	Unit	Description	Fping's flag	Defaults set by		Allowed limits by Zabbix	
retries	number	Number of times an attempt at pinging a target will be made, not including the first try	-r	3	1	0	unlimited
backoff factor	number	Number by which the wait time is multiplied on each successive request	-B	1.5	1.0	1.0	5.0

The defaults may differ slightly depending on the platform and version.

In addition, Zabbix uses fping options `-i interval ms` (do not mix up with the item parameter `interval` mentioned in the table above, which corresponds to fping option `-p`) and `-S source IP address` (or `-I` in older fping versions). These options are auto-detected by running checks with different option combinations. Zabbix tries to detect the minimum value in milliseconds that fping allows for the `-i` option by trying three values: 0, 1, and 10. The value that first succeeds is then used for subsequent ICMP checks. This process is done by each ICMP pinger process individually.

Auto-detected fping options are invalidated every hour and detected again on the next attempt to perform ICMP check. Set DebugLevel>=4 in order to view details of this process in the server or proxy log file.

Zabbix writes IP addresses to be checked by any of the `icmpping*` keys to a temporary file, which is then passed to fping. If items have different key parameters, only the ones with identical key parameters are written to a single file. All IP addresses written to the single file will be checked by fping in parallel, so Zabbix ICMP pinger process will spend fixed amount of time disregarding the number of IP addresses in the file.

VMware monitoring item keys

List of VMware monitoring item keys has been moved to VMware monitoring section.

SNMP agent

Overview

You may want to use SNMP monitoring on devices such as printers, network switches, routers or UPS that usually are SNMP-enabled and on which it would be impractical to attempt setting up complete operating systems and Zabbix agents.

To be able to retrieve data provided by SNMP agents on these devices, Zabbix server must be initially configured with SNMP support by specifying the `--with-net-snmp` flag. It is recommended to also install MIB files to ensure that item values are displayed in the correct format. Without the MIB files, formatting issues, such as displaying values in HEX instead of UTF-8 or vice versa, may occur.

SNMP checks are performed over the UDP protocol only.

Zabbix server and proxy daemons log lines similar to the following if they receive an incorrect SNMP response:

```
SNMP response from host "gateway" does not contain all of the requested variable bindings
```

While they do not cover all the problematic cases, they are useful for identifying individual SNMP devices for which combined requests should be disabled.

Zabbix server/proxy will retry up to 5 times for SNMP `walk` and `get` items. The retry mechanism doesn't apply to DNS resolution failures.

For legacy SNMP checks (single OID number or string) Zabbix server/proxy will retry at least one time after an unsuccessful query attempt: either through the SNMP library's retrying mechanism or through the internal combined processing mechanism.

Warning:
If monitoring SNMPv3 devices, make sure that msgAuthoritativeEngineID (also known as snmpEngineID or "Engine ID") is never shared by two devices. According to RFC 2571 (section 3.1.1.1) it must be unique for each device.

Warning:
RFC3414 requires the SNMPv3 devices to persist their engineBoots. Some devices do not do that, which results in their SNMP messages being discarded as outdated after being restarted. In such situation, SNMP cache needs to be manually cleared on a server/proxy (by using -R snmp_cache_reload) or the server/proxy needs to be restarted.

Note:
Zabbix caches SNMPv3 EngineID → IP mappings and reuses cached EngineIDs for subsequent checks instead of sending a probe each time, reducing network traffic. If an EngineID cannot be reused, a retry with a probe will be performed to discover the new EngineID.

Configuring SNMP monitoring

To start monitoring a device through SNMP, the following steps have to be performed:

Step 1

Find out the SNMP string (or OID) of the item you want to monitor.

To get a list of SNMP strings, use the `snmpwalk` command (part of net-snmp software which you should have installed as part of the Zabbix installation) or equivalent tool:

```
snmpwalk -v 2c -c public <host IP> .
```

As 2c here stands for SNMP version, you may also substitute it with 1, to indicate SNMP Version 1 on the device.

This should give you a list of SNMP strings and their last value. If it doesn't then it is possible that the SNMP 'community' is different from the standard 'public' in which case you will need to find out what it is.

You can then go through the list until you find the string you want to monitor, e.g., if you wanted to monitor the bytes coming in to your switch on port 3 you would use the `IF-MIB::ifHCInOctets.3` string from this line:

```
IF-MIB::ifHCInOctets.3 = Counter64: 3409739121
```

You may now use the `snmpget` command to find out the numeric OID for `IF-MIB::ifHCInOctets.3`:

```
snmpget -v 2c -c public -On <host IP> IF-MIB::ifHCInOctets.3
```

Note that the last number in the string is the port number you are looking to monitor. See also: Dynamic indexes.

This should give you something like the following:

```
.1.3.6.1.2.1.31.1.1.1.6.3 = Counter64: 3472126941
```

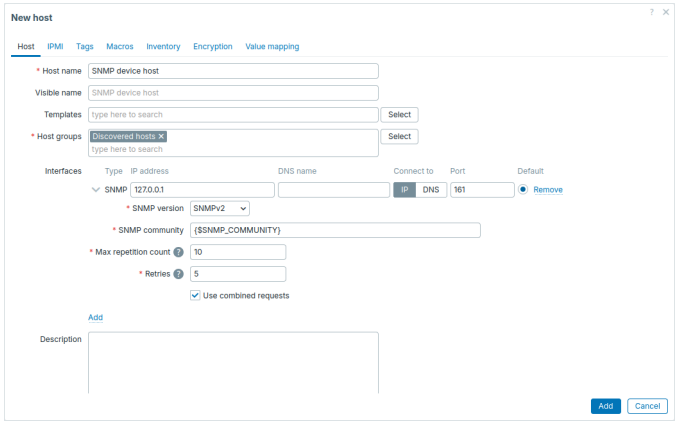
Again, the last number in the OID is the port number.

Note:
Some of the most used SNMP OIDs are translated automatically to a numeric representation by Zabbix.

In the last example above value type is "Counter64", which internally corresponds to ASN_COUNTER64 type. The full list of supported types is ASN_COUNTER, ASN_COUNTER64, ASN_UINTEGER, ASN_UNSIGNED64, ASN_INTEGER, ASN_INTEGER64, ASN_FLOAT, ASN_DOUBLE, ASN_TIMETICKS, ASN_GAUGE, ASN_IPADDRESS, ASN_OCTET_STR and ASN_OBJECT_ID. These types roughly correspond to "Counter32", "Counter64", "UInteger32", "INTEGER", "Float", "Double", "Timeticks", "Gauge32", "IpAddress", "OCTET STRING", "OBJECT IDENTIFIER" in **snmpget** output, but might also be shown as "STRING", "Hex-STRING", "OID" and other, depending on the presence of a display hint.

Step 2

Create a host corresponding to a device.



Add an SNMP interface for the host:

1. Enter the IP address/DNS name and port number.
2. Select the *SNMP version* from the dropdown.
3. Add interface credentials depending on the selected SNMP version:
    - SNMPv1 and SNMPv2 require only the community (the default value is `{$SNMP_COMMUNITY}`, a global macro that resolves to `public`).
    - SNMPv3 requires additional options; see SNMPv3.
4. Specify the *Max repetition count* value (default: 10) for native SNMP bulk requests (GetBulkRequest-PDUs); only for the `walk[]` item in SNMPv2 and SNMPv3. Note that setting this value too high may cause the SNMP agent check to exceed its timeout. User macros are supported.
5. Specify the maximum number of retries after a timeout in the *Retries* field for `walk[]` and `get[]` items (default: 5). For `walk[]`, retries are shared across the entire "walk". Setting this value to 0 disables retries. User macros are supported.
6. Mark the *Use combined requests* checkbox to allow combined processing of SNMP requests. This is not related to native SNMP bulk requests (`walk` and `get`).

You can use one of the provided SNMP templates that will automatically add a set of items. Before using a template, verify that it is compatible with the host.

Click *Add* to save the host.

SNMPv3

The following parameters are required for SNMPv3:

- *Context name* - Enter context name to identify item on SNMP subnet.
User macros are resolved in this field.
- *Security name* - Enter security name.
User macros are resolved in this field.
- *Security level* - Select security level:
    - *noAuthNoPriv* - no authentication nor privacy protocols are used;
    - *AuthNoPriv* - authentication protocol is used, privacy protocol is not;
    - *AuthPriv* - both authentication and privacy protocols are used.
- *Authentication protocol* - Select authentication protocol: *MD5*, *SHA1*; with net-snmp 5.8 and newer *SHA224*, *SHA256*, *SHA384*, or *SHA512*.
- *Authentication passphrase* - Enter authentication passphrase.
User macros are resolved in this field.
- *Privacy protocol* - Select privacy protocol: *DES*, *AES128*, *AES192*, *AES256*, *AES192C* (Cisco), or *AES256C* (Cisco).
See notes about privacy protocol support.
- *Privacy passphrase* - Enter privacy passphrase.
User macros are resolved in this field.

In case of wrong SNMPv3 credentials (security name, authentication protocol/passphrase, privacy protocol):

- Zabbix receives an ERROR from net-snmp, except for wrong *Privacy passphrase* in which case Zabbix receives a TIMEOUT error from net-snmp.
- SNMP interface availability will switch to red (unavailable).

Note:
Changes in *Authentication protocol*, *Authentication passphrase*, *Privacy protocol* or *Privacy passphrase*, made without changing the *Security name*, are normally applied automatically when the corresponding SNMPv3 interface is updated in Zabbix. In cases where *Security name* is also changed, all parameters will be updated immediately.

Privacy protocol support

Depending on your operating system and net-snmp configuration, some privacy protocols may not be available:

- On some newer operating systems (for example, RHEL9) support of DES is dropped for the `net-snmp` package.

- Encryption protocols AES192 and stronger are not supported out-of-the-box on the operating systems older then RHEL 8, CentOS 8, Oracle Linux 8, Debian 12, Ubuntu LTS 22.04, openSUSE Leap 15.5.

To check whether net-snmp library supports AES192+, use one of the following options:

1. `net-snmp-config`:

```
net-snmp-config --configure-options
```

If the output contains `--enable-blumenthal-aes`, AES192+ is supported.

Note that `net-snmp-config` is part of the development package for SNMP (`libsnmp-dev` for Debian/Ubuntu, `net-snmp-devel` for CentOS/RHEL/OL/SUSE) and may not be installed by default.

2. `snmpget`:

```
snmpget -v 3 -x AES-256
```

If the output contains `Invalid privacy protocol specified after -3x flag: AES-256`, AES192+ is not supported. If the output contains `No hostname specified.`, AES192+ is not supported.

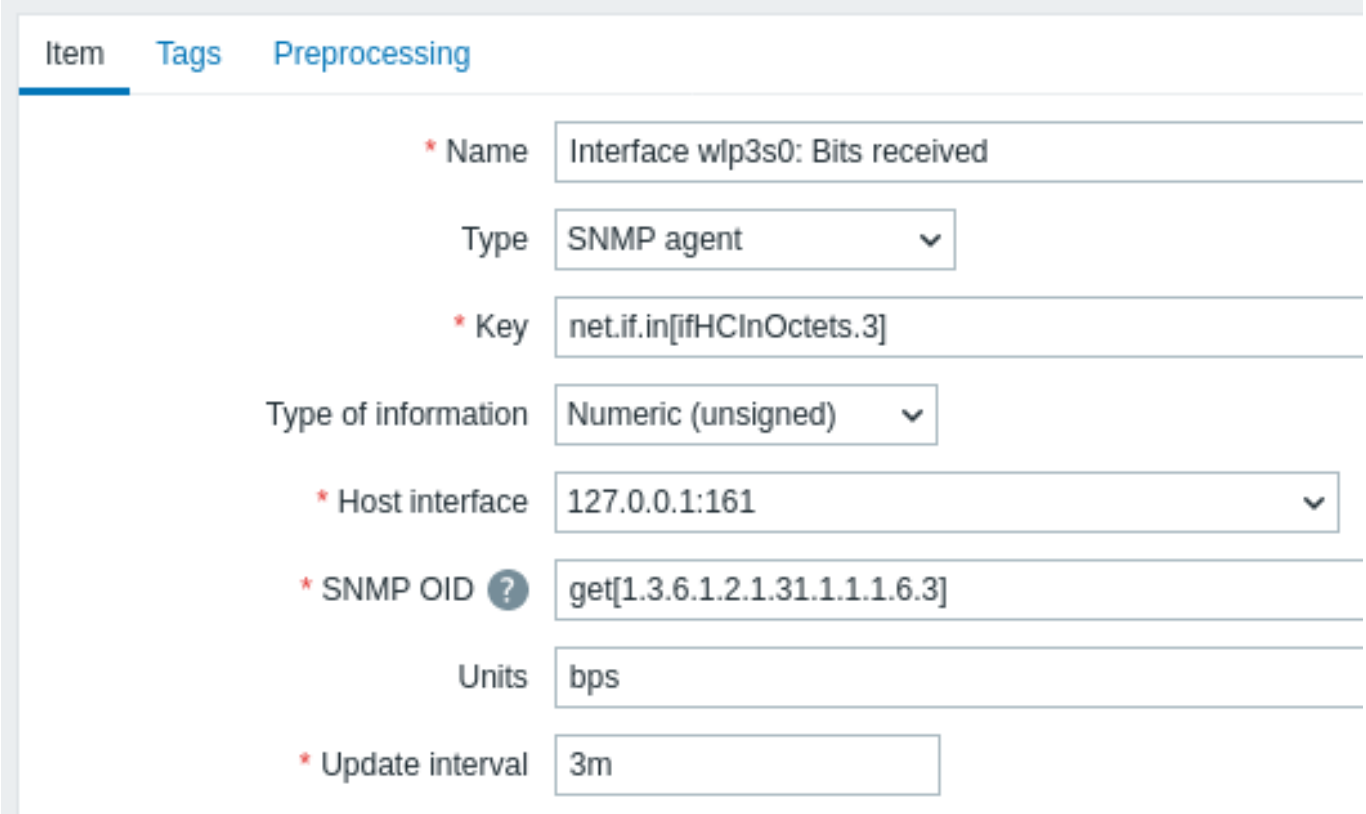
If your net-snmp library does not support AES192 and higher protocols, recompile net-snmp with `--enable-blumenthal-aes` option, then recompile Zabbix server specifying the option `--with-net-snmp=/home/user/yourcustomnetsnmp/bin/net-snmp-config`.

Step 3

Create an item for monitoring.

So, now go back to Zabbix and click *Items* for the SNMP host you created earlier. Depending on whether you used a template or not when creating your host, you will have either a list of SNMP items associated with your host or just an empty list. We will work on the assumption that you are going to create the item yourself using the information you have just gathered using snmpwalk and snmpget, so click *Create item*.

Fill in the required parameters in the new item form:



Parameter	Description
Name	Enter the item name.
Type	Select **SNMP agent** here.
Key	Enter the key as something meaningful.
Host interface	Make sure to select the SNMP interface, e.g., of your switch/router.

Parameter	Description
SNMP OID	Use one of the supported formats to enter OID value(s): **walk[OID1,OID2,...]** - retrieve a subtree of values. For example: `walk[1.3.6.1.2.1.2.2.1.2,1.3.6.1.2.1.2.2.1.3]`. This option makes use of native SNMP bulk requests (GetBulkRequest-PDUs) asynchronously. The timeout settings for this item can be set in the item configuration form. Consider setting a low timeout value to avoid long delays if the device is unreachable, as up to 5 retries will be made if earlier ones time out or fail (e.g., 3-second timeout can result in 15-second wait time). You may use this as the master item, with dependent items that extract data from the master using preprocessing. It is possible to specify multiple OIDs in a single snmp walk, such as `walk[OID1,OID2,...]` to asynchronously process one OID at a time. If the bulk request returns no results then it is attempted to retrieve a single record without bulk request. MIB names are supported as parameters; thus `walk[1.3.6.1.2.1.2.2.1.2]` and `walk[ifDescr]` will return the same output. If several OIDs/MIBs are specified, i.e. `walk[ifDescr,ifType,ifPhysAddress]`, then the output is a concatenated list. GetBulk requests are used with SNMPv2 and v3 interfaces and GetNext for SNMPv1 interfaces; the maximum number of repetitions for bulk requests is configured at the interface level, using the *Max repetition count* parameter. The parameter affects bulk requests by determining the maximum number of OIDs returned in a single bulk response. A higher value results in larger bulk responses, reducing the number of transmissions required. However, not all devices might support very high values, which could cause issues. This item returns the output of the snmpwalk utility with -Oe -Ot -On parameters. You may use this item as a master item in SNMP discovery. **get[OID]** - retrieve a *single* value asynchronously. For example: `get[1.3.6.1.2.1.31.1.1.1.6.3]` The timeout settings for this item can be set in the item configuration form. Consider setting a low timeout value to avoid long delays if the device is unreachable, as up to 5 retries will be made if earlier ones time out or fail (e.g., 3-second timeout can result in 15-second wait time). **OID** - (legacy) enter a single textual or numeric OID to retrieve a single value synchronously, optionally combined with other values. For example: `1.3.6.1.2.1.31.1.1.1.6.3`. For this option, the item check timeout will be equal to the value set in the server configuration file. It is **recommended** to use `walk[OID]` and `get[OID]` items for better performance. All `walk[OID]` and `get[OID]` items are executed asynchronously - it is not required to receive the response to one request before other checks are started. DNS resolving is asynchronous as well. The maximum concurrency of asynchronous checks is 1000 (defined by MaxConcurrentChecksPerPoller). The number of asynchronous SNMP pollers is defined by the StartSNMPPollers parameter. Note that for network traffic statistics, returned by any of the methods, a *Change per second* step must be added in the *Preprocessing* tab; otherwise, you will get the cumulative value from the SNMP device instead of the latest change.

All mandatory input fields are marked with a red asterisk.

Now save the item and go to *Monitoring* > *Latest data* for your SNMP data.

Example 1

General example:

Parameter	Description
OID	1.2.3.45.6.7.8.0 (or .1.2.3.45.6.7.8.0)

Parameter	Description
Key	<Unique string to be used as reference to triggers> For example, ”my_param”.

Note that OID can be given in either numeric or string form. However, in some cases, string OID must be converted to numeric representation. Utility snmpget may be used for this purpose:

```
snmpget -On localhost public enterprises.ucdavis.memory.memTotalSwap.0
```

Example 2

Monitoring of uptime:

Parameter	Description
OID	MIB::sysUpTime.0
Key	router.uptime
Value type	Float
Units	uptime
Preprocessing step: Custom multiplier	0.01

Native SNMP bulk requests

The **walk[OID1,OID2,...]** item allows the use of native SNMP functionality for bulk requests (GetBulkRequest-PDUs), available in SNMP versions 2/3.

A GetBulk request in SNMP executes multiple GetNext requests and returns the result in a single response. This may be used for regular SNMP items as well as for SNMP discovery to minimize network roundtrips.

The SNMP **walk[OID1,OID2,...]** item may be used as the master item that collects data in one request with dependent items that parse the response as needed using preprocessing.

Note that using native SNMP bulk requests is not related to the option of combining SNMP requests, which is Zabbix own way of combining multiple SNMP requests (see next section).

Up to five retries will occur for SNMP bulk items to avoid failure if one of the packets is lost. The timeout for SNMP items with `get` and `walk` (set in the item configuration form) is set for a whole session. Timeout is applied irrespective of whether data is retrieved completely; if data is received partially (for example, data is successfully collected for only one of multiple OIDs), then the item becomes unsupported with a message ”Only partial data received”. If the timeout is reached then a retry will occur, the timeout will be reset, and the last request will be resent allowing the session to continue from the last request if a single packet is lost or arrives too late. Consider setting a low timeout value to avoid long delays if the device is unreachable, as up to 5 retries will be made if earlier ones time out or fail (e.g., 3-second timeout can result in 15-second wait time).

Internal workings of combined processing

Zabbix server and proxy may query SNMP devices for multiple values in a single request. This affects several types of SNMP items:

- regular SNMP items
- SNMP items with dynamic indexes
- SNMP low-level discovery rules

All SNMP items on a single interface with identical parameters are scheduled to be queried at the same time. The first two types of items are taken by pollers in batches of at most 128 items, whereas low-level discovery rules are processed individually, as before.

On the lower level, there are two kinds of operations performed for querying values: getting multiple specified objects and walking an OID tree.

For ”getting”, a GetRequest-PDU is used with at most 128 variable bindings. For ”walking”, a GetNextRequest-PDU is used for SNMPv1 and GetBulkRequest with ”max-repetitions” field of at most 128 is used for SNMPv2 and SNMPv3.

Thus, the benefits of combined processing for each SNMP item type are outlined below:

- regular SNMP items benefit from ”getting” improvements;
- SNMP items with dynamic indexes benefit from both ”getting” and ”walking” improvements: ”getting” is used for index verification and ”walking” for building the cache;
- SNMP low-level discovery rules benefit from ”walking” improvements.

However, there is a technical issue that not all devices are capable of returning 128 values per request. Some always return a proper response, but others either respond with a ”tooBig(1)” error or do not respond at all once the potential response is over a certain limit.

In order to find an optimal number of objects to query for a given device, Zabbix uses the following strategy. It starts cautiously with querying 1 value in a request. If that is successful, it queries 2 values in a request. If that is successful again, it queries 3 values in a request and continues similarly by multiplying the number of queried objects by 1.5, resulting in the following sequence of request sizes: 1, 2, 3, 4, 6, 9, 13, 19, 28, 42, 63, 94, 128.

However, once a device refuses to give a proper response (for example, for 42 variables), Zabbix does two things.

First, for the current item batch it halves the number of objects in a single request and queries 21 variables. If the device is alive, then the query should work in the vast majority of cases, because 28 variables were known to work and 21 is significantly less than that. However, if that still fails, then Zabbix falls back to querying values one by one. If it still fails at this point, then the device is definitely not responding and request size is not an issue.

The second thing Zabbix does for subsequent item batches is it starts with the last successful number of variables (28 in our example) and continues incrementing request sizes by 1 until the limit is hit. For example, assuming the largest response size is 32 variables, the subsequent requests will be of sizes 29, 30, 31, 32, and 33. The last request will fail and Zabbix will never issue a request of size 33 again. From that point on, Zabbix will query at most 32 variables for this device.

If large queries fail with this number of variables, it can mean one of two things. The exact criteria that a device uses for limiting response size cannot be known, but we try to approximate that using the number of variables. So the first possibility is that this number of variables is around the device's actual response size limit in the general case: sometimes response is less than the limit, sometimes it is greater than that. The second possibility is that a UDP packet in either direction simply got lost. For these reasons, if Zabbix gets a failed query, it reduces the maximum number of variables to try to get deeper into the device's comfortable range, but only up to two times.

In the example above, if a query with 32 variables happens to fail, Zabbix will reduce the count to 31. If that happens to fail, too, Zabbix will reduce the count to 30. However, Zabbix will not reduce the count below 30, because it will assume that further failures are due to UDP packets getting lost, rather than the device's limit.

If, however, a device cannot handle combined requests properly for other reasons and the heuristic described above does not work, there is a "Use combined requests" setting for each interface that allows you to disable combined requests for that device.

If combined requests cause partial or malformed responses that result in incorrect per-second (delta) calculations (for example, apparent spikes in interface counters), disable *Use combined requests* for the affected interface to force separate per-item queries; this often prevents false spikes. Alternatively, consider using asynchronous `get[]` or `walk[]` items, which are executed asynchronously and are not subject to the per-interface *Use combined requests* batching—they can be used instead of legacy synchronous OID checks to avoid combined-request related issues. Look for server/proxy log entries similar to the one shown in the Overview section to help identify affected devices.

Additionally, if the interface frequently becomes unavailable, it may be necessary to increase the `UnavailableDelay` parameter in the Zabbix server or Zabbix proxy configuration files to reduce the frequency of requests. Items may become unsupported if partial data is received during discovery or OID walks.

Dynamic indexes

Overview

While you may find the required index number (for example, of a network interface) among the SNMP OIDs, sometimes you may not completely rely on the index number always staying the same.

Index numbers may be dynamic - they may change over time and your item may stop working as a consequence.

To avoid this scenario, it is possible to define an OID which takes into account the possibility of an index number changing.

For example, if you need to retrieve the index value to append to **ifInOctets** that corresponds to the **GigabitEthernet0/1** interface on a Cisco device, use the following OID:

```
ifInOctets["index","ifDescr","GigabitEthernet0/1"]
```

The syntax

A special syntax for OID is used:

<OID of data>["index","<base OID of index>","<string to search for>"]

Parameter	Description
OID of data	Main OID to use for data retrieval on the item.
index	Method of processing. Currently one method is supported: **index** - search for index and append it to the data OID
base OID of index	This OID will be looked up to get the index value corresponding to the string.

Parameter	Description
string to search for	The string to use for an exact match with a value when doing lookup. Case sensitive.

Example

Getting memory usage of *apache* process.

If using this OID syntax:

```
HOST-RESOURCES-MIB::hrSWRunPerfMem["index","HOST-RESOURCES-MIB::hrSWRunPath", "/usr/sbin/apache2"]
```

the index number will be looked up here:

```
...
HOST-RESOURCES-MIB::hrSWRunPath.5376 = STRING: "/sbin/getty"
HOST-RESOURCES-MIB::hrSWRunPath.5377 = STRING: "/sbin/getty"
HOST-RESOURCES-MIB::hrSWRunPath.5388 = STRING: "/usr/sbin/apache2"
HOST-RESOURCES-MIB::hrSWRunPath.5389 = STRING: "/sbin/sshd"
...
```

Now we have the index, 5388. The index will be appended to the data OID in order to receive the value we are interested in:

```
HOST-RESOURCES-MIB::hrSWRunPerfMem.5388 = INTEGER: 31468 KBytes
```

Index lookup caching

When a dynamic index item is requested, Zabbix retrieves and caches whole SNMP table under base OID for index, even if a match would be found sooner. This is done in case another item would refer to the same base OID later - Zabbix would look up index in the cache, instead of querying the monitored host again. Note that each poller process uses separate cache.

In all subsequent value retrieval operations only the found index is verified. If it has not changed, value is requested. If it has changed, cache is rebuilt - each poller that encounters a changed index walks the index SNMP table again.

Special OIDs

Some of the most used SNMP OIDs are translated automatically to a numeric representation by Zabbix. For example, **ifIndex** is translated to **1.3.6.1.2.1.2.2.1.1**, **ifIndex.0** is translated to **1.3.6.1.2.1.2.2.1.1.0**.

The table contains list of the special OIDs.

Special OID	Identifier	Description
ifIndex	1.3.6.1.2.1.2.2.1.1	A unique value for each interface.
ifDescr	1.3.6.1.2.1.2.2.1.2	A textual string containing information about the interface.This string should include the name of the manufacturer, the product name and the version of the hardware interface.
ifType	1.3.6.1.2.1.2.2.1.3	The type of interface, distinguished according to the physical/link protocol(s) immediately 'below' the network layer in the protocol stack.
ifMtu	1.3.6.1.2.1.2.2.1.4	The size of the largest datagram which can be sent / received on the interface, specified in octets.
ifSpeed	1.3.6.1.2.1.2.2.1.5	An estimate of the interface's current bandwidth in bits per second.
ifPhysAddress	1.3.6.1.2.1.2.2.1.6	The interface's address at the protocol layer immediately 'below' the network layer in the protocol stack.
ifAdminStatus	1.3.6.1.2.1.2.2.1.7	The current administrative state of the interface.
ifOperStatus	1.3.6.1.2.1.2.2.1.8	The current operational state of the interface.
ifInOctets	1.3.6.1.2.1.2.2.1.10	The total number of octets received on the interface, including framing characters.
ifInUcastPkts	1.3.6.1.2.1.2.2.1.11	The number of subnetwork-unicast packets delivered to a higher-layer protocol.
ifInNUcastPkts	1.3.6.1.2.1.2.2.1.12	The number of non-unicast (i.e., subnetwork- broadcast or subnetwork-multicast) packets delivered to a higher-layer protocol.
ifInDiscards	1.3.6.1.2.1.2.2.1.13	The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space.

Special OID	Identifier	Description
ifInErrors	1.3.6.1.2.1.2.2.1.14	The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
ifInUnknownProtos	1.3.6.1.2.1.2.2.1.15	The number of packets received via the interface which were discarded because of an unknown or unsupported protocol.
ifOutOctets	1.3.6.1.2.1.2.2.1.16	The total number of octets transmitted out of the interface, including framing characters.
ifOutUcastPkts	1.3.6.1.2.1.2.2.1.17	The total number of packets that higher-level protocols requested be transmitted, and which were not addressed to a multicast or broadcast address at this sub-layer, including those that were discarded or not sent.
ifOutNUcastPkts	1.3.6.1.2.1.2.2.1.18	The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast or broadcast address at this sub-layer, including those that were discarded or not sent.
ifOutDiscards	1.3.6.1.2.1.2.2.1.19	The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space.
ifOutErrors	1.3.6.1.2.1.2.2.1.20	The number of outbound packets that could not be transmitted because of errors.
ifOutQLen	1.3.6.1.2.1.2.2.1.21	The length of the output packet queue (in packets).

MIB files

Introduction

MIB stands for the Management Information Base. MIB files enable the use of textual representations of OIDs (Object Identifiers). It is possible to use raw OIDs when monitoring SNMP devices with Zabbix, but if you feel more comfortable using textual representation, you need to install MIB files.

For example,

`ifHCOutOctets`

is textual representation of the OID

`1.3.6.1.2.1.31.1.1.1.10`

Installing MIB files

On Debian-based systems:

```
apt install snmp-mibs-downloader
download-mibs
```

On RedHat-based systems:

```
dnf install net-snmp-libs
```

Enabling MIB files

On RedHat-based systems, MIB files should be enabled by default. On Debian-based systems, you have to edit the file `/etc/snmp/snmp.conf` and comment out the line that says `mibs :`

```
# As the snmp packages come without MIB files due to license reasons, loading
# of MIBs is disabled by default. If you added the MIBs you can re-enable
# loading them by commenting out the following line.
mibs :
```

Testing MIB files

Testing SNMP MIBs can be done using `snmpwalk` utility. If you don't have it installed, use the following instructions.

On Debian-based systems:

```
apt install snmp
```

On RedHat-based systems:

```
dnf install net-snmp-utils
```

After that, the following command must not give error when you query a network device:

```
$ snmpwalk -v 2c -c public <NETWORK DEVICE IP> ifInOctets
IF-MIB::ifInOctets.1 = Counter32: 176137634
IF-MIB::ifInOctets.2 = Counter32: 0
IF-MIB::ifInOctets.3 = Counter32: 240375057
IF-MIB::ifInOctets.4 = Counter32: 220893420
[...]
```

Using MIBs in Zabbix

The most important to keep in mind is that Zabbix processes do not get informed of the changes made to MIB files. So after every change you must restart Zabbix server or proxy, e. g.:

```
systemctl restart zabbix-server
```

After that, the changes made to MIB files are in effect.

Using custom MIB files

There are standard MIB files coming with every GNU/Linux distribution. But some device vendors provide their own.

Let's say, you would like to use CISCO-SMI MIB file. The following instructions will download and install it:

```
wget ftp://ftp.cisco.com/pub/mibs/v2/CISCO-SMI.my -P /tmp
mkdir -p /usr/local/share/snmp/mibs
grep -q '^mibdirs +/usr/local/share/snmp/mibs' /etc/snmp/snmp.conf 2>/dev/null || echo "mibdirs +/usr/loca
cp /tmp/CISCO-SMI.my /usr/local/share/snmp/mibs
```

Now you should be able to use it. Try to translate the name of the object *ciscoProducts* from the MIB file to OID:

```
snmptranslate -IR -On CISCO-SMI::ciscoProducts
.1.3.6.1.4.1.9.1
```

If you receive errors instead of the OID, ensure all the previous commands did not return any errors.

The object name translation worked, you are ready to use custom MIB file. Note the MIB name prefix (*CISCO-SMI::*) used in the query. You will need this when using command-line tools as well as Zabbix.

Don't forget to restart Zabbix server/proxy before using this MIB file in Zabbix.

Attention:
Keep in mind that MIB files can have dependencies. That is, one MIB may require another. In order to satisfy these dependencies you have to install all the affected MIB files.

SNMP trap

Overview

Receiving SNMP traps is the opposite to querying SNMP-enabled devices.

In this case, the information is sent from an SNMP-enabled device to `snmptrapd` and is collected or "trapped" by Zabbix server or Zabbix proxy from a file.

Usually, traps are sent upon some condition change and the agent connects to the server on port 162 (as opposed to port 161 on the agent side that is used for queries). Using traps may detect some short problems that occur amidst the query interval and may be missed by the query data.

Receiving SNMP traps in Zabbix is designed to work with `snmptrapd` and one of the mechanisms for passing the traps to Zabbix - either a Bash or Perl script or SNMPTT.

Note:
The simplest way to set up trap monitoring after configuring Zabbix is to use the Bash script solution, because Perl and SNMPTT are often missing in modern distributions and require more complex configuration. However, this solution uses a script configured as `traphandle`. For better performance on production systems, use the embedded Perl solution (either script with `do perl` option or SNMPTT).

The workflow of receiving a trap:

1. snmptrapd receives a trap.
2. snmptrapd passes the trap to the receiver script (Bash, Perl) or SNMPTT.
3. The receiver parses, formats and writes the trap to a file.
4. Zabbix SNMP trapper reads and parses the trap file.
5. For each trap Zabbix finds all SNMP trapper items with host interfaces matching the received trap address. Note that only the selected *IP* or *DNS* in host interface is used during the matching.
6. For each found item, the trap is compared to regexp in `snmptrap[regexp]`. The trap is set as the value of **all** matched items. If no matching item is found and there is an `snmptrap.fallback` item, the trap is set as the value of that.
7. If the trap was not set as the value of any item, Zabbix by default logs the unmatched trap. (This is configured by *Log unmatched SNMP traps* in *Administration* > *General* > *Other*.)

Notes on HA failover

During high-availability (HA) node switch, Zabbix will continue processing after the last record within the last ISO 8601 timestamp; if the same record is not found then only the timestamp will be used to identify last position.

Configuring SNMP traps

This item type requires the following frontend configuration:

Step 1: Create an SNMP interface for the host

1. In *Data collection* > *Hosts*, create a host or edit the existing one.
2. In the *Interfaces* field, select the *SNMP* interface type.
3. Enter the IP address/DNS name and port number.
4. Select the *SNMP version* from the dropdown: (*SNMPv1*, *SNMPv2*, or *SNMPv3*) and add interface credentials depending on the selected SNMP version.

The address from each received trap will be compared to the IP and DNS addresses of all SNMP interfaces to find the corresponding hosts.

Step 2: Configure the item

1. For the host, create an item or edit the existing one.
2. In the *Key* field, use one of the SNMP trap keys:
   - `snmptrap[regexp]` - catches all SNMP traps that match the regular expression specified in the `regexp` parameter; if `regexp` is unspecified, catches any trap.
 User macros and global regular expressions are supported in the parameter.
 Return value: SNMP trap.
 This item can be set only for SNMP interfaces.
   - `snmptrap.fallback` - catches all SNMP traps that were not caught by any of the `snmptrap[]` items for that interface.
 Return value: SNMP trap.
 This item can be set only for SNMP interfaces.

Note:
Multiline regular expression matching is not supported at this time.

3. Set the *Type of information* to *Log* for the timestamps to be parsed. Other formats such as *Numeric* are also acceptable but might require a custom trap handler.

Setting up SNMP trap monitoring

Configuring Zabbix server/proxy

To read the traps, Zabbix server or proxy must be configured to start the SNMP trapper process and point to the trap file that is being written by SNMPTT or a Bash/Perl trap receiver. To do that, edit the configuration file (`zabbix_server.conf` or `zabbix_proxy.conf`):

```
StartSNMPTrapper=1
SNMPTrapperFile=[TRAP FILE]
```

Warning:
If systemd parameter `PrivateTmp` is used, this file is unlikely to work in `/tmp`.

Configuring Bash trap receiver

Requirements: only snmptrapd.

A Bash trap receiver script can be used to pass traps to Zabbix server from snmptrapd using trapper file. To configure it, add the `traphandle` option to snmptrapd configuration file (`snmptrapd.conf`), see example.

Note:
snmptrapd might need to be restarted to pick up changes to its configuration.

Configuring Perl trap receiver

Requirements: Perl, Net-SNMP compiled with `--enable-embedded-perl` (done by default since Net-SNMP 5.4)

A Perl trap receiver (look for `misc/snmptrap/zabbix_trap_receiver.pl`) can be used to pass traps to Zabbix server directly from snmptrapd. To configure it:

- Add the Perl script to the snmptrapd configuration file (`snmptrapd.conf`), e.g.:

```
perl do "[FULL PATH TO PERL RECEIVER SCRIPT]";
```

- Configure the receiver, e.g:

```
$SNMPTrapperFile = '[TRAP FILE]';
$DateTimeFormat = '[DATE TIME FORMAT]';
```

Note:
snmptrapd might need to be restarted to pick up changes to its configuration.

Note:
If the script name is not quoted, snmptrapd will refuse to start up with messages, similar to these:


```
Regexp modifiers "/l" and "/a" are mutually exclusive at (eval 2) line 1, at end of line
Regexp modifier "/l" may not appear twice at (eval 2) line 1, at end of line
```

Configuring SNMPTT

At first, snmptrapd should be configured to use SNMPTT.

Note:
For the best performance, SNMPTT should be configured as a daemon using `snmptthandler-embedded` to pass the traps to it. See instructions for configuring SNMPTT.

When SNMPTT is configured to receive the traps, configure `snmptt.ini`:

1. Enable the use of the Perl module from the `NET-SNMP` package:

```
net_snmp_perl_enable = 1
```

2. Log traps to the trap file which will be read by Zabbix:

```
log_enable = 1
log_file = [TRAP FILE]
```

3. Set the date-time format:

```
date_time_format = %Y-%m-%dT%H:%M:%S%z
```

Warning:
The `net-snmp-perl` package has been removed in RHEL 8.0-8.2; re-added in RHEL 8.3. For more information, see the known issues.

Now format the traps for Zabbix to recognize them (edit `snmptt.conf`):

1. Each `FORMAT` statement should start with `ZBXTRAP [address]`, where `[address]` will be compared to IP and DNS addresses of SNMP interfaces on Zabbix. E.g.:

```
EVENT coldStart .1.3.6.1.6.3.1.1.5.1 "Status Events" Normal
FORMAT ZBXTRAP $aA Device reinitialized (coldStart)
```

2. See more about SNMP trap format below.

Attention:
Do not use unknown traps - Zabbix will not be able to recognize them. Unknown traps can be handled by defining a general event in `snmptt.conf`:


```
EVENT general .* "General event" Normal
```

SNMP trap format

All customized Perl trap receivers and SNMPTT trap configuration must format the trap in the following way:

```
[timestamp] [the trap, part 1] ZBXTRAP [address] [the trap, part 2]
```

where

- `[timestamp]` - the timestamp in `%Y-%m-%dT%H:%M:%S%z` format.
- `ZBXTRAP` - header that indicates that a new trap starts in this line.
- `[address]` - IP address used to find the host for this trap.

Note that `ZBXTRAP` and `[address]` will be cut out from the message during processing. If the trap is formatted otherwise, Zabbix might parse the traps unexpectedly.

Example trap:

```
2024-01-11T15:28:47+0200 .1.3.6.1.6.3.1.1.5.3 Normal "Status Events" localhost - ZBXTRAP 192.168.1.1 Link
```

This will result in the following trap for SNMP interface with IP=`192.168.1.1`:

```
2024-01-11T15:28:47+0200 .1.3.6.1.6.3.1.1.5.3 Normal "Status Events"
localhost - Link down on interface 2. Admin state: 1. Operational state: 2
```

System requirements

Note:
It is recommended to install MIB files to ensure that item values are displayed in the correct format. Without the MIB files, formatting issues, such as displaying values in HEX instead of UTF-8 or vice versa, may occur.

Large file support

Zabbix has large file support for SNMP trapper files. The maximum file size that Zabbix can read is 2^63 (8 EiB). Note that the filesystem may impose a lower limit on the file size.

Log rotation

Zabbix does not provide any log rotation system - that should be handled by the user. The log rotation should first rename the old file and only later delete it so that no traps are lost:

1. Zabbix opens the trap file at the last known location and goes to step 3.
2. Zabbix checks if the currently opened file has been rotated by comparing the inode number to the defined trap file's inode number. If there is no opened file, Zabbix resets the last location and goes to step 1.
3. Zabbix reads the data from the currently opened file and sets the new location.
4. The new data is parsed. If this was the rotated file, the file is closed and goes back to step 2.
5. If there was no new data, Zabbix sleeps for 1 second and goes back to step 2.

File system

Because of the trap file implementation, Zabbix needs the file system to support inodes to differentiate files (the information is acquired by a `stat()` call).

Setup examples using different SNMP protocol versions

This example uses snmptrapd and a Bash receiver script to pass traps to Zabbix server.

Setup:

1. Configure Zabbix to start SNMP trapper and set the trap file. Add to `zabbix_server.conf`:

```
StartSNMPTrapper=1
SNMPTrapperFile=/var/lib/zabbix/snmptraps/snmptraps.log
```

2. Download the Bash script to `/usr/sbin/zabbix_trap_handler.sh`:

```
curl -o /usr/sbin/zabbix_trap_handler.sh https://raw.githubusercontent.com/zabbix/zabbix-docker/trunk/temp
```

If necessary, adjust the ZABBIX_TRAPS_FILE variable in the script. To use the default value, create the parent directory first:

```
mkdir -p /var/lib/zabbix/snmptraps
```

3. Add the following to snmtrapd.conf (refer to working example)

```
traphandle default /bin/bash /usr/sbin/zabbix_trap_handler.sh
```

Note:
snmptrapd might need to be restarted to pick up changes to its configuration.

4. Create an SNMP item TEST (keep in mind the initial configuration requirements):

   Type: *SNMP trap*
 Type of information: *Log*
 Host interface: *SNMP 127.0.0.1*
 Key: snmptrap["linkUp"]
 Log time format: yyyyMMdd.hhmmss

Note that the ISO 8601 date and time format is used.

5. Next we will configure snmptrapd for our chosen SNMP protocol version and send test traps using the snmptrap utility.

SNMPv1, SNMPv2

SNMPv1 and SNMPv2 protocols rely on "community string" authentication. In the example below we will use "secret" as community string. It must be set to the same value on SNMP trap senders.

Please note that while still widely used in production environments, SNMPv2 doesn't offer any encryption and real sender authentication. The data is sent as plain text and therefore these protocol versions should only be used in secure environments such as private network and should never be used over any public or third-party network.

SNMP version 1 isn't really used these days since it doesn't support 64-bit counters and is considered a legacy protocol.

To enable accepting SNMPv1 or SNMPv2 traps you should add the following line to snmptrapd.conf. Replace secret with the SNMP community string configured on SNMP trap senders:

```
authCommunity log,execute,net secret
```

Next we can send a test trap using snmptrap. We will use the common "link up" OID in this example:

```
snmptrap -v 2c -c secret localhost '' linkUp.0
```

SNMPv3

SNMPv3 addresses SNMPv1/v2 security issues and provides authentication and encryption. You can use the MD5 or multiple SHA authentication methods and DES/multiple AES as cipher.

Configure snmptrapd to accept SNMPv3 traps by adding the following lines to snmptrapd.conf:

```
##### Create an SNMPv3 USM user with authentication and privacy
createUser -e <engineID> <securityName> SHA <authKey> AES <privKey>
##### Allow the user to log traps and execute trap handler scripts
authUser log,execute <securityName>
```

For example:

```
createUser -e 0x8000000001020304 traptest SHA myAuthPassword AES myPrivPassword
authUser log,execute traptest
```

Attention:
Please note the "execute" keyword that allows you to execute scripts for this user security model.

To test the SNMPv3 trap configuration, send a test trap using the following command:

```
snmptrap -v 3 -n "" -a SHA -A myAuthPassword -x AES -X myPrivPassword -l authPriv -u traptest -e 0x8000000
```

Warning:
If you wish to use strong encryption methods such as AES192 or AES256, please use net-snmp starting with version 5.8. You might have to recompile it with configure option: --enable-blumenthal-aes. Older versions of net-snmp do not support AES192/AES256. See also: Strong Authentication or Encryption.

Verification

In both examples you will see similar lines in your /var/lib/zabbix/snmptraps/snmptraps.log:

```
2024-01-30T10:04:23+0200 ZBXTRAP 127.0.0.1
UDP: [127.0.0.1]:56585->[127.0.0.1]:162
DISMAN-EVENT-MIB::sysUpTimeInstance = 2538834
SNMPv2-MIB::snmpTrapOID.0 = IF-MIB::linkUp.0
```

The item value in Zabbix will be:

```
2024-01-30 10:04:23 2024-01-30 10:04:21

2024-01-30T10:04:21+0200 UDP: [127.0.0.1]:56585->[127.0.0.1]:162
DISMAN-EVENT-MIB::sysUpTimeInstance = 2538834
SNMPv2-MIB::snmpTrapOID.0 = IF-MIB::linkUp.0
```

Example with Perl:

```
2024-01-30T11:42:54+0200 ZBXTRAP 127.0.0.1
PDU INFO:
  receivedfrom                   UDP: [127.0.0.1]:58649->[127.0.0.1]:162
  notificationtype               TRAP
  version                        1
  community                      public
  errorstatus                    0
  transactionid                  1
  requestid                      2101882550
  messageid                      0
  errorindex                     0
VARBINDS:
  DISMAN-EVENT-MIB::sysUpTimeInstance type=67 value=Timeticks: (457671) 1:16:16.71
  SNMPv2-MIB::snmpTrapOID.0       type=6  value=OID: IF-MIB::linkUp.0
```

See also

- Zabbix blog article on SNMP traps
- Configuring snmptrapd (official net-snmp documentation)
- Configuring snmptrapd to receive SNMPv3 notifications (official net-snmp documentation)

Zabbix internal

Overview

Internal checks allow you to monitor the internal processes of Zabbix. In other words, you can monitor what goes on with Zabbix server or Zabbix proxy.

Internal checks are calculated:

- on Zabbix server - if the host is monitored by server
- on Zabbix proxy - if the host is monitored by proxy

Internal checks are processed by server or proxy regardless of the host maintenance status.

To use this item, choose the **Zabbix internal** item type.

> **Note:**
> Internal checks are processed by Zabbix pollers.

Performance

Using some internal items may negatively affect performance. These items are:

- `zabbix[host,,items]`
- `zabbix[host,,items_unsupported]`
- `zabbix[hosts]`
- `zabbix[items]`
- `zabbix[items_unsupported]`
- `zabbix[queue,,]`
- `zabbix[requiredperformance]`
- `zabbix[stats,,,queue,,]`

- zabbix[triggers]

The System information and Queue frontend sections are also affected.

Supported checks

The item keys are listed without customizable parameters and additional information. Click the item key to see the full details.

Item key	Description
zabbix[apm,]	The statistics of the telemetry data received by the APM collector.
zabbix[boottime]	The startup time of Zabbix server or Zabbix proxy process in seconds.
zabbix[cep]	Complex event processing statistics.
zabbix[cluster,discovery,nodes]	Discovers the high availability cluster nodes.
zabbix[connector_queue]	The count of values enqueued in the connector queue.
zabbix[discovery_queue]	The count of network checks enqueued in the discovery queue.
zabbix[host„items]	The number of enabled items (supported and not supported) on the host.
zabbix[host„items_unsupported]	The number of enabled unsupported items on the host.
zabbix[host„maintenance]	The current maintenance status of the host.
zabbix[host,active_agent,available]	The availability of active agent checks on the host.
zabbix[host,discovery,interfaces]	The details of all configured interfaces of the host in Zabbix frontend.
zabbix[host„available]	The availability of the main interface of a particular type of checks on the host.
zabbix[hosts]	The number of monitored hosts.
zabbix[items]	The number of enabled items (supported and not supported).
zabbix[items_unsupported]	The number of unsupported items.
zabbix[java„]	The information about Zabbix Java gateway.
zabbix[lld_queue]	The count of values enqueued in the low-level discovery processing queue.
zabbix[preprocessing]	Statistics of the values received by the preprocessing manager.
zabbix[preprocessing_queue]	The count of values enqueued in the preprocessing queue.
zabbix[process„„]	The percentage of time a particular Zabbix process or a group of processes (identified by <type> and <mode>) spent in <state>.
zabbix[proxy„]	The information about Zabbix proxy.
zabbix[proxy,discovery]	The list of Zabbix proxies.
zabbix[proxy group„available]	The number of online proxies in a proxy group.
zabbix[proxy group„pavailable]	The percentage of online proxies in a proxy group.
zabbix[proxy group„proxies]	The list of Zabbix proxies in a proxy group.
zabbix[proxy group„state]	The state of a proxy group.
zabbix[proxy group,discovery]	Returns a list of proxy groups with configuration data and real-time data.
zabbix[proxy_buffer,buffer,]	Returns the proxy memory buffer usage statistics.
zabbix[proxy_buffer,state,changes]	Returns the number of state changes between disk/memory buffer modes since start.
zabbix[proxy_buffer,state,current]	Returns the current working state where the new data is being stored.
zabbix[proxy_history]	The number of values in the proxy history table waiting to be sent to the server.
zabbix[queue„]	The number of monitored items in the queue which are delayed at least by <from> seconds, but less than <to> seconds.
zabbix[rcache„]	The availability statistics of the Zabbix configuration cache.
zabbix[requiredperformance]	The required performance of Zabbix server or Zabbix proxy, in new values per second expected.
zabbix[stats„]	Returns the internal metrics of a Zabbix server or proxy. If <ip> and <port> are specified, metrics are retrieved from the remote instance; otherwise, from the local instance.
zabbix[stats„,queue„]	Returns the internal queue metrics of a Zabbix server or proxy. If <ip> and <port> are specified, metrics are retrieved from the remote instance; otherwise, from the local instance.
zabbix[tcache„]	The effectiveness statistics of the Zabbix trend function cache.
zabbix[triggers]	The number of enabled triggers in Zabbix database, with all items enabled on enabled hosts.
zabbix[uptime]	The uptime of the Zabbix server or proxy process in seconds.
zabbix[vcache,buffer,]	The availability statistics of the Zabbix value cache.
zabbix[vcache,cache,]	The effectiveness statistics of the Zabbix value cache.
zabbix[version]	The version of Zabbix server or proxy.
zabbix[vmware,buffer,]	The availability statistics of the Zabbix vmware cache.
zabbix[vps,written]	The total number of history values written to database.
zabbix[wcache„]	The statistics and availability of the Zabbix write cache.

Item key details

- Parameters without angle brackets are mandatory and must be used *as is* (for example, ”host” and ”available” in `zabbix[host,<type>,available]`).
- Parameters with angle brackets **< >** must be replaced with a valid value. If a parameter has a default value, it can be omitted.
- Values for items and item parameters labeled ”not supported on proxy” can only be retrieved if the host is monitored by server. Conversely, values ”not supported on server” can only be retrieved if the host is monitored by proxy.

zabbix[apm,<type>]

 The statistics of the telemetry data received by the APM collector:

- *requests* - the number of requests received by the APM collector (counter):
  - *accepted* - the number of requests accepted for processing;
  - *dropped* - the number of requests that were dropped;
- *rows* - the number of telemetry data rows written to the data source, per data type (counter).

Return value: *JSON object*.

Parameters:

- **type** - *traces*, *metrics*, or *logs*. If not specified, statistics for all types are returned.

Comments:

- This item is not supported on Zabbix server.

Example of return values for `zabbix[apm]`:

```
{
  "requests": {
    "accepted": 15230,
    "dropped": 12
  },
  "rows": {
    "traces": 48710,
    "metrics_gauge": 31250,
    "metrics_sum": 18400,
    "metrics_histogram": 5120,
    "metrics_exponential_histogram": 640,
    "metrics_summary": 210,
    "logs": 22870
  }
}
```

zabbix[boottime]

 The startup time of Zabbix server or Zabbix proxy process in seconds.
 Return value: *Integer*.

zabbix[cep]

 Complex event processing statistics.
 Return value: *JSON*.

zabbix[cluster,discovery,nodes]

 Discovers the high availability cluster nodes.
 Return value: *JSON object*.

Comments:

- This item can be used in low-level discovery.
- Returned fields mostly match properties of the High availability node object.
- Additional fields:
  - `db_timestamp` - current server time (Unix timestamp) included in the response
  - `lastaccess_age` - seconds since the last node heartbeat (db_timestamp - lastaccess)
  - `status` - node state:
    - 0 - standby
    - 1 - stopped manually
    - 2 - unavailable
    - 3 - active

Example of return values:

```
[
  {
    "id": "ckvupihk70001z8mkpw5cg0u3",
    "name": "zabbix-prod-01",
    "status": 3,
    "address": "10.0.4.12:10051",
    "port": 10051,
    "lastaccess": 1756115995,
    "db_timestamp": 1756116000,
    "lastaccess_age": 5
  },
  {
    "id": "ckvx2a9k70004b1nq2hz9d7f",
    "name": "zabbix-standby-02",
    "status": 0,
    "address": "10.0.4.13:10051",
    "port": 10051,
    "lastaccess": 1756115550,
    "db_timestamp": 1756116000,
    "lastaccess_age": 450
  },
  {
    "id": "ckw0bq3l70007y4r1a0m5kz8",
    "name": "zabbix-backup-eu1",
    "status": 1,
    "address": "backup.example.com:10051",
    "port": 10051,
    "lastaccess": 1756105080,
    "db_timestamp": 1756116000,
    "lastaccess_age": 10920
  }
]
```


zabbix[connector_queue]

 The count of values enqueued in the connector queue.
 Return value: *Integer*.

zabbix[discovery_queue]

 The count of network checks enqueued in the discovery queue.
 Return value: *Integer*.

zabbix[host,,items]

 The number of enabled items (supported and not supported) on the host.
 Return value: *Integer*.

zabbix[host,,items_unsupported]

 The number of enabled unsupported items on the host.
 Return value: *Integer*.

zabbix[host,,maintenance]

 The current maintenance status of the host.
 Return values: *0* - normal state; *1* - maintenance with data collection; *2* - maintenance without data collection.

Comments:

- This item is always processed by Zabbix server regardless of the host location (on server or proxy). The proxy will not receive this item with configuration data.
- The second parameter must be empty and is reserved for future use.

zabbix[host,active_agent,available]

 The availability of active agent checks on the host.
 Return values: *0* - unknown; *1* - available; *2* - not available.

zabbix[host,discovery,interfaces]

 The details of all configured interfaces of the host in Zabbix frontend.
 Return value: *JSON object*.

Comments:

- This item can be used in low-level discovery.

- This item is not supported on Zabbix proxy.

zabbix[host,<type>,available]

 The availability of the main interface of a particular type of checks on the host.
 Return values: *0* - not available; *1* - available; *2* - unknown.

Parameters:

- **type** - *agent*, *snmp*, *ipmi*, or *jmx*.

Comments:

- The item value is calculated according to the configuration parameters regarding host unreachability/unavailability.

zabbixhosts

 The number of monitored hosts.
 Return value: *Integer*.

zabbixitems

 The number of enabled items (supported and not supported).
 Return value: *Integer*.

zabbix[items_unsupported]

 The number of unsupported items.
 Return value: *Integer*.

zabbix[java,,<param>]

 The information about Zabbix Java gateway.
 Return values: *1* - if `<param>` is *ping*; *Java gateway version* - if `<param>` is *version* (for example: "8.2.0").

Parameters:

- **param** - *ping* or *version*.

Comments:

- This item can be used to check Java gateway availability using the `nodata()` trigger function.
- The second parameter must be empty and is reserved for future use.

zabbix[lld_queue]

 The count of values enqueued in the low-level discovery processing queue.
 Return value: *Integer*.

Comments:

- This item can be used to monitor the low-level discovery processing queue length.

zabbixpreprocessing

 Statistics of the values received by the preprocessing manager:

- *queued* - the number and size of queued values that require preprocessing (counter)
- *direct* - the number and size of queued values that do not require preprocessing (counter)
- *queue* - the count of values enqueued in the preprocessing queue (same as `zabbix[preprocessing_queue]`)

Return value: *JSON*.

Example of return values:

```
{"data":
    {
    "queued": {
        "count": 106,
        "size": 58620
    },
    "direct": {
        "count": 395,
        "size": 33843
    },
    "queue": 0
    }
}
```

zabbix[preprocessing_queue]

 The count of values enqueued in the preprocessing queue.
 Return value: *Integer*.

Comments:

- This item can be used to monitor the preprocessing queue length.

zabbix[process,<type>,<mode>,<state>]

 The percentage of time a particular Zabbix process or a group of processes (identified by <type> and <mode>) spent in <state>. It is calculated for the last minute only.
 Return value: *Float*.

Parameters:

- **type**:
  - for server processes: *agent poller*, *alert manager*, *alert syncer*, *alerter*, *availability manager*, *browser poller*, *configuration syncer*, *configuration syncer worker*, *connector manager*, *connector worker*, *discovery manager*, *discovery worker*, *escalator*, *ha manager*, *history poller*, *history syncer*, *housekeeper*, *http agent poller*, *http poller*, *icmp pinger*, *internal poller ipmi manager*, *ipmi poller*, *java poller*, *lld manager*, *lld worker*, *odbc poller*, *poller*, *preprocessing manager*, *preprocessing worker*, *proxy group manager*, *proxy poller*, *self-monitoring*, *service manager*, *snmp poller*, *snmp trapper*, *task manager*, *timer*, *trapper*, *trigger housekeeper*, *unreachable poller*, *vmware collector*;
  - for proxy processes: *agent poller*, *apm manager*, *apm worker*, *availability manager*, *browser poller*, *configuration syncer*, *data sender*, *discovery manager*, *discovery worker*, *history syncer*, *housekeeper*, *http agent poller*, *http poller*, *icmp pinger*, *internal poller ipmi manager*, *ipmi poller*, *java poller*, *odbc poller*, *poller*, *preprocessing manager*, *preprocessing worker*, *self-monitoring*, *snmp poller*, *snmp trapper*, *task manager*, *trapper*, *unreachable poller*, *vmware collector*;
- **mode**:
  - *avg* - average value for all processes of a given type (default);
  - *count* - returns number of forks for a given process type, <state> should not be specified;
  - *max* - maximum value;
  - *min* - minimum value;
  - *<process number>* - process number (between 1 and the number of pre-forked instances; for example, if 4 trappers are running, the value is between 1 and 4);
- **state**:
  - *busy* - process is in busy state, for example, the processing request (default);
  - *idle* - process is in idle state doing nothing.

Comments:

- If <mode> is a Zabbix process number that is not running (for example, with 5 pollers running the <mode> is specified to be 6), such an item will turn unsupported.
- Minimum and maximum refers to the usage percentage for a single process. So if in a group of 3 pollers usage percentages per process were 2, 18 and 66, min would return 2 and max would return 66.
- Processes report what they are doing in shared memory and the self-monitoring process summarizes that data each second. State changes (busy/idle) are registered upon change - thus a process that becomes busy registers as such and doesn't change or update the state until it becomes idle. This ensures that even fully hung processes will be correctly registered as 100% busy.
- Currently, "busy" means "not sleeping", but in the future additional states might be introduced - waiting for locks, performing database queries, etc. Note that asynchronous pollers are considered busy if they have reached the limit set by the MaxConcurrentChecksPerPoller server/proxy configuration parameter.
- On Linux and most other systems, resolution is 1/100 of a second.

Examples:

```
zabbix[process,poller,avg,busy] #the average time of poller processes spent doing something during the las
zabbix[process,"icmp pinger",max,busy] #the maximum time spent doing something by any ICMP pinger process
zabbix[process,"history syncer",2,busy] #the time spent doing something by history syncer number 2 during
zabbix[process,trapper,count] #the amount of currently running trapper processes
```

zabbix[proxy,<name>,<param>]

 The information about Zabbix proxy.
 Return value: *Integer*.

Parameters:

- **name** - the proxy name;
- **param**:
  - *lastaccess* - the timestamp of the last heartbeat message received from proxy;
  - *delay* - how long the collected values are unsent; calculated as "proxy delay" + ("current server time" - "proxy lastaccess"), where "proxy delay" is the difference between the current proxy time and the timestamp of the oldest unsent value on proxy.

Comments:

- This item is not supported on hosts monitored by a proxy or proxy group.
- The `fuzzytime()` function can be used to check the availability of proxy.

Example:

```
zabbix[proxy,"Germany",lastaccess] #the timestamp of the last heartbeat message received from "Germany" pr
```

zabbix[proxy,discovery]

 The list of Zabbix proxies with name, mode, encryption, compression, version, last seen, host count, item count, required values per second (vps), version status (current/outdated/unsupported), timeouts by item type, proxy group name (if proxy belongs to group), state (unknown/offline/online).
 Return value: *JSON object*.

Comment:

- This item is not supported on hosts monitored by a proxy or proxy group.

Example of return values:

```
[
  {
    "name":"Zabbix proxy",
    "passive":false,
    "unencrypted":true,
    "psk":false,
    "cert":false,
    "items":22,
    "compression":true,
    "version":"8.2.0",
    "compatibility":1,
    "last_seen":1,
    "hosts":1,
    "requiredperformance":0.333333,
    "timeouts": {
        "zabbix_agent":3,
        "simple_check":3,
        "snmp_check":3,
        "external_check":3,
        "db_monitor":3,
        "http_agent":3,
        "ssh_agent":3,
        "telnet_agent":3,
        "script":3,
        "browser":60
    },
    "state":"online"
  }
]
```

zabbix[proxy group,<name>,available]

 The number of online proxies in a proxy group.
 Return value: *Integer*.

Parameters:

- **name** - the proxy group name.

Comment:

- This item is not supported on hosts monitored by a proxy or proxy group.

zabbix[proxy group,<name>,pavailable]

 The percentage of online proxies in a proxy group.
 Return value: *Float*.

Parameters:

- **name** - the proxy group name.

Comment:

- This item is not supported on hosts monitored by a proxy or proxy group.

zabbix[proxy group,<name>,proxies]

 The list of Zabbix proxies in a proxy group with name, mode, encryption, compression, version, last seen, host count, item count, required values per second (vps), version status (current/outdated/unsupported), timeouts, proxy group name, state (unknown/offline/online).
 Return value: *JSON*.

Parameters:

- **name** - the proxy group name.

Comment:

- This item is not supported on hosts monitored by a proxy or proxy group.

zabbix[proxy group,<name>,state]

 The state of a proxy group.
 Return value: *0* - unknown; *1* - offline; *2* - recovering; *3* - online; *4* - degrading.

Parameters:

- **name** - the proxy group name.

Comment:

- This item is not supported on hosts monitored by a proxy or proxy group.

zabbix[proxy group,discovery]

 Returns a list of proxy groups with configuration data and real-time data. Configuration data include the proxy group name, failover delay, and the minimum number of online proxies required. Real-time data include the proxy group state (see comments for details), the number of online proxies, and the percentage of online proxies.
 Return value: *JSON*.

Comments:

- This item is not supported on hosts monitored by a proxy or proxy group.
- This item does not return groupless proxies.
- If there is an invalid value for `failover_delay` or `min_online`, then a special value −1 is reported to indicate that. Invalid values might occur if macros are used for configuration and the macros cannot be expanded to a valid value.
- The proxy group state is reported as an integer: *0* - unknown; *1* - offline; *2* - recovering; *3* - online; *4* - degrading.

Example of return values:

```
{
    "groups": [
       { "name": "Riga", "failover_delay": 60, "min_online": 1 },
       { "name": "Tokyo", "failover_delay": 60, "min_online": 2 },
       { "name": "Porto Alegre", "failover_delay": 60, "min_online": 3 }
    ],
    "details": {
        "Riga": { "state": 3, "available": 10, "pavailable": 20 },
        "Tokyo": { "state": 3, "available": 10, "pavailable": 20 },
        "Porto Alegre": { "state": 1, "available": 0, "pavailable": 0 }
    }
}
```

zabbix[proxy_buffer,buffer,<mode>]

 The proxy memory buffer usage statistics.
 Return values: *Integer* (for size); *Float* (for percentage).

Parameters:

- **mode**:
  - *total* - the total size of buffer (can be used to check if memory buffer is enabled);
  - *free* - the size of free buffer;
  - *pfree* - the percentage of free buffer;
  - *used* - the size of used buffer;
  - *pused* - the percentage of used buffer.

Comments:

- Returns a 'Proxy memory buffer is disabled' error when the memory buffer is disabled;
- This item is not supported on the Zabbix server.

zabbix[proxy_buffer,state,changes]

 Returns the number of state changes between disk/memory buffer modes since start.
 Return values: *Integer*; *0* - the memory buffer is disabled.

Comments:

- Frequent state changes indicate that either the memory buffer size or age must be increased.
- If the memory buffer state is monitored infrequently (for example, once a minute) then the buffer might flip its state without it being registered.

zabbix[proxy_buffer,state,current]

 Returns the current working state where the new data is being stored.
 Return values: *0* - disk; *1* - memory.

Comment:

- 0 is also returned when the memory buffer is disabled.

zabbix[proxy_history]

 The number of values in the proxy history table waiting to be sent to the server.
 Return values: *Integer*.

Comment:

- This item is not supported on the Zabbix server.

zabbix[queue,<from>,<to>]

 The number of monitored items in the queue which are delayed at least by `<from>` seconds, but less than `<to>` seconds.
 Return value: *Integer*.

Parameters:

- **from** - delayed by at least (default is 6 seconds);
- **to** - delayed by at most (default is infinity).

Comments:

- Time suffixes (`s,m,h,d,w`) are supported in the parameters.

Attention:
The Zabbix proxy is not aware of maintenance periods; see Calculation of queues during maintenance for details.

zabbix[rcache,<cache>,<mode>]

 The availability statistics of the Zabbix configuration cache.
 Return values: *Integer* (for size); *Float* (for percentage).

Parameters:

- **cache** - *buffer*;
- **mode**:
  - *total* - the total size of buffer;
  - *free* - the size of free buffer;
  - *pfree* - the percentage of free buffer;
  - *used* - the size of used buffer;
  - *pused* - the percentage of used buffer.

zabbix[requiredperformance]

 The required performance of Zabbix server or Zabbix proxy, in new values per second expected.
 Return value: *Float*.

Comments:

- Approximately correlates with "Required server performance, new values per second" in *Reports* > *System information*.

zabbix[stats,<ip>,<port>]

 Returns the internal metrics of a Zabbix server or proxy. If `<ip>` and `<port>` are specified, metrics are retrieved from the remote instance; otherwise, from the local instance.
 Return values: *JSON object*.

Parameters:

- **ip** - the IP/DNS/network mask list of servers/proxies to be remotely queried (default is 127.0.0.1);
- **port** - the port of server/proxy to be remotely queried (default is 10051).

Comments:

- The stats request will only be accepted from the addresses listed in the `StatsAllowedIP` server/proxy parameter on the target instance.
- A selected set of internal metrics is returned by this item. For details, see Remote monitoring of Zabbix stats.

zabbix[stats,<ip>,<port>,queue,<from>,<to>]

 Returns the internal queue metrics (see `zabbix[queue,<from>,<to>]`) of a Zabbix server or proxy. If `<ip>` and `<port>` are specified, metrics are retrieved from the remote instance; otherwise, from the local instance.
 Return values: *JSON object*.

Parameters:

- **ip** - the IP/DNS/network mask list of servers/proxies to be remotely queried (default is `127.0.0.1`);
- **port** - the port of server/proxy to be remotely queried (default is 10051);
- **from** - delayed by at least (default is 6 seconds);
- **to** - delayed by at most (default is infinity).

Comments:

- The stats request will only be accepted from the addresses listed in the `StatsAllowedIP` server/proxy parameter on the target instance.
- A selected set of internal metrics is returned by this item. For details, see Remote monitoring of Zabbix stats.

Attention:
The Zabbix proxy is not aware of maintenance periods; see Calculation of queues during maintenance for details.

zabbix[tcache,cache,<parameter>]

 The effectiveness statistics of the Zabbix trend function cache.
 Return values: *Integer* (for size); *Float* (for percentage).

Parameters:

- **parameter**:
  - *all* - total cache requests (default);
  - *hits* - cache hits;
  - *phits* - percentage of cache hits;
  - *misses* - cache misses;
  - *pmisses* - percentage of cache misses;
  - *items* - the number of cached items;
  - *requests* - the number of cached requests;
  - *pitems* - percentage of cached items from cached items + requests. Low percentage most likely means that the cache size can be reduced.
  - *paccessed* - percentage of accessed entries in last 24 hours.

Comments:

- This item is not supported on Zabbix proxy.

zabbixtriggers

 The number of enabled triggers in Zabbix database, with all items enabled on enabled hosts.
 Return value: *Integer*.

Comments:

- This item is not supported on Zabbix proxy.

zabbix[uptime]

 The uptime of the Zabbix server or proxy process in seconds.
 Return value: *Integer*.

zabbix[vcache,buffer,<mode>]

 The availability statistics of the Zabbix value cache.
 Return values: *Integer* (for size); *Float* (for percentage).

Parameters:

- **mode**:
  - *total* - the total size of buffer;
  - *free* - the size of free buffer;
  - *pfree* - the percentage of free buffer;
  - *used* - the size of used buffer;
  - *pused* - the percentage of used buffer.

Comments:

- This item is not supported on Zabbix proxy.

zabbix[vcache,cache,<parameter>]

 The effectiveness statistics of the Zabbix value cache.
 Return value: *Integer*. If <parameter> is set to *mode*, the return value is: *0* - normal mode; *1* - low memory mode.

Parameters:

- **parameter**:
  - *requests* - the total number of requests;
  - *hits* - the number of cache hits (history values taken from the cache);
  - *misses* - the number of cache misses (history values taken from the database);
  - *mode* - the value cache operating mode.

Comments:

- Once the low-memory mode has been switched on, the value cache will remain in this state for 24 hours, even if the problem that triggered this mode is resolved sooner.
- You may use this key with the *Change per second* preprocessing step in order to get values-per-second statistics.
- This item is not supported on Zabbix proxy.

zabbixversion

 The version of Zabbix server or proxy.
 Return value: *String*. For example: `8.2.0`.

zabbix[vmware,buffer,<mode>]

 The availability statistics of the Zabbix vmware cache.
 Return values: *Integer* (for size); *Float* (for percentage).

Parameters:

- **mode**:
  - *total* - the total size of buffer;
  - *free* - the size of free buffer;
  - *pfree* - the percentage of free buffer;
  - *used* - the size of used buffer;
  - *pused* - the percentage of used buffer.

zabbix[vps,written]

 The total number of history values written to database.
 Return value: *Integer*.

zabbix[wcache,<cache>,<mode>]

 The statistics and availability of the Zabbix write cache.
 Return values: *Integer* (for number/size); *Float* (for percentage).

Parameters:

- **cache** - *values*, *history*, *index*, or *trend*;
- **mode**:
  - with *values* cache parameter:
    - *all* - the total number of values processed by Zabbix server/proxy, except unsupported items (counter) (default);
    - *float* - the number of processed float values (counter);
    - *uint* - the number of processed unsigned integer values (counter);
    - *str* - the number of processed character/string values (counter);
    - *log* - the number of processed log values (counter);
    - *text* - the number of processed text values (counter);
    - *bin* - the number of processed binary values (counter);
    - *json* - the number of processed JSON values (counter);
    - *not supported* - the number of times item processing resulted in item becoming unsupported or keeping that state (counter);
  - with *history*, *index*, *trend* cache parameters:
    - *pfree* - the percentage of free buffer (default);
    - *total* - the total size of buffer;
    - *free* - the size of free buffer;
    - *used* - the size of used buffer;
    - *pused* - the percentage of used buffer.

Comments:

- The `trend` cache parameter is not supported with Zabbix proxy.
- The history cache is used to store item values. A low number indicates performance problems on the database side.

- The history index cache is used to index the values stored in the history cache.
- After the history cache is filled and then cleared, the history index cache will still keep some data. This behavior is expected and helps the system run more efficiently by avoiding the extra processing required to constantly resize the memory.
- The trend cache stores the aggregate for the current hour for all items that receive data.
- You may use the `zabbix[wcache,values]` key with the *Change per second* preprocessing step in order to get values-per-second statistics.

Zabbix trapper

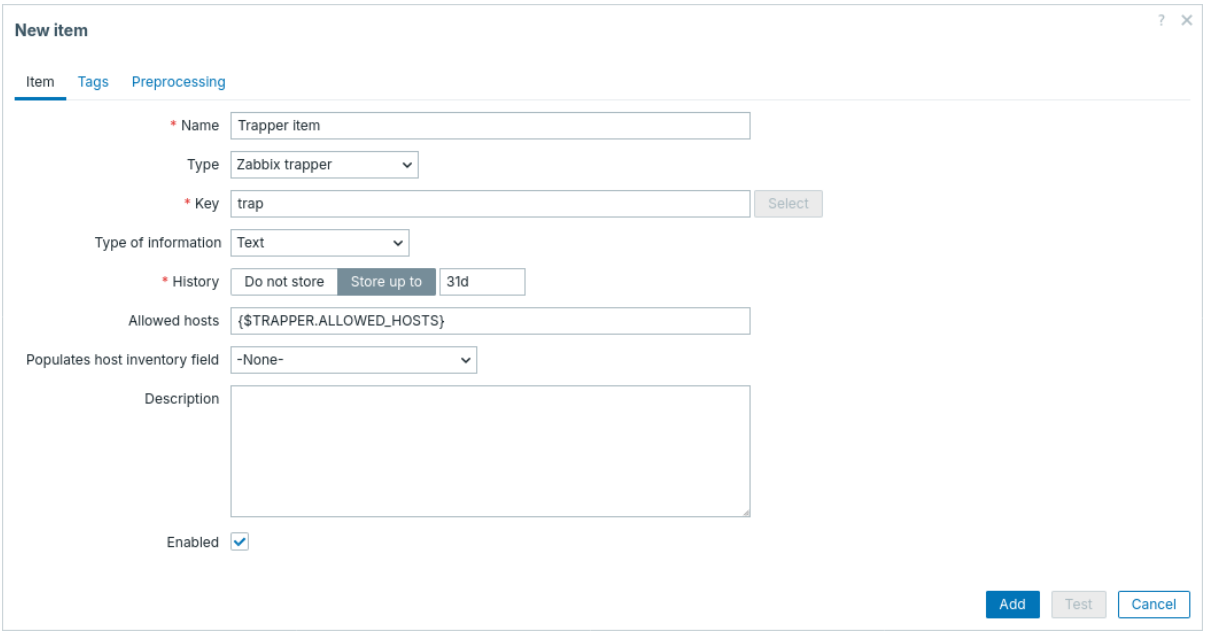
Overview

Trapper items accept incoming data instead of querying for it. This is useful for any data you want to send to Zabbix.

Configuration

To configure a trapper item:

1. Go to *Data collection* > *Hosts*.
2. Click *Items* in the row of the host.
3. Click *Create item*.
4. Enter item parameters in the configuration form.



All mandatory input fields are marked with a red asterisk.

The fields that require specific information for trapper items are:

Type	Select "Zabbix trapper".
Key	Enter a key that will be used to recognize the item when sending data to Zabbix server.
Type of information	Select the type of information that will correspond to the format of data that will be sent.

Allowed hosts	List of comma-delimited IP addresses (optionally in CIDR notation) or DNS names. The default value is the user macro {$TRAPPER.ALLOWED_HOSTS}. By default, {$TRAPPER.ALLOWED_HOSTS} is 127.0.0.1,::1. You can either enter addresses directly in the field or change the macro value. If left empty, incoming connections will be denied from all hosts. If specified, incoming connections will be accepted only from the hosts listed here. If IPv6 support is enabled then '127.0.0.1', '::127.0.0.1', '::ffff:127.0.0.1' are treated equally and '::/0' will allow any IPv4 or IPv6 address. '0.0.0.0/0' can be used to allow any IPv4 address. Note that "IPv4-compatible IPv6 addresses" (0000::/96 prefix) are supported but deprecated by RFC4291. Example: 127.0.0.1, 192.168.1.0/24, 192.168.3.1-255, 192.168.1-10.1-255, ::1,2001:db8::/32, mysqlserver1, zabbix.example.com, {HOST.HOST} Spaces, user macros, and host macros {HOST.HOST}, {HOST.NAME}, {HOST.IP}, {HOST.DNS}, {HOST.CONN} are supported.

Note:
Before sending values, you may have to wait up to 60 seconds after saving the item until Zabbix server picks up the changes from a configuration cache update.

Sending data

Sending data to Zabbix server or proxy is possible using the Zabbix sender utility or Zabbix sender protocol. Sending data to Zabbix server is also possible using the `history.push` API method.

Zabbix sender

For sending data to Zabbix server or proxy using the Zabbix sender utility, you could run the following command to send the "test value":

```
zabbix_sender -z <server IP address> -p 10051 -s "New host" -k trap -o "test value"
```

To send the "test value", the following command options are used:

- `-z` to specify Zabbix server IP address
- `-p` to specify Zabbix server port number (10051 by default)
- `-s` to specify the host (make sure to use the technical instead of the visible host name)
- `-k` to specify the key of the item configured in the trapper item
- `-o` to specify the value to send

Attention:
The Zabbix trapper process does not expand macros used in the item key to check the corresponding item key existence for the targeted host.

For more information on the communication between Zabbix sender and Zabbix server or proxy, see *Zabbix sender protocol*.

history.push

For sending data to Zabbix server using the `history.push` API method, you could make the following HTTP POST request containing some test values:

```
curl --request POST \
  --url 'https://example.com/zabbix/api_jsonrpc.php' \
  --header 'Authorization: Bearer 0424bd59b807674191e7d77572075f33' \
  --header 'Content-Type: application/json-rpc' \
  --data '{"jsonrpc":"2.0","method":"history.push","params":[{"itemid":10600,"value":"test value 1"},{"ite
```

If the request is correct, the response returned by API could look as follows:

```
{
    "jsonrpc": "2.0",
    "result": {
        "response": "success",
```

```
        "data": [
            {
                "itemid": "10600"
            },
            {
                "itemid": "10601",
                "error": "Item is disabled."
            },
            {
                "error": "No permissions to referred object or it does not exist."
            }
        ]
    },
    "id": 1
}
```


Errors in response data indicate that sending data for specific items has failed validation by Zabbix server. This can happen for the following reasons:

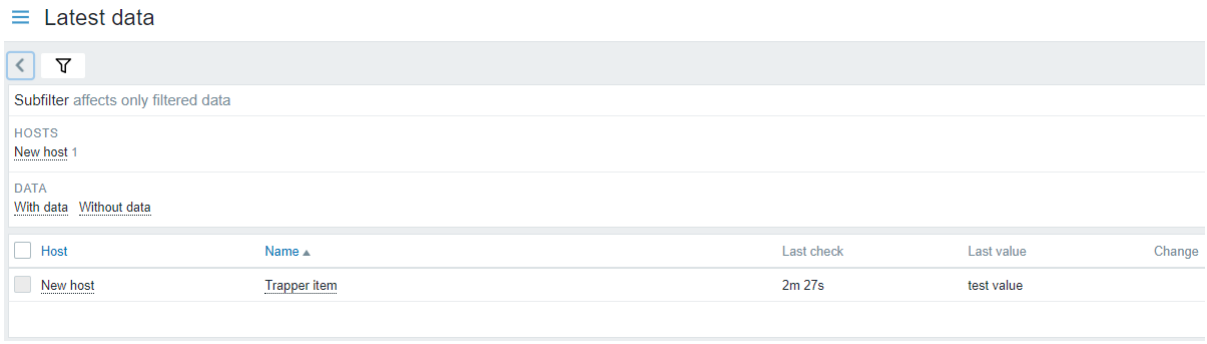
- the user sending the data has no *read* permission to the item's host.
- the host is disabled or in maintenance without data collection.
- the item does not exist or is not yet included in the server configuration cache.
- the item is disabled or its type is other than Zabbix trapper or HTTP agent (with trapping enabled).
- the user's IP or DNS is not set in the item's *Allowed hosts* list.
- another item has a value with a duplicate timestamp on the nanosecond level.

The absence of errors indicates that the values sent have been accepted for processing, which includes preprocessing (if any), trigger processing, and saving to the database. Note that the processing of an accepted value may also fail (for example, during preprocessing), resulting in the value being discarded.

For more information on how to work with Zabbix API, see API.

Displaying data

Once data is sent, you can navigate to *Monitoring* > *Latest data* to see the result:



Note:
If a single numeric value is sent, the data graph will show a horizontal line to the left and right of the value's time point.

External check

Overview

An external check is a check executed by Zabbix server by running a shell script or a binary. However, when hosts are monitored by a Zabbix proxy, the external checks are executed by the proxy.

External checks do not require any agent running on a host being monitored.

The syntax of the item key is:

`script[<parameter1>,<parameter2>,...]`

Where:

ARGUMENT	DEFINITION
script	Name of a shell script or a binary.
parameter(s)	Optional command line parameters.

If you don't want to pass any parameters to the script you may use:

```
script[] or
script
```

Zabbix server or proxy will search the directory specified for external scripts and execute the command (see `ExternalScripts` parameter in Zabbix server/proxy configuration file). The command will be executed under the same user as Zabbix server/proxy, so any access permissions or environment variables should be handled in a wrapper script, if necessary. Permissions on the command should also allow that user to execute it. Only commands in the specified directory are available for execution.

Warning:
Do not overuse external checks, as each script requires starting a fork process by Zabbix server/proxy, and running many scripts can significantly decrease Zabbix performance.

Usage example

Executing the script `check_oracle.sh` with the first parameters '-h'. The second parameter will be replaced by IP address or DNS name, depending on the selection in the host properties.

```
check_oracle.sh["-h","{HOST.CONN}"]
```

Assuming host is configured to use IP address, Zabbix server/proxy will execute:

```
check_oracle.sh '-h' '192.168.1.4'
```

External check result

The return value of an external check is a standard output together with a standard error produced by the check.

Attention:
An item that returns text (character, log, or text type of information) will not become unsupported in case of a standard error output.

The return value is limited to 16MB (including trailing whitespace that is truncated); database limits also apply.

If the requested script is not found or Zabbix server/proxy has no permissions to execute it, the item will become unsupported and a corresponding error message will be displayed.

In case of a timeout, the item will become unsupported, a corresponding error message will be displayed, and the process forked for the script will be terminated.

Database monitor

Overview

The *Database monitor* item type in the Zabbix frontend is used for ODBC monitoring (ODBC checks).

ODBC is a C programming language middle-ware API for accessing database management systems (DBMS). The ODBC concept was developed by Microsoft and later ported to other platforms.

Zabbix may query any database, which is supported by ODBC. To do that, Zabbix does not directly connect to the databases, but uses the ODBC interface and drivers set up in ODBC. This allows for more efficient monitoring of different databases for multiple purposes (for example, checking specific database queues, usage statistics, etc.).

Zabbix supports unixODBC, which is one of the most commonly used open source ODBC API implementations.

Attention:
See also: known issues for ODBC checks.

Installing unixODBC

The recommended way of installing unixODBC is to use the Linux operating system default package repositories. In the most popular Linux distributions, unixODBC is included in the package repository by default. If packages are not available, the source files can be obtained at the unixODBC homepage: http://www.unixodbc.org/download.html.

To install unixODBC, use the package manager for the system of your choice:

```
##### For Ubuntu/Debian systems:
apt install unixodbc unixodbc-dev

##### For RedHat/Fedora-based systems:
dnf install unixODBC unixODBC-devel

##### For SUSE-based systems:
zypper in unixODBC-devel
```

Attention:
The `unixodbc-dev` or `unixODBC-devel` package is necessary to compile Zabbix with unixODBC support. To enable ODBC support, Zabbix should be compiled with the following configuration option:

`--with-unixodbc[=ARG] # Use ODBC driver against unixODBC package.`

Installing unixODBC drivers

The unixODBC database driver should be installed for the database that will be monitored. For a list of supported databases and drivers, see the unixODBC homepage: http://www.unixodbc.org/drivers.html.

Note:
In some Linux distributions, database drivers are included in package repositories.

MySQL

To install the MySQL unixODBC database driver, use the package manager for the system of your choice:

```
##### For Ubuntu/Debian systems:
apt install odbc-mariadb

##### For RedHat/Fedora-based systems:
dnf install mariadb-connector-odbc

##### For SUSE-based systems:
zypper install mariadb-connector-odbc
```

To install the database driver without a package manager, please refer to MySQL documentation for `mysql-connector-odbc`, or MariaDB documentation for `mariadb-connector-odbc`.

PostgreSQL

To install the PostgreSQL unixODBC database driver, use the package manager for the system of your choice:

```
##### For Ubuntu/Debian systems:
apt install odbc-postgresql

##### For RedHat/Fedora-based systems:
dnf install postgresql-odbc

##### For SUSE-based systems:
zypper install psqlODBC
```

To install the database driver without a package manager, please refer to PostgreSQL documentation.

Oracle

To install the unixODBC database driver, please refer to Oracle documentation.

MSSQL

To install the MSSQL unixODBC database driver, use the package manager for the system of your choice:

```
##### For Ubuntu/Debian systems:
apt install tdsodbc

##### For RedHat/Fedora-based systems (EPEL packages: https://docs.fedoraproject.org/en-US/epel/):
dnf install epel-release
dnf install freetds
```

```
##### For SUSE-based systems:
zypper install libtdsodbc0
```

To install the database driver without a package manager, please refer to FreeTDS user guide.

Configuring unixODBC

To configure unixODBC, you must edit the `odbcinst.ini` and `odbc.ini` files. You can verify the location of these files by executing the following command:

```
odbcinst -j
```

The command result should contain information that is similar to the following:

```
unixODBC 2.3.9
DRIVERS............: /etc/odbcinst.ini
SYSTEM DATA SOURCES: /etc/odbc.ini
FILE DATA SOURCES..: /etc/ODBCDataSources
```

odbcinst.ini

The `odbcinst.ini` file lists the installed ODBC database drivers. If `odbcinst.ini` is missing, it is necessary to create it manually.

```
[TEST_MYSQL]
Description=ODBC for MySQL
Driver=/usr/lib/libmyodbc5.so
FileUsage=1
```

Parameter	Description
TEST_MYSQL	Database driver name.
Description	Database driver description.
Driver	Database driver library location.
FileUsage	Determines whether the database driver supports connecting to a database server without the support for accessing local files (0); supports reading data from files (1); supports writing data to files (2).
Threading	Thread serialization level. Supported for PostgreSQL. Since 1.6, if the driver manager is built with thread support, you may add another driver entry.

odbc.ini

The `odbc.ini` file is used to configure data sources. Note that the list of supported parameters depends on the database driver (for example, Oracle databases may use ServerName instead of Server, etc.).

```
[TEST_MYSQL]
Description=MySQL Test Database
Driver=mysql
Server=127.0.0.1
User=root
Password=
Port=3306
Socket=
Database=zabbix
```

Parameter	Description
TEST_MYSQL	Data source name (DSN).
Description	Data source description.
Driver	Database driver name (as specified in `odbcinst.ini`).
Server	Database server IP/DNS.
User	Database user for connection.
Password	Database user password.
Port	Database connection port.
Socket	Database connection socket.
Database	Database name.

For other possible configuration parameter options, see MySQL documentation.

The `odbc.ini` file for PostgreSQL may contain additional parameters:

```
[TEST_PSQL]
Description=PostgreSQL Test Database
Driver=postgresql
Username=zbx_test
Password=zabbix
Servername=127.0.0.1
Database=zabbix
Port=5432
ReadOnly=No
Protocol=8.0+
ShowOidColumn=No
FakeOidIndex=No
RowVersioning=No
ShowSystemTables=No
Fetch=Yes
BoolsAsChar=Yes
SSLmode=Require
ConnSettings=
```

Parameter	Description
ReadOnly	Specifies whether the database connection allows only read operations (`SELECT` queries) and restricts modifications (`INSERT`, `UPDATE`, and `DELETE` statements); useful for scenarios where data should remain unchanged.
Protocol	PostgreSQL backend protocol version (ignored when using SSL connections).
ShowOidColumn	Specifies whether to include Object ID (OID) in SQLColumns.
FakeOidIndex	Specifies whether to create a fake unique index on OID.
RowVersioning	Specifies whether to enable applications to detect if data has been modified by other users while you are attempting to update a row. Note that this parameter can speed up the update process, since, to update a row, every single column does not need to be specified in the `WHERE` clause.
ShowSystemTables	Specifies whether the database driver should treat system tables as regular tables in SQLTables; useful for accessibility, allowing visibility into system tables.
Fetch	Specifies whether the driver should automatically use declare cursor/fetch to handle `SELECT` statements and maintain a cache of 100 rows.
BoolsAsChar	Controls the mapping of Boolean types. If set to "Yes", Bools are mapped to `SQL_CHAR`; otherwise, they are mapped to `SQL_BIT`.
SSLmode	Specifies the SSL mode for the connection.
ConnSettings	Additional settings sent to the backend on connection.

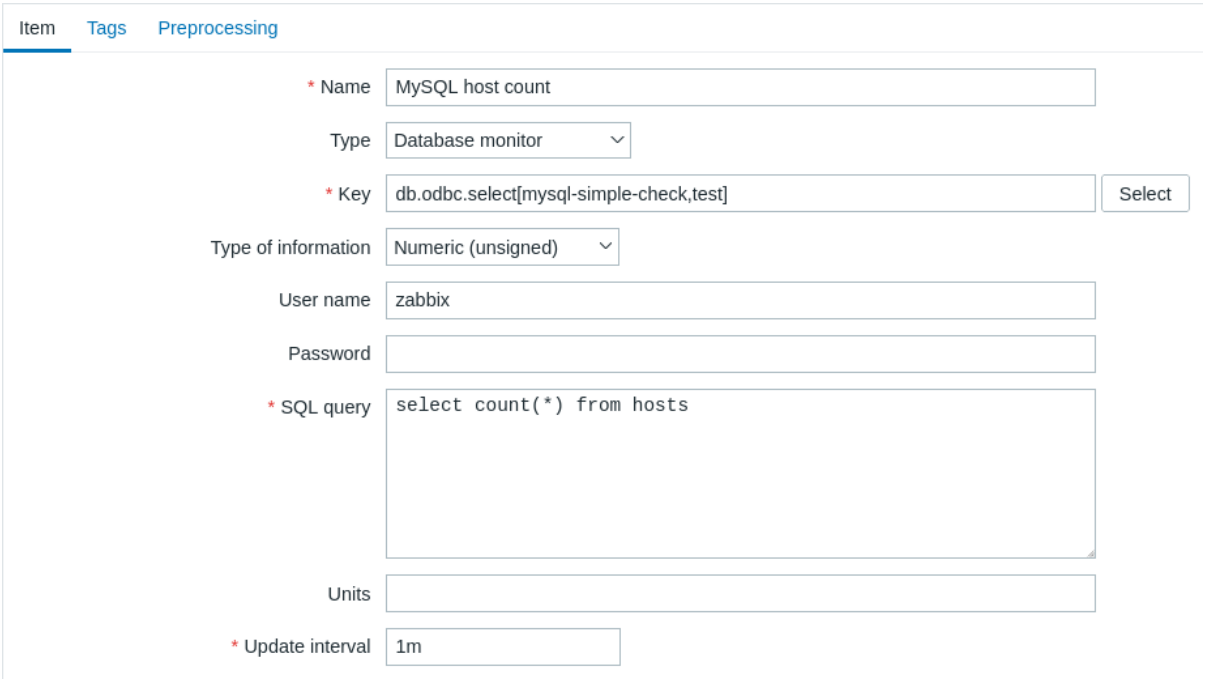
Testing ODBC connection

To test if the ODBC connection is working successfully, you can use the `isql` utility (included in the `unixODBC` package):

```
isql test
+---------------------------------------+
| Connected!                            |
|                                       |
| sql-statement                         |
| help [tablename]                      |
| quit                                  |
|                                       |
+---------------------------------------+
```

Item configuration in Zabbix frontend

Configure a **Database monitoring** item.



All mandatory input fields are marked with a red asterisk.

For database monitoring items, you must specify:

Type	Select "Database monitor" here.
Key	Enter one of the supported item keys: **db.odbc.select[]** - this item returns one value (the first column of the first row of the SQL query result); **db.odbc.get[]** - this item returns multiple rows/columns in JSON format; **db.odbc.discovery[]** - this item returns low-level discovery data.
User name	Enter the database user name (up to 255 characters). This parameter is optional if the database user name is specified in the `odbc.ini` file. If a connection string is used, and the *User name* field is not empty, then it is appended to the connection string as `UID=<user>`.
Password	Enter the database user password (up to 255 characters). This parameter is optional if the password is specified in the `odbc.ini` file. If a connection string is used, and the *Password* field is not empty, then it is appended to the connection string as `PWD=<password>`. Special characters are supported in this field. For all databases except Oracle, passwords containing % must be enclosed in curly brackets {}. The password will be appended to connection string after the username as, for example, `UID=<username>;PWD=P?;)*word`. To test the resulting string, you can run the following command: `isql -v -k 'Driver=libmaodbc.so;Database=zabbix;UID=zabbix;PWD=P?;)*word'`
SQL query	Enter the SQL query. Note that with `db.odbc.select[]`, the query must return one value only.
Type of information	Select the type of information that will be returned by the query here. If the type of information is selected incorrectly, the item will become unsupported.

Important notes

- Database monitoring items will become unsupported if no *odbc poller* processes are started in the server or proxy configuration. To activate ODBC pollers, set the `StartODBCPollers` parameter in Zabbix server configuration file or, for checks performed by proxy, in Zabbix proxy configuration file.
- The *Timeout* parameter value in the item configuration form is used as the ODBC login timeout and the query execution timeout. Note that these timeout settings might be ignored if the installed ODBC driver does not support them.
- The SQL command must return a result set like any query using the `select` statement. The query syntax will depend on the RDBMS which will process them. The syntax of request to a storage procedure must be started with the `call` keyword.

Item key details

Parameters without angle brackets are mandatory. Parameters marked with angle brackets **< >** are optional.

db.odbc.select[<unique short description>,<dsn>,<connection string>]

 Returns one value, that is, the first column of the first row of the SQL query result.
 Return value: Depending on the SQL query.

Parameters:

- **unique short description** - a unique short description to identify the item (for use in triggers, etc.);
- **dsn** - the data source name (as specified in `odbc.ini`);
- **connection string** - the connection string (may contain driver-specific arguments).

Comments:

- Although `dsn` and `connection string` are optional parameters, at least one of them is required; if both are defined, `dsn` will be ignored.
- If a query returns more than one column, only the first column is read. If a query returns more than one line, only the first line is read.

db.odbc.get[<unique short description>,<dsn>,<connection string>]

 Transforms the SQL query result into a JSON array.
 Return value: *JSON object*.

Parameters:

- **unique short description** - a unique short description to identify the item (for use in triggers, etc.);
- **dsn** - the data source name (as specified in `odbc.ini`);
- **connection string** - the connection string (may contain driver-specific arguments).

Comments:

- Although `dsn` and `connection string` are optional parameters, at least one of them is required; if both are defined, `dsn` will be ignored.
- Multiple rows/columns in JSON format may be returned. This item may be used as a master item that collects all data in one system call, while JSONPath preprocessing may be used in dependent items to extract individual values. For more information, see an example of the returned format, used in low-level discovery.

Example:

```
# Connection for MySQL ODBC driver 5:
db.odbc.get[MySQL example,,"Driver=/usr/local/lib/libmyodbc5a.so;Database=master;Server=127.0.0.1;Port=330
```

db.odbc.discovery[<unique short description>,<dsn>,<connection string>]

 Transforms the SQL query result into a JSON array, used for low-level discovery. The column names from the query result are turned into low-level discovery macro names paired with the discovered field values. These macros can be used in creating item, trigger, etc. prototypes.
 Return value: *JSON object*.

Parameters:

- **unique short description** - a unique short description to identify the item (for use in triggers, etc.);
- **dsn** - the data source name (as specified in `odbc.ini`);
- **connection string** - the connection string (may contain driver-specific arguments).

Comments:

- Although `dsn` and `connection string` are optional parameters, at least one of them is required; if both are defined, `dsn` will be ignored.

Error messages

ODBC error messages are structured into fields to provide detailed information. For example, an error message might look like this:

```
Cannot execute ODBC query: [SQL_ERROR]:[42601][7][ERROR: syntax error at or near ";"; Error while executin
```

- "`Cannot execute ODBC query`" - Zabbix message
- "`[SQL_ERROR]`" - ODBC return code
- "`[42601]`" - SQLState
- "`[7]`" - Native error code
- "`[ERROR: syntax error at or near ";"; Error while executing the query]`" - Native error message

Note that the error message length is limited to 2048 bytes, so the message can be truncated. If there is more than one ODBC diagnostic record, Zabbix tries to concatenate them (separated with |) as far as the length limit allows.

HTTP agent

Overview

This item type allows data polling using the HTTP/HTTPS protocol. Trapping is also possible using the Zabbix sender utility or Zabbix sender protocol (for sending data to Zabbix server or proxy), or using the `history.push` API method (for sending data to Zabbix server).

HTTP item checks are executed by Zabbix server. However, when hosts are monitored by a Zabbix proxy, HTTP item checks are executed by the proxy.

HTTP item checks do not require any agent running on a host being monitored.

HTTP agent supports both HTTP and HTTPS. Zabbix will optionally follow redirects (see the *Follow redirects* option below). Maximum number of redirects is hard-coded to 10 (using cURL option `CURLOPT_MAXREDIRS`).

> **Attention:**
> Zabbix server/proxy must be initially configured with cURL (libcurl) support.

HTTP checks are executed asynchronously - it is not required to receive the response to one request before other checks are started. DNS resolving is asynchronous as well.

The maximum concurrency of asynchronous checks is 1000 (defined by MaxConcurrentChecksPerPoller).

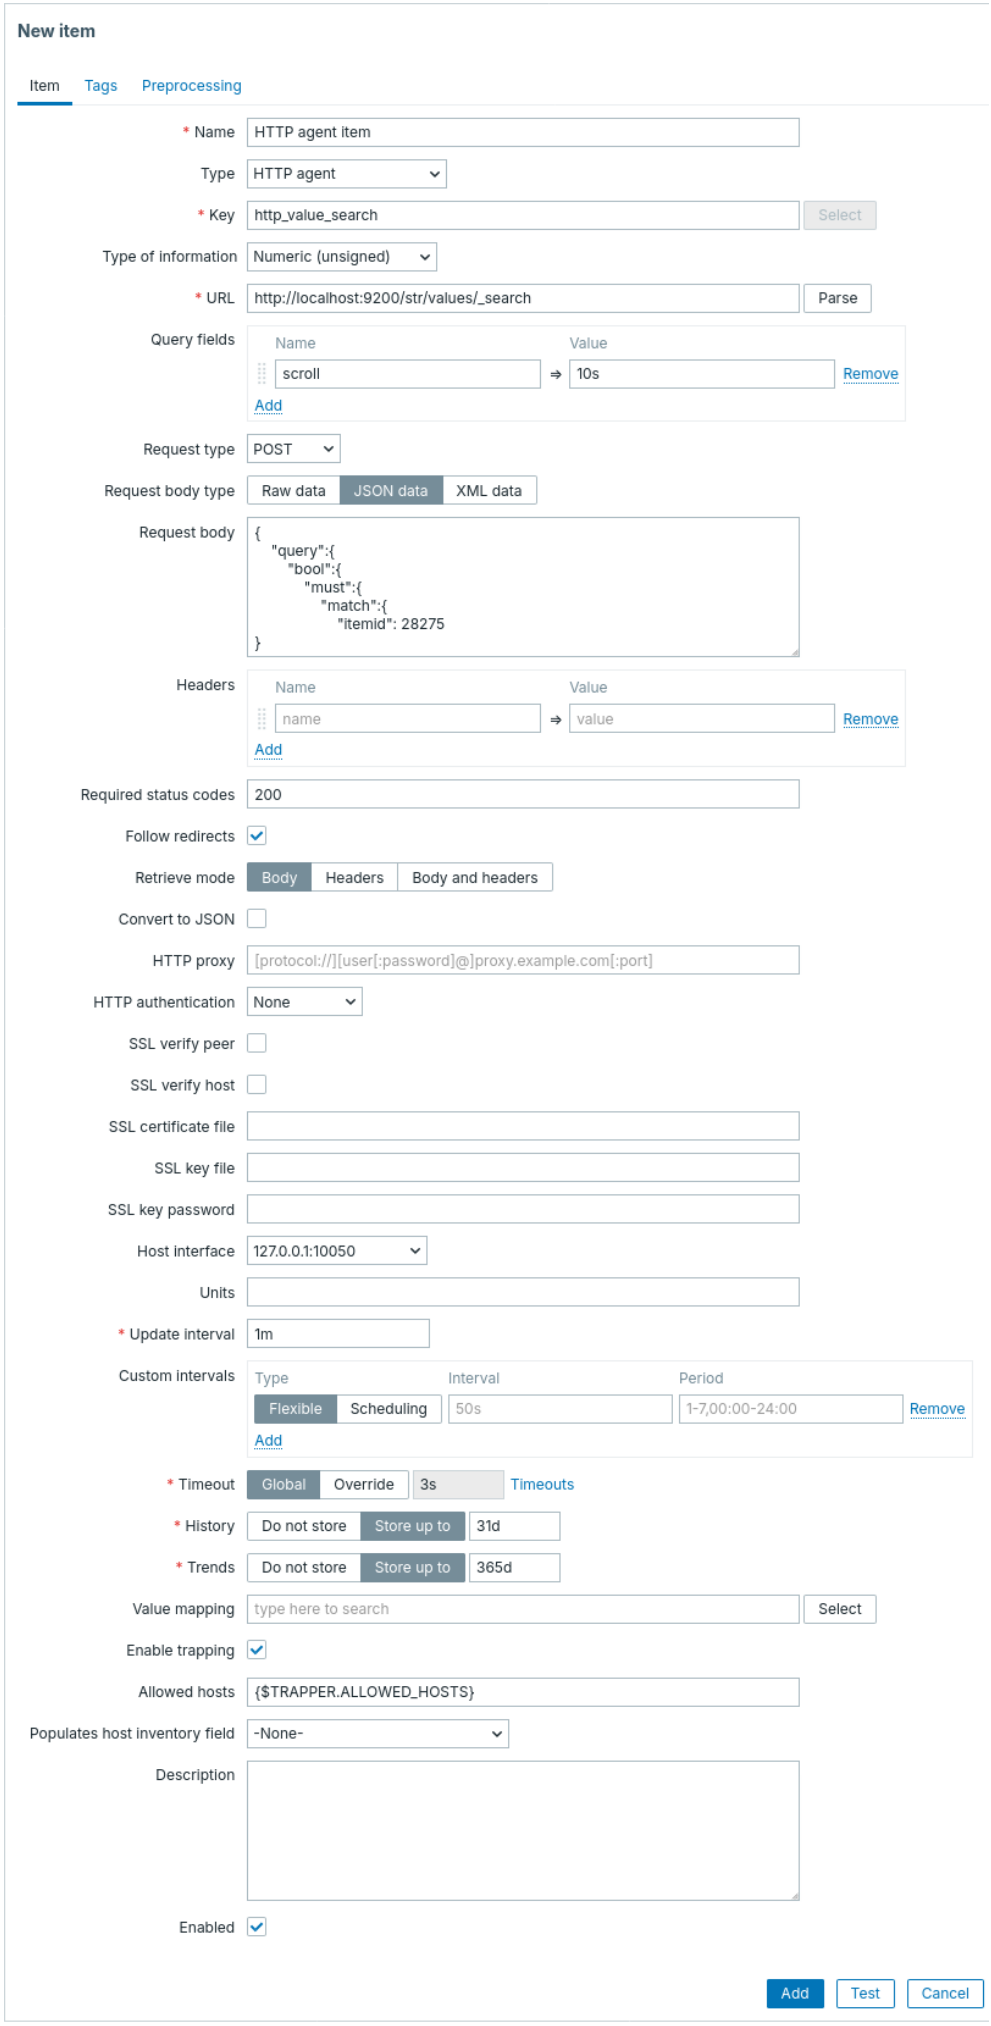
The number of asynchronous HTTP agent pollers is defined by the StartHTTPAgentPollers parameter.

The persistent connections cURL feature has been added to HTTP agent checks since Zabbix 7.0.

Configuration

To configure an HTTP item:

1. Go to *Data collection* > *Hosts*.
2. Click *Items* in the row of the host.
3. Click *Create item*.
4. Enter parameters of the item in the form.


New item
Item
Tags
Preprocessing
* Name
HTTP agent item
Type
HTTP agent
* Key
http_value_search
Select
Type of information
Numeric (unsigned)
* URL
http://localhost:9200/str/values/_search
Parse
Query fields
Name
Value
scroll
10s
Remove
Add
Request type
POST
Request body type
Raw data
JSON data
XML data
Request body
{
"query":{
"bool":{
"must":{
"match":{
"itemid": 28275
}
Headers
Name
Value
name
value
Remove
Add
Required status codes
200
Follow redirects
Retrieve mode
Body
Headers
Body and headers
Convert to JSON
HTTP proxy
[protocol://][user[:password]@]proxy.example.com[:port]
HTTP authentication
None
SSL verify peer
SSL verify host
SSL certificate file
SSL key file
SSL key password
Host interface
127.0.0.1:10050
Units
* Update interval
1m
Custom intervals
Type
Interval
Period
Flexible
Scheduling
50s
1-7,00:00-24:00
Remove
Add
* Timeout
Global
Override
3s
Timeouts
* History
Do not store
Store up to
31d
* Trends
Do not store
Store up to
365d
Value mapping
type here to search
Select
Enable trapping
Allowed hosts
{$TRAPPER.ALLOWED_HOSTS}
Populates host inventory field
-None-
Description
Enabled
Add
Test
Cancel

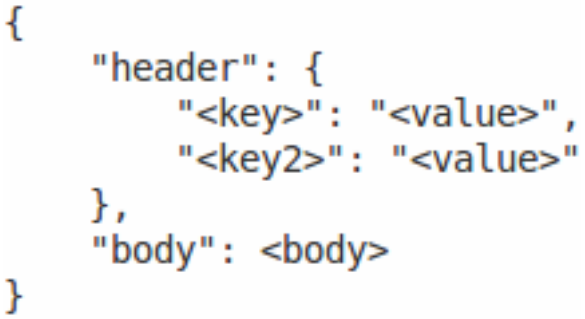
All mandatory input fields are marked with a red asterisk.

The fields that require specific information for HTTP items are:

Parameter	Description
Type	Select **HTTP agent** here.
Key	Enter a unique item key.

URL	URL to connect to and retrieve data. For example:
	`https://www.example.com`
	`http://www.example.com/download`
	Domain names can be specified in Unicode characters. They are automatically punycode-converted to ASCII when executing the HTTP check.
	The *Parse* button can be used to separate optional query fields (like `?key=value&limit=10`) from the URL, moving the attributes and values into *Query fields* for automatic URL-encoding.
	For security reasons, do not include actual usernames or passwords in URLs. Use user macros instead to avoid exposing credentials.
	Limited to 2048 characters.
	Supported macros: {HOST.IP}, {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.NAME}, {HOST.PORT}, {ITEM.ID}, {ITEM.KEY}, {ITEM.KEY.ORIG}, user macros, low-level discovery macros.
	This sets the CURLOPT_URL cURL option.

Query fields	Variables for the URL (see above). Specified as attribute and value pairs. Values are URL-encoded automatically. Values from macros are resolved and then URL-encoded automatically. Supported macros: {HOST.IP}, {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.NAME}, {HOST.PORT}, {ITEM.ID}, {ITEM.KEY}, {ITEM.KEY.ORIG}, user macros, low-level discovery macros. This sets the CURLOPT_URL cURL option.
Request type	Select request method type: *GET*, *POST*, *PUT* or *HEAD*.
Request body type	Select the request body type: **Raw data** - custom HTTP request body, macros are substituted but no encoding is performed. **JSON data** - HTTP request body in JSON format. Macros can be used as string, number, true and false; macros used as strings must be enclosed in double quotes. Values from macros are resolved and then escaped automatically. If "Content-Type" is not specified in headers then it will default to "Content-Type: application/json". **XML data** - HTTP request body in XML format. Macros can be used as a text node, attribute or CDATA section. Values from macros are resolved and then escaped automatically in a text node and attribute. If "Content-Type" is not specified in headers then it will default to "Content-Type: application/xml".
Request body	Enter the request body. Supported macros: {HOST.IP}, {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.NAME}, {HOST.PORT}, {ITEM.ID}, {ITEM.KEY}, {ITEM.KEY.ORIG}, user macros, low-level discovery macros.
Headers	Custom HTTP headers that will be sent when performing a request. Specified as attribute and value pairs. Supported macros: {HOST.IP}, {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.NAME}, {HOST.PORT}, {ITEM.ID}, {ITEM.KEY}, {ITEM.KEY.ORIG}, user macros, low-level discovery macros. This sets the CURLOPT_HTTPHEADER cURL option.
Required status codes	List of expected HTTP status codes. If Zabbix gets a code which is not in the list, the item will become unsupported. If empty, no check is performed. For example: 200,201,210-299. Supported macros in the list: user macros, low-level discovery macros. This uses the CURLINFO_RESPONSE_CODE cURL option.
Follow redirects	Mark the checkbox to follow HTTP redirects. This sets the CURLOPT_FOLLOWLOCATION cURL option.
Retrieve mode	Select the part of response that must be retrieved: **Body** - body only **Headers** - headers only **Body and headers** - body and headers

Parameter	Description
Convert to JSON	Headers are saved as attribute and value pairs under the "header" key if *Retrieve mode* is not set to *Body*. If 'Content-Type: application/json' is encountered then body is saved as an object; otherwise, it is stored as string, for example: 
HTTP proxy	You can specify an HTTP proxy to use, using the format `[protocol://][username[:password]@]proxy.example.com[:port]`. The optional `protocol://` prefix may be used to specify alternative proxy protocols (e.g., `https`, `socks4`, `socks5`; see documentation; the protocol prefix support was added in cURL 7.21.7). With no protocol specified, the proxy will be treated as an HTTP proxy. If you specify the wrong protocol, the connection will fail and the item will become unsupported. By default, 1080 port will be used. If specified, the proxy will overwrite proxy related environment variables like `http_proxy`, `HTTPS_PROXY`. If not specified, the proxy will not overwrite proxy-related environment variables. The entered value is passed on "as is", no sanity checking takes place. Note that only simple authentication is supported with HTTP proxy. Supported macros: `{HOST.IP}`, `{HOST.CONN}`, `{HOST.DNS}`, `{HOST.HOST}`, `{HOST.NAME}`, `{HOST.PORT}`, `{ITEM.ID}`, `{ITEM.KEY}`, `{ITEM.KEY.ORIG}`, user macros, low-level discovery macros. This sets the `CURLOPT_PROXY` cURL option.
HTTP authentication	Select the authentication option: **None** - no authentication used; **Basic** - basic authentication is used; **NTLM** - NTLM (Windows NT LAN Manager) authentication is used; **Kerberos** - Kerberos authentication is used (see also: Configuring Kerberos with Zabbix); **Digest** - Digest authentication is used. This sets the `CURLOPT_HTTPAUTH` cURL option.
User name	Enter the user name (up to 255 characters). This field is available if *HTTP authentication* is set to *Basic*, *NTLM*, *Kerberos*, or *Digest*. User macros and low-level discovery macros are supported.
Password	Enter the user password (up to 255 characters). This field is available if *HTTP authentication* is set to *Basic*, *NTLM*, *Kerberos*, or *Digest*. User macros and low-level discovery macros are supported.
SSL verify peer	Mark the checkbox to verify the SSL certificate of the web server. The server certificate is verified against the trusted certificate authority (CA) certificates from the system-wide CA location. You can override the location of CA files using Zabbix server or proxy configuration parameter SSLCALocation. This sets the `CURLOPT_SSL_VERIFYPEER` cURL option.
SSL verify host	Mark the checkbox to verify that the *Common Name* field or the *Subject Alternate Name* field of the web server certificate matches. This sets the `CURLOPT_SSL_VERIFYHOST` cURL option.
SSL certificate file	Name of the SSL certificate file used for client authentication. The certificate file must be in PEM[1] format. If the certificate file contains also the private key, leave the *SSL key file* field empty. If the key is encrypted, specify the password in *SSL key password* field. The directory containing this file is specified by Zabbix server or proxy configuration parameter `SSLCertLocation`. Supported macros: `{HOST.IP}`, `{HOST.CONN}`, `{HOST.DNS}`, `{HOST.HOST}`, `{HOST.NAME}`, `{HOST.PORT}`, `{ITEM.ID}`, `{ITEM.KEY}`, `{ITEM.KEY.ORIG}`, user macros, low-level discovery macros. This sets the `CURLOPT_SSLCERT` cURL option.

Parameter	Description
SSL key file	Name of the SSL private key file used for client authentication. The private key file must be in PEM[1] format. The directory containing this file is specified by Zabbix server or proxy configuration parameter `SSLKeyLocation`. Supported macros: `{HOST.IP}`, `{HOST.CONN}`, `{HOST.DNS}`, `{HOST.HOST}`, `{HOST.NAME}`, `{HOST.PORT}`, `{ITEM.ID}`, `{ITEM.KEY}`, `{ITEM.KEY.ORIG}`, user macros, low-level discovery macros. This sets the `CURLOPT_SSLKEY` cURL option.
SSL key password	SSL private key file password. Supported macros: user macros, low-level discovery macros. This sets the `CURLOPT_KEYPASSWD` cURL option.
Timeout	Zabbix will not spend more than the set amount of time on processing the URL (1-600 seconds). Actually this parameter defines the maximum time for making a connection to the URL and maximum time for performing an HTTP request. Therefore, Zabbix will not spend more than 2 x *Timeout* seconds on one check. This sets the `CURLOPT_TIMEOUT` cURL option. For more information on the *Timeout* parameter, see general item attributes.
Enable trapping	With this checkbox marked, the item will also function as a trapper item and will accept data sent to Zabbix server or proxy using the Zabbix sender utility or Zabbix sender protocol, or will accept data sent to Zabbix server using the `history.push` API method. For more information about sending data, see: Trapper items.
Allowed hosts	Visible only if *Enable trapping* checkbox is marked. List of comma-delimited IP addresses (optionally in CIDR notation) or DNS names. The default value is the user macro `{$TRAPPER.ALLOWED_HOSTS}`. By default, `{$TRAPPER.ALLOWED_HOSTS}` is `127.0.0.1,::1`. You can either enter addresses directly in the field or change the macro value. If left empty, incoming connections will be denied from all hosts. If specified, incoming connections will be accepted only from the hosts listed here. If IPv6 support is enabled then `127.0.0.1`, `::127.0.0.1`, `::ffff:127.0.0.1` are treated equally and `::/0` will allow any IPv4 or IPv6 address. `0.0.0.0/0` can be used to allow any IPv4 address. Note that "IPv4-compatible IPv6 addresses" (`0000::/96` prefix) are supported but deprecated by RFC4291. Example: `127.0.0.1`, `192.168.1.0/24`, `192.168.3.1-255`, `192.168.1-10.1-255`, `::1,2001:db8::/32`, `mysqlserver1`, `zabbix.example.com`, `{HOST.HOST}` Spaces and user macros are allowed in this field. Host macros: `{HOST.HOST}`, `{HOST.NAME}`, `{HOST.IP}`, `{HOST.DNS}`, `{HOST.CONN}` are allowed in this field.

Note:
If the *HTTP proxy* field is left empty, another way for using an HTTP proxy is to set proxy-related environment variables. For HTTP - set the `http_proxy` environment variable for the Zabbix server user. For example:
`http_proxy=http://proxy_ip:proxy_port`.
For HTTPS - set the `HTTPS_PROXY` environment variable. For example:
`HTTPS_PROXY=http://proxy_ip:proxy_port`.
More details are available by running a shell command: *# man curl*.

Attention:
[1] Zabbix supports certificate and private key files in PEM format only. In case you have your certificate and private key data in PKCS #12 format file (usually with extension `*.p12` or `*.pfx`) you may generate the PEM file from it using the following commands:

```
openssl pkcs12 -in ssl-cert.p12 -clcerts -nokeys -out ssl-cert.pem
openssl pkcs12 -in ssl-cert.p12 -nocerts -nodes  -out ssl-cert.key
```

Examples

Example 1

Send simple GET requests to retrieve data from services such as Elasticsearch:

- Create a GET item with URL: `localhost:9200/?pretty`
- Notice the response:

```
{
  "name" : "YQ2VAY-",
  "cluster_name" : "elasticsearch",
  "cluster_uuid" : "kH4CYqh5QfqgeTsjh2F9zg",
  "version" : {
    "number" : "6.1.3",
    "build_hash" : "af51318",
    "build_date" : "2018-01-26T18:22:55.523Z",
    "build_snapshot" : false,
    "lucene_version" : "7.1.0",
    "minimum_wire_compatibility_version" : "5.6.0",
    "minimum_index_compatibility_version" : "5.0.0"
  },
  "tagline" : "You know, for search"
}
```


- Now extract the version number using a JSONPath preprocessing step: `$.version.number`

Example 2

Send simple POST requests to retrieve data from services such as Elasticsearch:

- Create a POST item with URL: `http://localhost:9200/str/values/_search?scroll=10s`
- Configure the following POST body to obtain the processor load (1 min average per core)

```
{
    "query": {
        "bool": {
            "must": [{
                "match": {
                    "itemid": 28275
                }
            }],
            "filter": [{
                "range": {
                    "clock": {
                        "gt": 1517565836,
                        "lte": 1517566137
                    }
                }
            }]
        }
    }
}
```

- Received:


```
{
    "_scroll_id": "DnF1ZXJ5VGhlbkZldGNoBQAAAAAAAAAkFllRMlZBWS1UU1pxTmdEeGVwQjRBTFEAAAAAAAAAJRZZUTJWQVk
    "took": 18,
    "timed_out": false,
    "_shards": {
        "total": 5,
        "successful": 5,
        "skipped": 0,
        "failed": 0
    },
    "hits": {
        "total": 1,
        "max_score": 1.0,
        "hits": [{
            "_index": "dbl",
            "_type": "values",
            "_id": "dqX9VWEBV6sEKSMyk6sw",
            "_score": 1.0,
            "_source": {
```

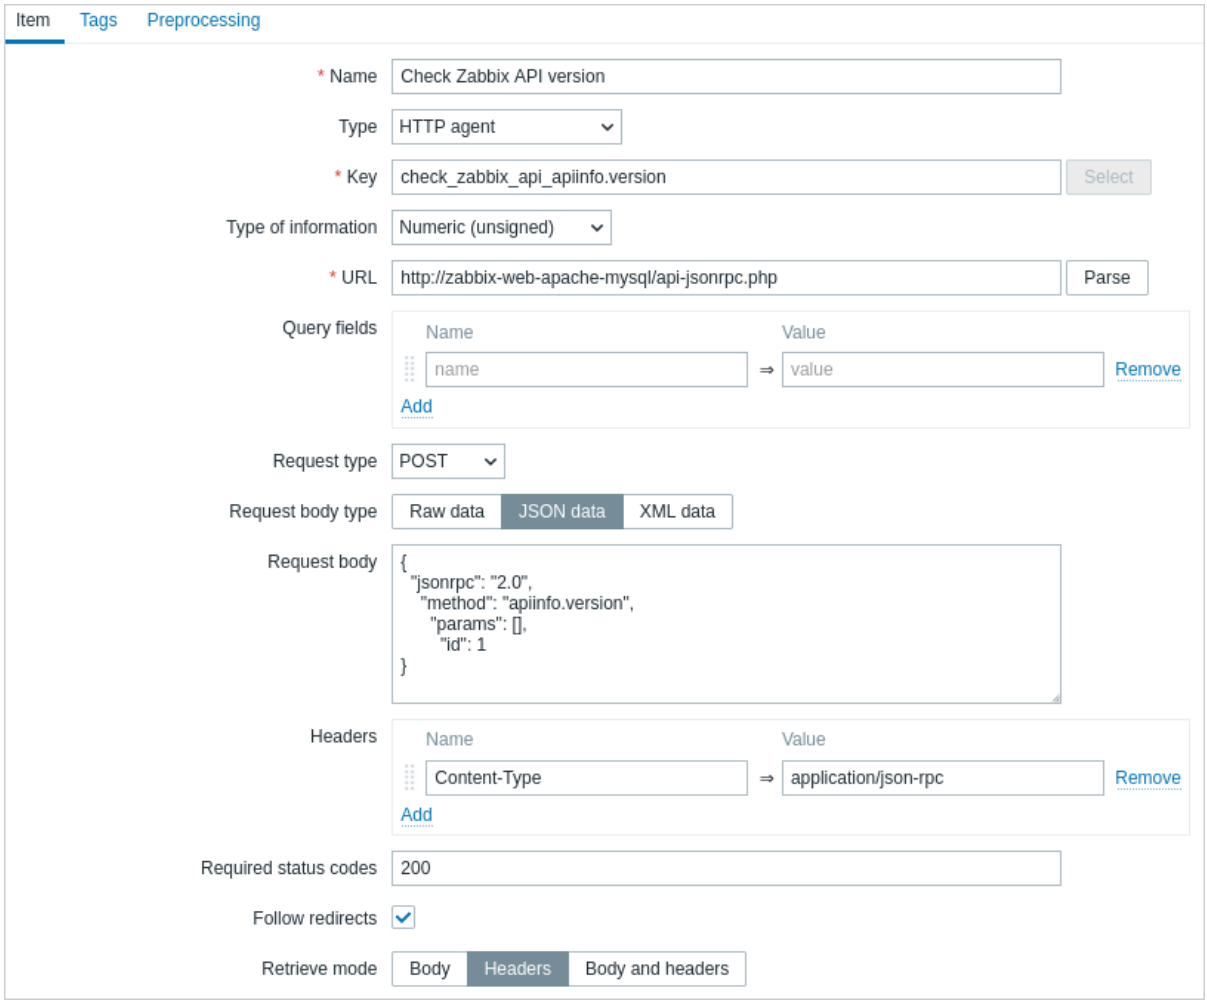
```
                "itemid": 28275,
                "value": "0.138750",
                "clock": 1517566136,
                "ns": 25388713,
                "ttl": 604800
            }
        }]
    }
}
```

- Now use a JSONPath preprocessing step to get the item value: `$.hits.hits[0]._source.value`

Example 3

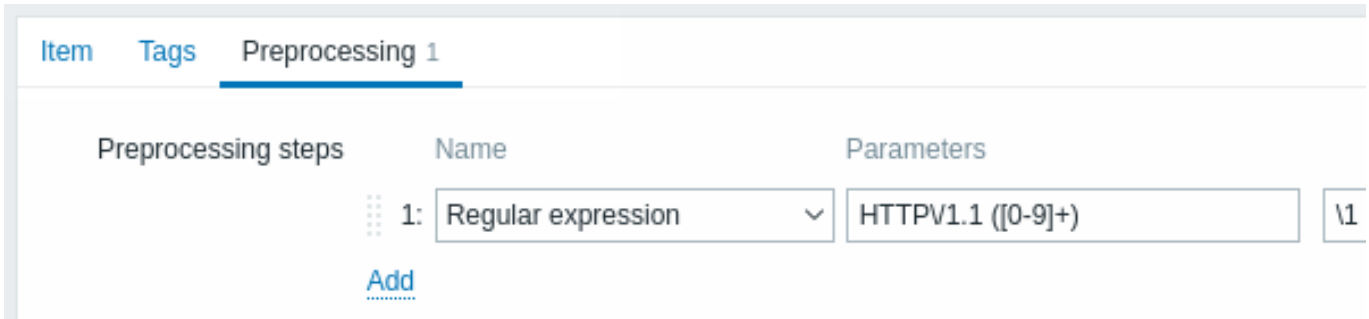
Checking if Zabbix API is alive, using apiinfo.version.

- Item configuration:

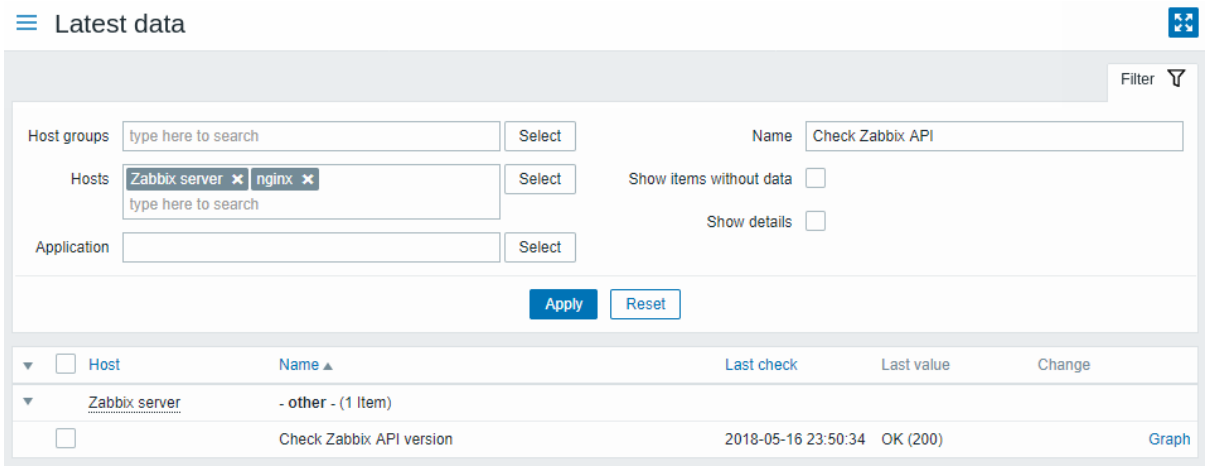


Note the use of the POST method with JSON data, setting request headers and asking to return headers only:

- Item value preprocessing with regular expression to get HTTP code:



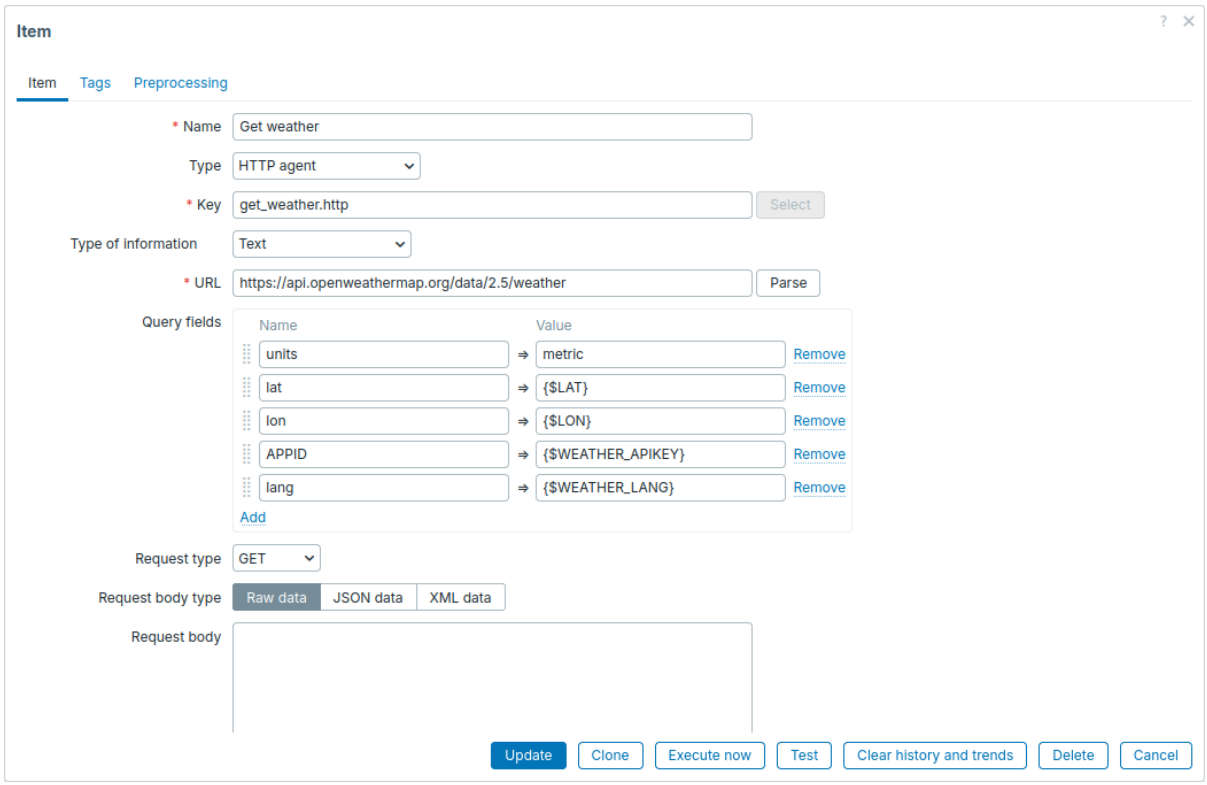
- Checking the result in *Latest data*:



Example 4

Retrieving weather information by connecting to the Openweathermap public service.

- Configure a master item for bulk data collection in a single JSON:



Note the usage of macros in query fields. Refer to the Openweathermap API for how to fill them.

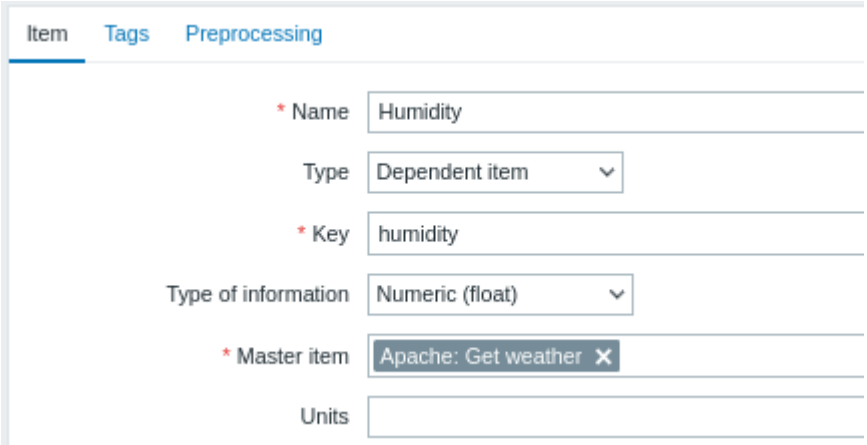
Sample JSON returned in response to HTTP agent:

```
{
    "body": {
        "coord": {
            "lon": 40.01,
            "lat": 56.11
        },
        "weather": [{
            "id": 801,
            "main": "Clouds",
            "description": "few clouds",
            "icon": "02n"
        }],
        "base": "stations",
        "main": {
            "temp": 15.14,
```

```
            "pressure": 1012.6,
            "humidity": 66,
            "temp_min": 15.14,
            "temp_max": 15.14,
            "sea_level": 1030.91,
            "grnd_level": 1012.6
        },
        "wind": {
            "speed": 1.86,
            "deg": 246.001
        },
        "clouds": {
            "all": 20
        },
        "dt": 1526509427,
        "sys": {
            "message": 0.0035,
            "country": "RU",
            "sunrise": 1526432608,
            "sunset": 1526491828
        },
        "id": 487837,
        "name": "Stavrovo",
        "cod": 200
    }
}
```

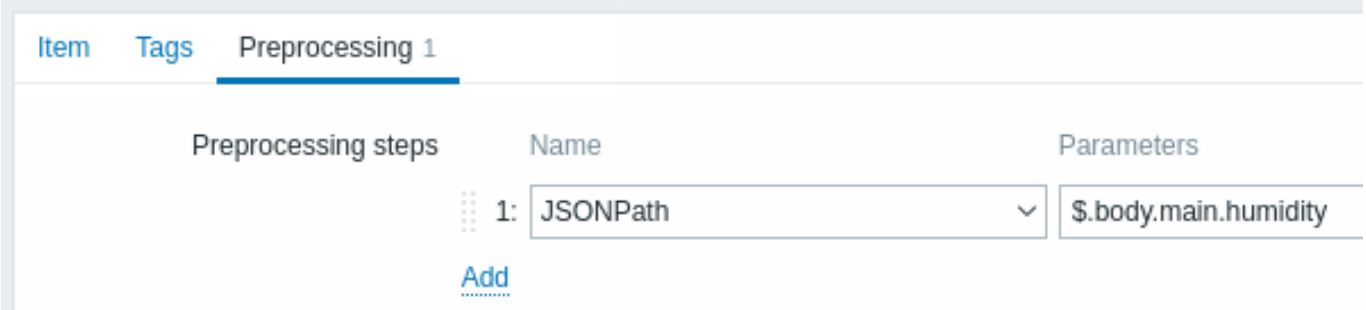

The next task is to configure dependent items that extract data from the JSON.

- Configure a sample dependent item for humidity:

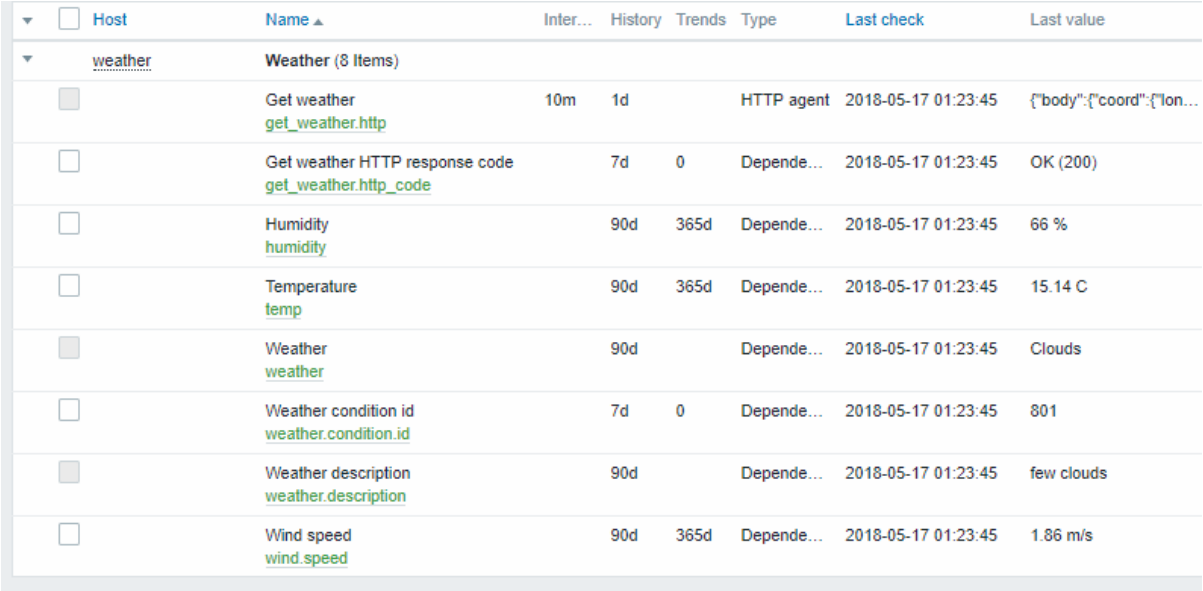


Other weather metrics such as 'Temperature' are added in the same manner.

- Sample dependent item value preprocessing with JSONPath:



- Check the result of weather data in *Latest data*:

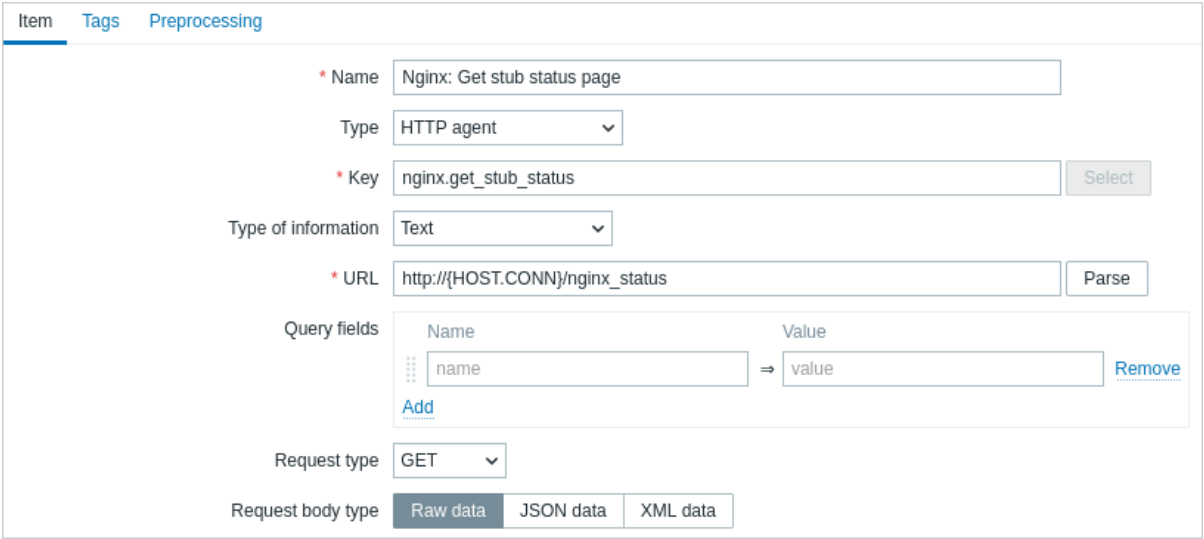


Host	Name ▲	Inter...	History	Trends	Type	Last check	Last value
weather	**Weather** (8 Items)						
	Get weather get_weather.http	10m	1d		HTTP agent	2018-05-17 01:23:45	{"body":{"coord":{"lon...
	Get weather HTTP response code get_weather.http_code		7d	0	Depende...	2018-05-17 01:23:45	OK (200)
	Humidity humidity		90d	365d	Depende...	2018-05-17 01:23:45	66 %
	Temperature temp		90d	365d	Depende...	2018-05-17 01:23:45	15.14 C
	Weather weather		90d		Depende...	2018-05-17 01:23:45	Clouds
	Weather condition id weather.condition.id		7d	0	Depende...	2018-05-17 01:23:45	801
	Weather description weather.description		90d		Depende...	2018-05-17 01:23:45	few clouds
	Wind speed wind.speed		90d	365d	Depende...	2018-05-17 01:23:45	1.86 m/s

Example 5

Connecting to NGINX status page and getting its metrics in bulk.

- Configure NGINX following the official guide.
- Configure a master item for bulk data collection:

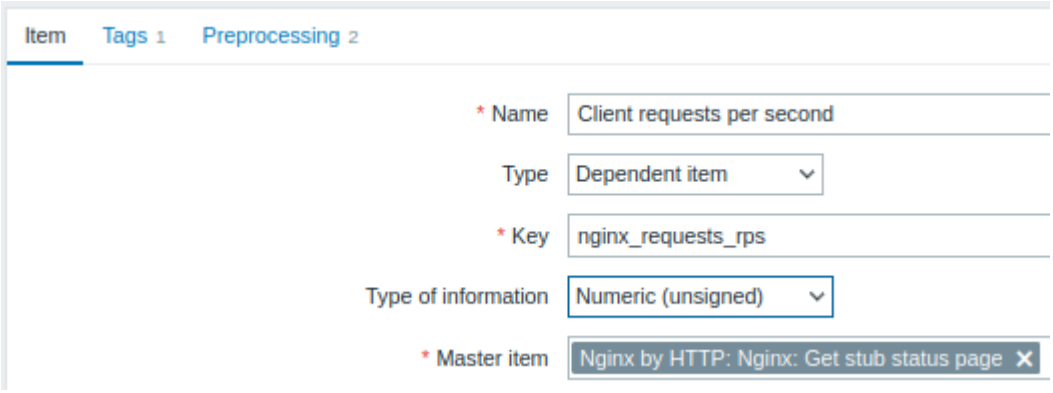


Sample NGINX stub status output:

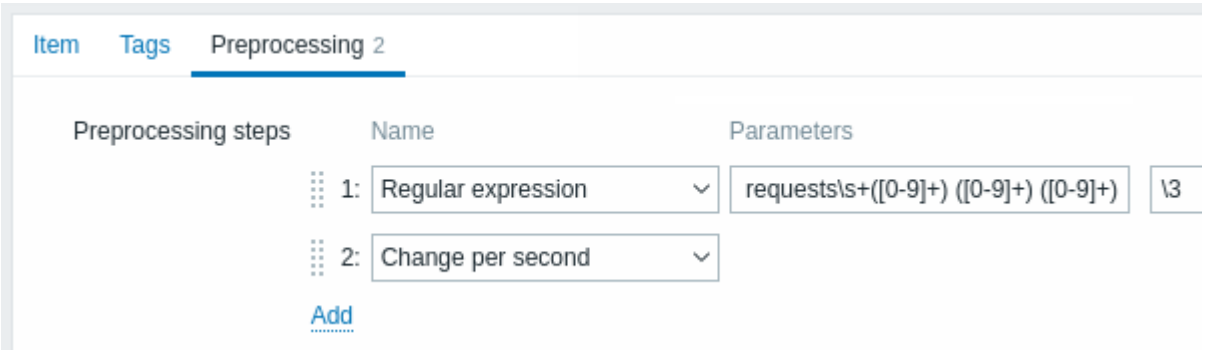
```
Active connections: 1 Active connections:
server accepts handled requests
 52 52 52
Reading: 0 Writing: 1 Waiting: 0
```

The next task is to configure dependent items that extract data.

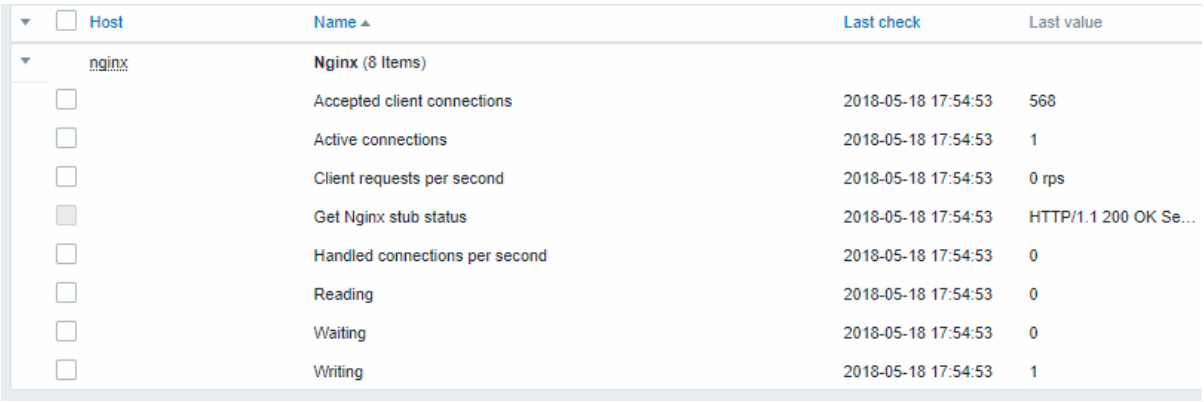
- Configure a sample dependent item for requests per second:



- Sample dependent item value preprocessing with regular expression `server accepts handled requests\s+([0-9]+) ([0-9]+) ([0-9]+)`:



- Check the complete result from stub module in *Latest data*:



Prometheus check

Overview

Zabbix can query metrics exposed in the Prometheus line format.

Two steps are required to start gathering Prometheus data:

- an HTTP master item pointing to the appropriate data endpoint, e.g., `https://<prometheus host>/metrics`
- dependent items using a Prometheus preprocessing option to query required data from the metrics gathered by the master item

There are two Prometheus data preprocessing options:

- *Prometheus pattern* - used in normal items to query Prometheus data
- *Prometheus to JSON* - used in normal items and for low-level discovery. In this case queried Prometheus data is returned in a JSON format.

Bulk processing

Bulk processing is supported for dependent items. To enable caching and indexing, the *Prometheus pattern* preprocessing must be the **first** preprocessing step. When *Prometheus pattern* is first preprocessing step then the parsed Prometheus data is cached and indexed by the first `<label>==<value>` condition in the *Prometheus pattern* preprocessing step. This cache is reused when

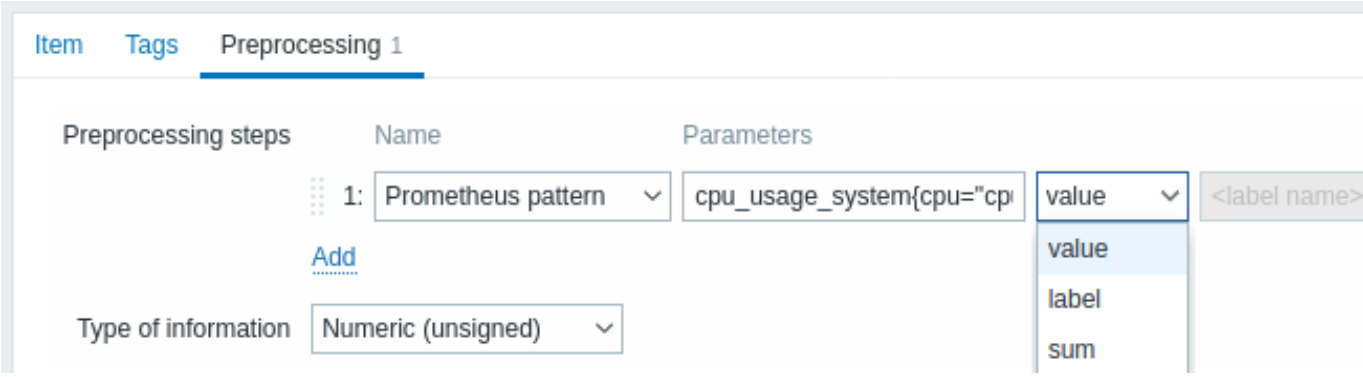
processing other dependent items in this batch. For optimal performance, the first label should be the one with most different values.

If there is other preprocessing to be done before the first step, it should be moved either to the master item or to a new dependent item which would be used as the master item for the dependent items.

Configuration

Providing you have the HTTP master item configured, you need to create a dependent item that uses a Prometheus preprocessing step:

1. Enter general dependent item parameters in the configuration form.
2. Go to the *Preprocessing* tab.
3. Select a Prometheus preprocessing option (*Prometheus pattern* or *Prometheus to JSON*).



The following parameters are specific to the *Prometheus pattern* preprocessing option:

Parameter	Description	Examples
Pattern	To define the required data pattern you may use a query language that is similar to Prometheus query language (see comparison table), e.g.: <metric name> - select by metric name {__name__="<metric name>"} - select by metric name {__name__=~"<regex>"} - select by metric name matching a regular expression {<label name>="<label value>",...} - select by label name {<label name>=~"<regex>",...} - select by label name matching a regular expression {__name__=~".*"}==<value> - select by metric value Or a combination of the above: <metric name>{<label1 name>="<label1 value>",<label2 name>=~"<regex>",...}==<value> Label value can be any sequence of UTF-8 characters, but the backslash, double-quote and line feed characters have to be escaped as \\, \" and \n respectively; other characters shall not be escaped.	*wmi_os_physical_memory_free_bytes* *cpu_usage_system{cpu="cpu-total"}* *cpu_usage_system{cpu=".*"}* *cpu_usage_system{cpu="cpu-total",host="* *wmi_service_state{name="dhcp"}==1* *wmi_os_timezone{timezone=~".*"}==1*

Parameter	Description	Examples
Result processing	Specify whether to return the value, the label or apply the appropriate function (if the pattern matches several lines and the result needs to be aggregated): **value** - return metric value (error if multiple lines matched) **label** - return value of the label specified in the *Label* field (error if multiple metrics are matched) **sum** - return the sum of values **min** - return the minimum value **max** - return the maximum value **avg** - return the average value **count** - return the count of values This field is only available for the *Prometheus pattern* option.	See also examples of using parameters below.
Output	Define label name (optional). In this case the value corresponding to the label name is returned. This field is only available for the *Prometheus pattern* option, if 'Label' is selected in the *Result processing* field.	

Examples of using parameters

1. The most common use case is to return the **value**. To return the value of `/var/db` from:

   `node_disk_usage_bytes{path="/var/cache"} 2.1766144e+09`
`node_disk_usage_bytes{path="/var/db"} 20480`
`node_disk_usage_bytes{path="/var/dpkg"} 8192`
`node_disk_usage_bytes{path="/var/empty"} 4096`

use the following parameters:

- *Pattern* - `node_disk_usage_bytes{path="/var/db"}`
- *Result processing* - select 'value'

2. You may also be interested in the **average** value of all `node_disk_usage_bytes` parameters:

- *Pattern* - `node_disk_usage_bytes`
- *Result processing* - select 'avg'

3. While Prometheus supports only numerical data, a workaround is commonly used to return the relevant textual description as well. This can be accomplished with a filter and specifying the label. So, to return the value of the 'color' label from:

   `elasticsearch_cluster_health_status{cluster="elasticsearch",color="green"} 1`
`elasticsearch_cluster 0`
`elasticsearch_cluster_health_status{cluster="elasticsearch",color="yellow"} 0`

use the following parameters:

- *Pattern* - `elasticsearch_cluster_health_status {cluster="elasticsearch"} == 1`
- *Result processing* - select 'label'
- *Label* - specify 'color'

The filter (based on the numeric value '1') will match the appropriate row, while the label will return the health status description (currently 'green'; but potentially also 'red' or 'yellow').

Prometheus to JSON

Data from Prometheus can be used for low-level discovery. In this case data in JSON format are needed and the *Prometheus to JSON* preprocessing option will return exactly that.

For more details, see Discovery using Prometheus data.

Query language comparison

The following table lists differences and similarities between PromQL and Zabbix Prometheus preprocessing query language.

PromQL instant vector selector	Zabbix Prometheus preprocessing
Differences	

PromQL instant vector selector		Zabbix Prometheus preprocessing
Query target	Prometheus server	Plain text in Prometheus exposition format
Returns	Instant vector	Metric or label value (Prometheus pattern) Array of metrics for single value in JSON (Prometheus to JSON)
Label matching operators	**=**, **!=**, **=~**, **!~**	**=**, **!=**, **=~**, **!~**
Regular expression used in label or metric name matching	RE2	PCRE
Comparison operators	See list	Only **==** (equal) is supported for value filtering
Similarities		
Selecting by metric name that equals string	<metric name> or {__name__="<metric name>"}	<metric name> or {__name__="<metric name>"}
Selecting by metric name that matches regular expression	{__name__=~"<regex>"}	{__name__=~"<regex>"}
Selecting by <label name> value that equals string	{<label name>="<label value>",...}	{<label name>="<label value>",...}
Selecting by <label name> value that matches regular expression	{<label name>=~"<regex>",...}	{<label name>=~"<regex>",...}
Selecting by value that equals string	{__name__=~".*"} == <value>	{__name__=~".*"} == <value>

IPMI agent

Overview

You can monitor the health and availability of Intelligent Platform Management Interface (IPMI) devices in Zabbix. To perform IPMI checks Zabbix server must be initially configured with IPMI support.

IPMI is a standardized interface for remote "lights-out" or "out-of-band" management of computer systems. It allows you to monitor hardware status directly from the so-called "out-of-band" management cards, independently from the operating system or whether the machine is powered on at all.

Zabbix IPMI monitoring works only for devices having IPMI support (HP iLO, DELL DRAC, IBM RSA, Sun SSP, etc).

An IPMI manager process schedules the IPMI checks by IPMI pollers. A host is always polled by only one IPMI poller at a time, reducing the number of open connections to BMC controllers. Thus it's safe to increase the number of IPMI pollers without worrying about BMC controller overloading. The IPMI manager process is automatically started when at least one IPMI poller is started.

See also known issues for IPMI checks.

Configuration

Host configuration

A host must be configured to process IPMI checks. An IPMI interface must be added, with the respective IP and port numbers, and IPMI authentication parameters must be defined.

See the configuration of hosts for more details.

Server configuration

By default, the Zabbix server is not configured to start any IPMI pollers, thus any added IPMI items won't work. To change this, open the Zabbix server configuration file (`zabbix_server.conf`) as root and look for the following line:

```
# StartIPMIPollers=0
```

Uncomment it and set poller count to, say, 3, so that it reads:

```
StartIPMIPollers=3
```

Save the file and restart zabbix_server afterwards.

Item configuration

When configuring an item on a host level:

- Select 'IPMI agent' as the *Type*
- Enter an item key that is unique within the host (say, ipmi.fan.rpm)

- For *Host interface* select the relevant IPMI interface (IP and port). Note that an IPMI interface must exist on the host.
- Specify the *IPMI sensor* (for example, 'FAN MOD 1A RPM' on Dell Poweredge) to retrieve the metric from. By default, the sensor ID should be specified. It is also possible to use prefixes before the value:
  - `id:` - to specify sensor ID;
  - `name:` - to specify sensor full name. This can be useful in situations when sensors can only be distinguished by specifying the full name.
- Select the respective type of information ('Numeric (float)' in this case; for discrete sensors - 'Numeric (unsigned)'), units (most likely 'rpm') and any other required item attributes

Supported checks

IPMI agent supports the built-in item **ipmi.get**, which returns IPMI-sensor related information and can be used for the discovery of IPMI sensors.
 Return value: *JSON object*

Timeout and session termination

IPMI message timeouts and retry counts are defined in OpenIPMI library. Due to the current design of OpenIPMI, it is not possible to make these values configurable in Zabbix, neither on interface nor item level.

IPMI session inactivity timeout for LAN is 60 +/-3 seconds. Currently it is not possible to implement periodic sending of Activate Session command with OpenIPMI. If there are no IPMI item checks from Zabbix to a particular BMC for more than the session timeout configured in BMC then the next IPMI check after the timeout expires will time out due to individual message timeouts, retries or receive error. After that a new session is opened and a full rescan of the BMC is initiated. A new UDP port may be opened to manage the new session. Inactivity is defined by the absence of both outgoing requests and incoming responses. If you want to avoid unnecessary rescans of the BMC it is advised to set the IPMI item polling interval below the IPMI session inactivity timeout configured in BMC.

Notes on IPMI discrete sensors

To find sensors on a host start Zabbix server with **DebugLevel=4** enabled. Wait a few minutes and find sensor discovery records in Zabbix server logfile:

```
$ grep 'Added sensor' zabbix_server.log
8358:20130318:111122.170 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:7 id:'CATERR' reading_type:
8358:20130318:111122.170 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:15 id:'CPU Therm Trip' read
8358:20130318:111122.171 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:17 id:'System Event Log' re
8358:20130318:111122.171 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:17 id:'PhysicalSecurity' re
8358:20130318:111122.171 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:14 id:'IPMI Watchdog' readi
8358:20130318:111122.171 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:16 id:'Power Unit Stat' rea
8358:20130318:111122.171 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:16 id:'P1 Therm Ctrl %' rea
8358:20130318:111122.172 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:16 id:'P1 Therm Margin' rea
8358:20130318:111122.172 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:13 id:'System Fan 2' readin
8358:20130318:111122.172 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:13 id:'System Fan 3' readin
8358:20130318:111122.172 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:14 id:'P1 Mem Margin' readi
8358:20130318:111122.172 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:17 id:'Front Panel Temp' re
8358:20130318:111122.173 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:15 id:'Baseboard Temp' read
8358:20130318:111122.173 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:9 id:'BB +5.0V' reading_typ
8358:20130318:111122.173 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:14 id:'BB +3.3V STBY' readi
8358:20130318:111122.173 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:9 id:'BB +3.3V' reading_typ
8358:20130318:111122.173 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:17 id:'BB +1.5V P1 DDR3' re
8358:20130318:111122.173 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:17 id:'BB +1.1V P1 Vccp' re
8358:20130318:111122.174 Added sensor: host:'192.168.1.12:623' id_type:0 id_sz:14 id:'BB +1.05V PCH' readi
```

To decode IPMI sensor types and states, a copy of IPMI 2.0 specifications is available (please note that no further updates to the IPMI specification are planned).

The first parameter to start with is "reading_type". Use "Table 42-1, Event/Reading Type Code Ranges" from the specifications to decode "reading_type" code. Most of the sensors in our example have "reading_type:0x1" which means "threshold" sensor. "Table 42-3, Sensor Type Codes" shows that "type:0x1" means temperature sensor, "type:0x2" - voltage sensor, "type:0x4" - Fan etc. Threshold sensors sometimes are called "analog" sensors as they measure continuous parameters like temperature, voltage, revolutions per minute.

Another example - a sensor with "reading_type:0x3". "Table 42-1, Event/Reading Type Code Ranges" says that reading type codes 02h-0Ch mean "Generic Discrete" sensor. Discrete sensors have up to 15 possible states (in other words - up to 15 meaningful bits). For example, for sensor 'CATERR' with "type:0x7" the "Table 42-3, Sensor Type Codes" shows that this type means "Processor" and the meaning of individual bits is: 00h (the least significant bit) - IERR, 01h - Thermal Trip etc.

There are few sensors with "reading_type:0x6f" in our example. For these sensors the "Table 42-1, Event/Reading Type Code Ranges" advises to use "Table 42-3, Sensor Type Codes" for decoding meanings of bits. For example, sensor 'Power Unit Stat' has

type "type:0x9" which means "Power Unit". Offset 00h means "PowerOff/Power Down". In other words if the least significant bit is 1, then server is powered off. To test this bit, the bitand function with mask '1' can be used. The trigger expression could be like

```
bitand(last(/www.example.com/Power Unit Stat,#1),1)=1
```

to warn about a server power off.

Notes on discrete sensor names in OpenIPMI-2.0.16, 2.0.17, 2.0.18 and 2.0.19

Names of discrete sensors in OpenIPMI-2.0.16, 2.0.17 and 2.0.18 often have an additional "0" (or some other digit or letter) appended at the end. For example, while ipmitool and OpenIPMI-2.0.19 display sensor names as "PhysicalSecurity" or "CATERR", in OpenIPMI-2.0.16, 2.0.17 and 2.0.18 the names are "PhysicalSecurity0" or "CATERR0", respectively.

When configuring an IPMI item with Zabbix server using OpenIPMI-2.0.16, 2.0.17 and 2.0.18, use these names ending with "0" in the *IPMI sensor* field of IPMI agent items. When your Zabbix server is upgraded to a new Linux distribution, which uses OpenIPMI-2.0.19 (or later), items with these IPMI discrete sensors will become "NOT SUPPORTED". You have to change their *IPMI sensor* names (remove the '0' in the end) and wait for some time before they turn "Enabled" again.

Notes on threshold and discrete sensor simultaneous availability

Some IPMI agents provide both a threshold sensor and a discrete sensor under the same name. The preference is always given to the threshold sensor.

Notes on connection termination

If IPMI checks are not performed (by any reason: all host IPMI items disabled/notsupported, host disabled/deleted, host in maintenance etc.) the IPMI connection will be terminated from Zabbix server or proxy in 3 to 4 hours depending on the time when Zabbix server/proxy was started.

SSH agent

Overview

SSH checks are performed as agent-less monitoring. Zabbix agent is not needed for SSH checks.

To perform SSH checks Zabbix server must be initially configured with SSH2 support (libssh or libssh2). See also: Requirements.

> **Attention:**
> Starting with RHEL 8, only libssh is supported. For other distributions, libssh is recommended over libssh2.

Configuration

Passphrase authentication

SSH checks provide two authentication methods - a user/password pair and key-file based.

If you do not intend to use keys, no additional configuration is required, besides linking libssh or libssh2 to Zabbix, if you're building from source.

Key file authentication

To use key based authentication for SSH items, certain changes to the server configuration are required.

Open the Zabbix server configuration file (zabbix_server.conf) as root and look for the following line:

```
##### SSHKeyLocation=
```

Uncomment it and set the full path to the folder where the public and private keys will be located:

```
SSHKeyLocation=/home/zabbix/.ssh
```

Save the file and restart Zabbix server afterwards.

The path /home/zabbix here is the home directory for the *zabbix* user account, and .ssh is a directory where by default public and private keys will be generated by an ssh-keygen command inside the home directory.

Usually installation packages of Zabbix server from different OS distributions create the *zabbix* user account with a home directory elsewhere, for example, /var/lib/zabbix (as for system accounts).

Before generating the keys, you could reallocate the home directory to /home/zabbix, so that it corresponds with the SSHKeyLocation Zabbix server configuration parameter mentioned above.

Note:
The following steps can be skipped if *zabbix* account has been added manually according to the installation section. In such a case the home directory for the *zabbix* account is most likely already `/home/zabbix`.

To change the home directory of the *zabbix* user account, all working processes which are using it have to be stopped:

```
systemctl stop zabbix-agent
systemctl stop zabbix-server
```

To change the home directory location with an attempt to move it (if it exists) the following command should be executed:

```
usermod -m -d /home/zabbix zabbix
```

It is also possible that a home directory did not exist in the old location, so it should be created at the new location. A safe attempt to do that is:

```
test -d /home/zabbix || mkdir /home/zabbix
```

To be sure that all is secure, additional commands could be executed to set permissions to the home directory:

```
chown zabbix:zabbix /home/zabbix
chmod 700 /home/zabbix
```

Previously stopped processes can now be started again:

```
systemctl start zabbix-agent
systemctl start zabbix-server
```

Now, the steps to generate the public and private keys can be performed with the following commands (for better readability, command prompts are commented out):

```
sudo -u zabbix ssh-keygen -t rsa
##### Generating public/private rsa key pair.
##### Enter file in which to save the key (/home/zabbix/.ssh/id_rsa):
/home/zabbix/.ssh/id_rsa
##### Enter passphrase (empty for no passphrase):
<Leave empty>
##### Enter same passphrase again:
<Leave empty>
##### Your identification has been saved in /home/zabbix/.ssh/id_rsa.
##### Your public key has been saved in /home/zabbix/.ssh/id_rsa.pub.
##### The key fingerprint is:
##### 90:af:e4:c7:e3:f0:2e:5a:8d:ab:48:a2:0c:92:30:b9 zabbix@it0
##### The key's randomart image is:
##### +--[ RSA 2048]----+
##### |                 |
##### |        .        |
##### |       o         |
##### | .      o        |
##### |+      . S       |
##### |.+    o =        |
##### |E .    * =       |
##### |=o . ..* .       |
##### |... oo.o+        |
##### +-----------------+
```

Note:
The public and private keys (`id_rsa.pub` and `id_rsa`) have been generated by default in the `/home/zabbix/.ssh` directory, which corresponds to the Zabbix server `SSHKeyLocation` configuration parameter.

Attention:
Key types other than "rsa" may be supported by the ssh-keygen tool and SSH servers but they may not be supported by libssh2 used by Zabbix.

Shell configuration form

This step should be performed only once for every host that will be monitored by SSH checks.

By using the following commands, the **public** key file can be installed on a remote host *10.10.10.10*, so that the SSH checks can be performed with a *root* account (for better readability, command prompts are commented out):

```
sudo -u zabbix ssh-copy-id root@10.10.10.10
##### The authenticity of host '10.10.10.10 (10.10.10.10)' can't be established.
##### RSA key fingerprint is 38:ba:f2:a4:b5:d9:8f:52:00:09:f7:1f:75:cc:0b:46.
##### Are you sure you want to continue connecting (yes/no)?
yes
##### Warning: Permanently added '10.10.10.10' (RSA) to the list of known hosts.
##### root@10.10.10.10's password:
<Enter root password>
##### Now try logging into the machine, with "ssh 'root@10.10.10.10'",
##### and check to make sure that only the key(s) you wanted were added.
```

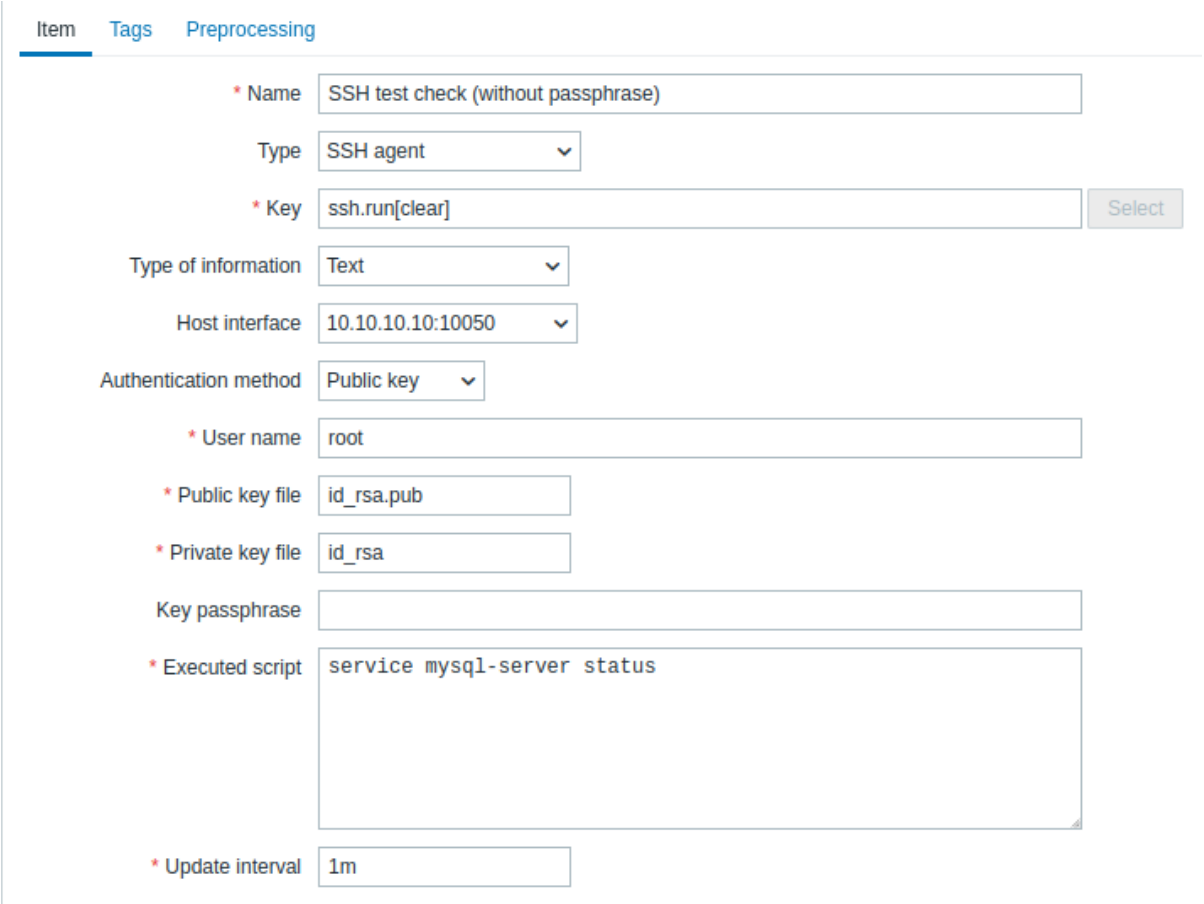
Now it is possible to check the SSH login using the default private key (`/home/zabbix/.ssh/id_rsa`) for the *zabbix* user account:

```
sudo -u zabbix ssh root@10.10.10.10
```

If the login is successful, then the configuration part in the shell is finished and the remote SSH session can be closed.

Item configuration

Actual command(s) to be executed must be placed in the *Executed script* field in the item configuration. Multiple commands can be executed one after another by placing them on a new line. In this case returned values will also be formatted as multilined.



All mandatory input fields are marked with a red asterisk.

The fields that require specific information for SSH items are:

Parameter	Description	Comments
Type	Select **SSH agent** here.	

Parameter	Description	Comments

Key	Unique (per host) item key in the format `ssh.run[unique short description,<ip>,<port>,<encoding>,<ssh options>,<subsystem>]`	**unique short description** is required and should be unique for each SSH item per host.
		
		Default port is 22, not the port specified in the interface to which this item is assigned.
		
		ssh options allow additional SSH options to be passed in the format `key1=value1;key2=value2,value3`. Multiple values for one key can be passed separated by comma (in this case, the parameter must be quoted); multiple option keys can be passed separated by semicolon.
		
		The following option keys are supported: `KexAlgorithms`, `HostkeyAlgorithms`, `Ciphers`, `MACs`, `PubkeyAcceptedKeyTypes`. Option key and value support depends on the SSH library (for example, `PubkeyAcceptedKeyTypes` is supported only with libssh); if an option is not supported, an error will be returned, and the item will become unsupported.
		
		Note that "+" sign for appending cipher settings and "!" for disabling specific cipher settings (as in GnuTLS and OpenSSL) are not supported.
		
		Examples:
		=>
		`ssh.run[KexAlgorithms,127.0.0.1,,,Ciphers=aes128`
		=>
		`ssh.run[KexAlgorithms,,,,"KexAlgorithms=diffie-h`
		=>
		`ssh.run[PubkeyAcceptedKeyTypes,127.0.0.1,,,Pubke`
		
		subsystem allows an SSH subsystem to be passed, limiting the SSH connection to specific operations permitted by the subsystem (e.g., file transfers using SFTP or network device management using NETCONF). Note that using a subsystem may also require the use of specific script syntax in the *Executed script* parameter.
		
		Examples:
		=>
		`ssh.run[SFTPBackup,192.0.2.18,,,sftp]`
		=>
		`ssh.run[Cisco1234,192.0.2.18,,,netconf]`

Authentication method	One of the "Password" or "Public key".	
User name	User name (up to 255 characters) to authenticate on remote host. Required.	
Public key file	File name of public key if *Authentication method* is "Public key". Required.	Example: *id_rsa.pub* - default public key file name generated by a command ssh-keygen.
Private key file	File name of private key if *Authentication method* is "Public key". Required.	Example: *id_rsa* - default private key file name.
Password or *Key passphrase*	Password (up to 255 characters) to authenticate or Passphrase **if** it was used for the private key.	Leave the *Key passphrase* field empty if passphrase was not used. See also known issues regarding passphrase usage.

Parameter	Description	Comments

Executed script	Executed shell command(s) using SSH remote session.	The return value of the executed shell command(s) is limited to 16MB (including trailing whitespace that is truncated); database limits also apply.
		
		Note that the libssh2 library may truncate executable scripts to ~32kB.
		
		Examples:
		`date +%s`
		`systemctl status mysql-server`
		`ps auxww | grep httpd | wc -l`
		
		Example (for NETCONF subsystem):
		`<rpc>`
		`<get-software-information/>`
		`</rpc>`
		`]]>]]>`
		`<rpc>`
		`<close-session/>`
		`</rpc>`
		`]]>]]>`

Telnet agent

Overview

Telnet checks are performed as agent-less monitoring. Zabbix agent is not needed for Telnet checks.

Configurable fields

Actual command(s) to be executed must be placed in the **Executed script** field in the item configuration.
 Multiple commands can be executed one after another by placing them on a new line. In this case returned value also will be formatted as multilined.

Supported characters that the shell prompt can end with:

- $
- #
- >
- %

Note:
A telnet prompt line which ended with one of these characters will be removed from the returned value, but only for the first command in the commands list, i.e. only at a start of the telnet session.

Key	Description
telnet.run[<unique short description>,<ip>,<port>,<encoding>]	Run a command on a remote device using telnet connection

Attention:
If a telnet check returns a value with non-ASCII characters and in non-UTF8 encoding then the *<encoding>* parameter of the key should be properly specified. See encoding of returned values page for more details.

JMX agent

Overview

JMX monitoring can be used to monitor JMX counters of a Java application.

JMX monitoring has native support in Zabbix in the form of a Zabbix daemon called "Zabbix Java gateway".

To retrieve the value of a particular JMX counter on a host, Zabbix server queries the Zabbix **Java gateway**, which in turn uses the JMX management API to query the application of interest remotely.

For more details and setup see the Zabbix Java gateway section.

> **Warning:**
> Communication between Java gateway and the monitored JMX application should not be firewalled.

Enabling remote JMX monitoring for Java application

A Java application does not need any additional software installed, but it needs to be started with the command-line options specified below to have support for remote JMX monitoring.

As a bare minimum, if you just wish to get started by monitoring a simple Java application on a local host with no security enforced, start it with these options:

```
java \
-Dcom.sun.management.jmxremote \
-Dcom.sun.management.jmxremote.port=12345 \
-Dcom.sun.management.jmxremote.authenticate=false \
-Dcom.sun.management.jmxremote.ssl=false \
-Dcom.sun.management.jmxremote.registry.ssl=false \
-jar /path/to/your/application.jar
```

This makes Java listen for incoming JMX connections on port 12345, from local host only, and tells it not to require authentication or SSL.

If you want to allow connections on another interface, set the -Djava.rmi.server.hostname parameter to the IP of that interface.

If you wish to be more stringent about security, there are many other Java options available to you. For instance, the next example starts the application with a more versatile set of options and opens it to a wider network, not just local host.

```
java \
-Djava.rmi.server.hostname=192.168.3.14 \
-Dcom.sun.management.jmxremote \
-Dcom.sun.management.jmxremote.port=12345 \
-Dcom.sun.management.jmxremote.authenticate=true \
-Dcom.sun.management.jmxremote.password.file=/etc/java-6-openjdk/management/jmxremote.password \
-Dcom.sun.management.jmxremote.access.file=/etc/java-6-openjdk/management/jmxremote.access \
-Dcom.sun.management.jmxremote.ssl=true \
-Dcom.sun.management.jmxremote.registry.ssl=true \
-Djavax.net.ssl.keyStore=$YOUR_KEY_STORE \
-Djavax.net.ssl.keyStorePassword=$YOUR_KEY_STORE_PASSWORD \
-Djavax.net.ssl.trustStore=$YOUR_TRUST_STORE \
-Djavax.net.ssl.trustStorePassword=$YOUR_TRUST_STORE_PASSWORD \
-Dcom.sun.management.jmxremote.ssl.need.client.auth=true \
-jar /path/to/your/application.jar
```

Most (if not all) of these settings can be specified in `$JRE/lib/management/management.properties` (or wherever that file is on your system).

Note that if you wish to use SSL, you have to modify startup.sh script by adding `-Djavax.net.ssl.*` options to Java gateway, so that it knows where to find key and trust stores.

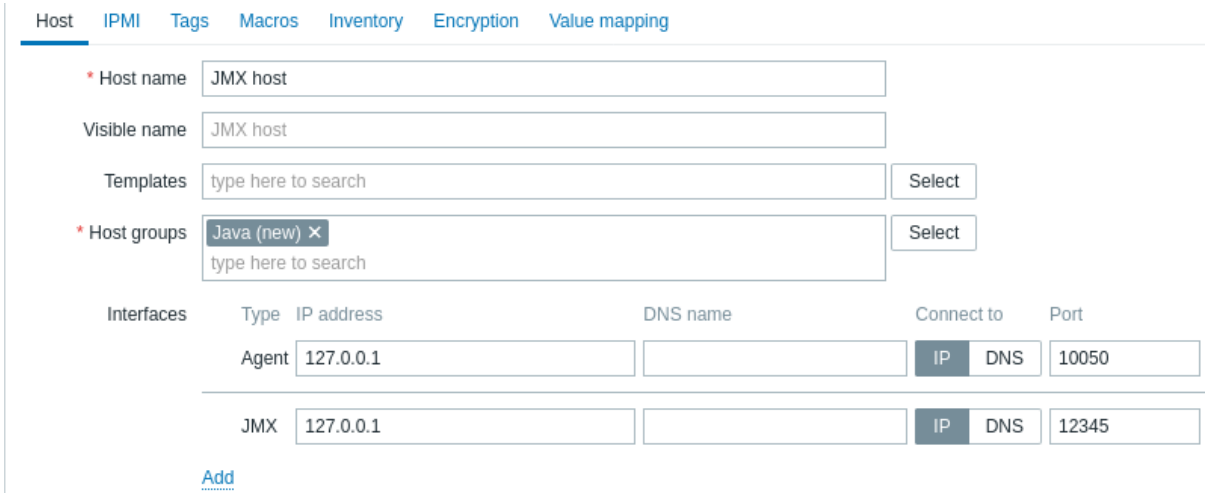
See Monitoring and Management Using JMX for a detailed description.

Configuring JMX interfaces and items in Zabbix frontend

With Java gateway running, server knowing where to find it and a Java application started with support for remote JMX monitoring, it is time to configure the interfaces and items in Zabbix GUI.

Configuring JMX interface

You begin by creating a JMX-type interface on the host of interest.

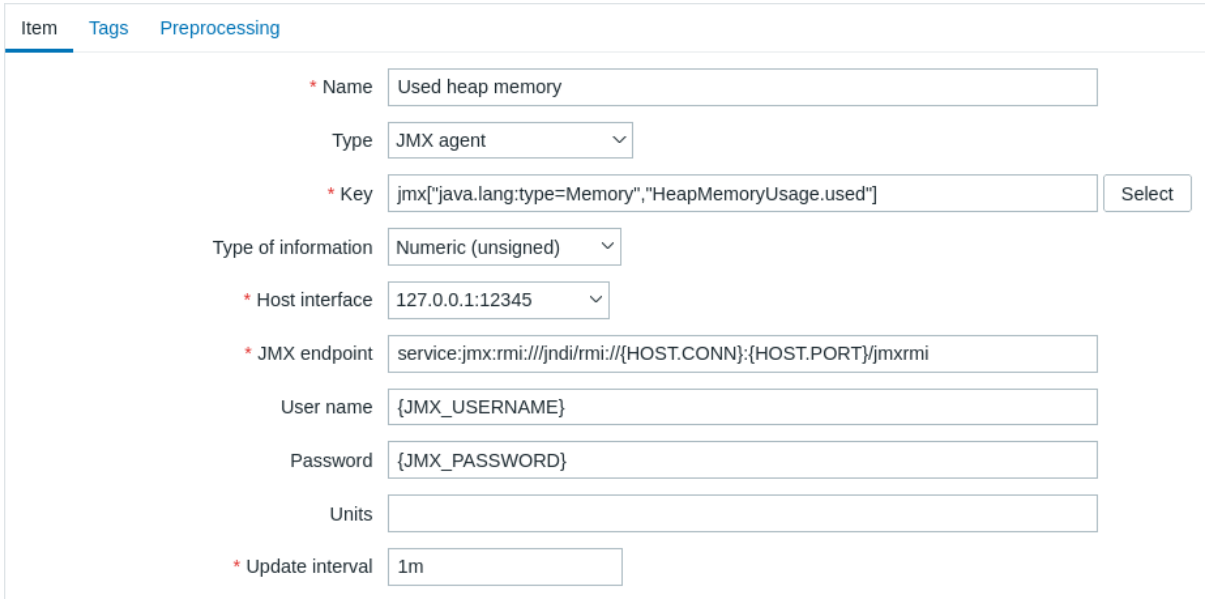


All mandatory input fields are marked with a red asterisk.

Adding JMX agent item

For each JMX counter you are interested in you add **JMX agent** item attached to that interface.

The key in the screenshot below says `jmx["java.lang:type=Memory","HeapMemoryUsage.used"]`.



All mandatory input fields are marked with a red asterisk.

The fields that require specific information for JMX items are:

Type	Set **JMX agent** here.
Key	The `jmx[]` item key contains three parameters: **object name** - the object name of an MBean **attribute name** - an MBean attribute name with optional composite data field names separated by dots **unique short description** - a unique description that allows multiple JMX items with the same object name and attribute name on the host (optional) See below for more detail on JMX item keys. You may discover MBeans and MBean attributes using a `jmx.discovery[]` low-level discovery item.
JMX endpoint	You may specify a custom JMX endpoint. Make sure that JMX endpoint connection parameters match the JMX interface. This can be achieved by using `{HOST.*}` macros as done in the default JMX endpoint. `{HOST.*}` macros and user macros are supported.
User name	Specify the user name (up to 255 characters), if you have configured authentication on your Java application. User macros are supported.

Password	Specify the password (up to 255 characters), if you have configured authentication on your Java application. User macros are supported.

If you wish to monitor a Boolean counter that is either "true" or "false", then you specify type of information as "Numeric (unsigned)" and select "Boolean to decimal" preprocessing step in the Preprocessing tab. Server will store Boolean values as 1 or 0, respectively.

JMX item keys in more detail

Simple attributes

An MBean object name is nothing but a string which you define in your Java application. An attribute name, on the other hand, can be more complex. In case an attribute returns primitive data type (an integer, a string etc.) there is nothing to worry about, the key will look like this:

```
jmx[com.example:Type=Hello,weight]
```

In this example the object name is "com.example:Type=Hello", the attribute name is "weight", and the returned value type should probably be "Numeric (float)".

Attributes returning composite data

It becomes more complicated when your attribute returns composite data. For example: your attribute name is "apple" and it returns a hash representing its parameters, like "weight", "color" etc. Your key may look like this:

```
jmx[com.example:Type=Hello,apple.weight]
```

This is how an attribute name and a hash key are separated, by using a dot symbol. Same way, if an attribute returns nested composite data the parts are separated by a dot:

```
jmx[com.example:Type=Hello,fruits.apple.weight]
```

Attributes returning tabular data

Tabular data attributes consist of one or multiple composite attributes. If such an attribute is specified in the attribute name parameter then this item value will return the complete structure of the attribute in JSON format. The individual element values inside the tabular data attribute can be retrieved using preprocessing.

Tabular data attribute example:

```
 jmx[com.example:type=Hello,foodinfo]
```

Item value:

```
[
  {
    "a": "apple",
    "b": "banana",
    "c": "cherry"
  },
  {
    "a": "potato",
    "b": "lettuce",
    "c": "onion"
  }
]
```

Problem with dots

So far so good. But what if an attribute name or a hash key contains dot symbol? Here is an example:

```
jmx[com.example:Type=Hello,all.fruits.apple.weight]
```

That's a problem. How to tell Zabbix that attribute name is "all.fruits", not just "all"? How to distinguish a dot that is part of the name from the dot that separates an attribute name and hash keys?

This is possible, all you need to do is to escape the dots that are part of the name with a backslash:

```
jmx[com.example:Type=Hello,all\.fruits.apple.weight]
```

Same way, if your hash key contains a dot you escape it:

```
jmx[com.example:Type=Hello,all\.fruits.apple.total\.weight]
```

Other issues

A backslash character in an attribute name should be escaped:

```
jmx[com.example:type=Hello,c:\\documents]
```

For handling any other special characters in JMX item key, please see the item key format section.

This is actually all there is to it. Happy JMX monitoring!

Non-primitive data types

It is possible to work with custom MBeans returning non-primitive data types, which override the **toString()** method.

Using custom endpoint with JBoss EAP 6.4

Custom endpoints allow working with different transport protocols other than the default RMI.

To illustrate this possibility, let's try to configure JBoss EAP 6.4 monitoring as an example. First, let's make some assumptions:

- You have already installed Zabbix Java gateway. If not, then you can do it in accordance with the documentation.
- Zabbix server and Java gateway are installed with the prefix `/usr/local/`
- JBoss is already installed in `/opt/jboss-eap-6.4/` and is running in standalone mode
- We shall assume that all these components work on the same host
- Firewall and SELinux are disabled (or configured accordingly)

Let's make some simple settings in `zabbix_server.conf`:

```
JavaGateway=127.0.0.1
StartJavaPollers=5
```

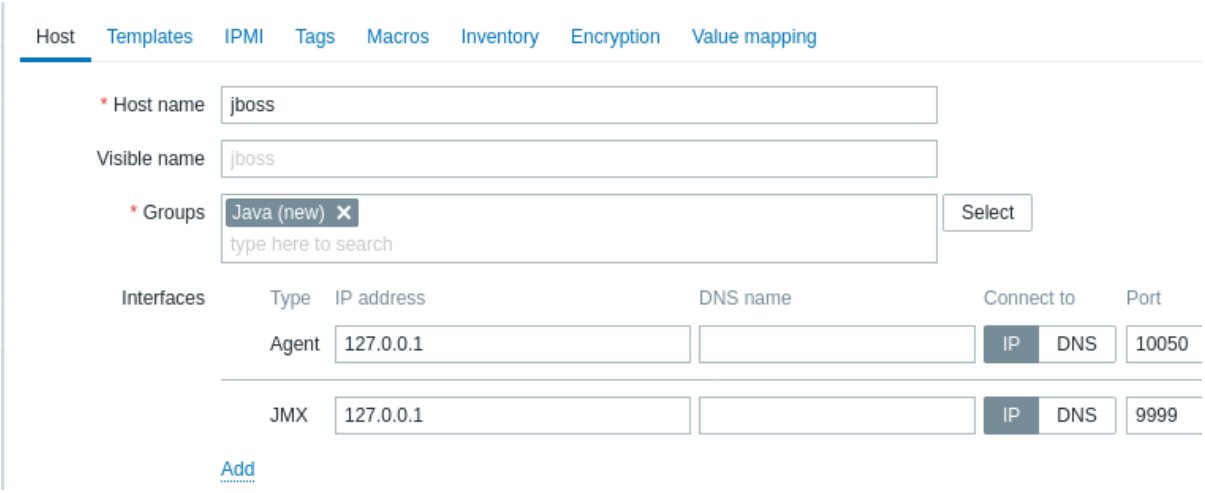
And in the `zabbix_java/settings.sh` configuration file (or `zabbix_java_gateway.conf`):

```
START_POLLERS=5
```

Check that JBoss listens to its standard management port:

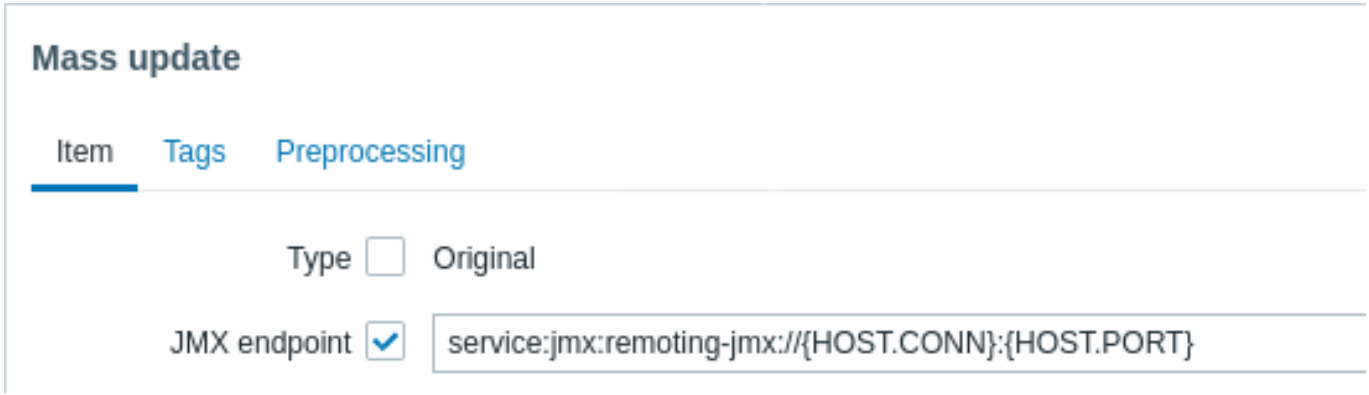
```
$ netstat -natp | grep 9999
tcp        0      0 127.0.0.1:9999          0.0.0.0:*               LISTEN      10148/java
```

Now let's create a host with JMX interface 127.0.0.1:9999 in Zabbix.



As we know that this version of JBoss uses the JBoss Remoting protocol instead of RMI, we may mass update the JMX endpoint parameter for items in our JMX template accordingly:

```
service:jmx:remoting-jmx://{HOST.CONN}:{HOST.PORT}
```



Let's update the configuration cache:

```
/usr/local/sbin/zabbix_server -R config_cache_reload
```

Note that you may encounter an error first.

```
3. mc [root@centos7-dev]:/home/vagrant/zabbix-3.2.6/src/zabbix_java (ssh)
com.zabbix.gateway.ZabbixException: java.net.MalformedURLException: Unsupported protocol: remoting-jmx
        at com.zabbix.gateway.JMXItemChecker.getValues(JMXItemChecker.java:97) ~[zabbix-java-gateway-3.4.2.jar:na]
        at com.zabbix.gateway.SocketProcessor.run(SocketProcessor.java:63) ~[zabbix-java-gateway-3.4.2.jar:na]
        at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1149) [na:1.8.0_144]
        at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:624) [na:1.8.0_144]
        at java.lang.Thread.run(Thread.java:748) [na:1.8.0_144]
Caused by: java.net.MalformedURLException: Unsupported protocol: remoting-jmx
        at javax.management.remote.JMXConnectorFactory.newJMXConnector(JMXConnectorFactory.java:359) ~[na:1.8.0_144]
        at javax.management.remote.JMXConnectorFactory.connect(JMXConnectorFactory.java:269) ~[na:1.8.0_144]
        at com.zabbix.gateway.ZabbixJMXConnectorFactory$1.run(ZabbixJMXConnectorFactory.java:76) ~[zabbix-java-gatewa
-3.4.2.jar:na]
        at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:511) ~[na:1.8.0_144]
        at java.util.concurrent.FutureTask.run(FutureTask.java:266) ~[na:1.8.0_144]
        ... 3 common frames omitted
2017-11-07 13:52:12.644 [pool-1-thread-1] WARN  com.zabbix.gateway.SocketProcessor - error processing request
com.zabbix.gateway.ZabbixException: java.net.MalformedURLException: Unsupported protocol: remoting-jmx
        at com.zabbix.gateway.JMXItemChecker.getValues(JMXItemChecker.java:97) ~[zabbix-java-gateway-3.4.2.jar:na]
        at com.zabbix.gateway.SocketProcessor.run(SocketProcessor.java:63) ~[zabbix-java-gateway-3.4.2.jar:na]
        at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1149) [na:1.8.0_144]
        at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:624) [na:1.8.0_144]
        at java.lang.Thread.run(Thread.java:748) [na:1.8.0_144]
Caused by: java.net.MalformedURLException: Unsupported protocol: remoting-jmx
        at javax.management.remote.JMXConnectorFactory.newJMXConnector(JMXConnectorFactory.java:359) ~[na:1.8.0_144]
        at javax.management.remote.JMXConnectorFactory.connect(JMXConnectorFactory.java:269) ~[na:1.8.0_144]
        at com.zabbix.gateway.ZabbixJMXConnectorFactory$1.run(ZabbixJMXConnectorFactory.java:76) ~[zabbix-java-gatewa
-3.4.2.jar:na]
        at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:511) ~[na:1.8.0_144]
        at java.util.concurrent.FutureTask.run(FutureTask.java:266) ~[na:1.8.0_144]
        ... 3 common frames omitted
2017-11-07 13:52:14.889 [Thread-0] INFO  com.zabbix.gateway.JavaGateway - Zabbix Java Gateway 3.4.2 (revision 72885)
as stopped
2017-11-07 13:52:26.167 [main] INFO  com.zabbix.gateway.JavaGateway - Zabbix Java Gateway 3.4.2 (revision 72885) has
tarted
```

"Unsupported protocol: remoting-jmx" means that Java gateway does not know how to work with the specified protocol. That can be fixed by creating a `~/needed_modules.txt` file with the following content:

```
jboss-as-remoting
jboss-logging
jboss-logmanager
jboss-marshalling
jboss-remoting
jboss-sasl
jcl-over-slf4j
jul-to-slf4j-stub
log4j-jboss-logmanager
remoting-jmx
slf4j-api
xnio-api
xnio-nio
```

and then executing the command:

```
for i in $(cat ~/needed_modules.txt); do find /opt/jboss-eap-6.4 -iname "${i}*.jar" -exec cp '{}' /usr/loc
```

Thus, Java gateway will have all the necessary modules for working with jmx-remoting. What's left is to restart the Java gateway, wait a bit and if you did everything right, see that JMX monitoring data begin to arrive in Zabbix (see also: Latest data).

Calculated item

Overview

This item type allows you to perform calculations on the metrics that already exist. A calculated item does not collect any data by itself.

For example, you may want to:

- Count two values.
- Calculate the average of several item values.
- Calculate the total value for a group of items.

The resulting value is stored in the Zabbix database as for any other item; both history and trend values are stored and graphs can be generated.

The calculation is based on a formula. The syntax of the formula is shared with trigger expressions.

In fact, it's possible to keep triggers simple by having complex data-shaping logic in a calculated item—then reference the calculated item in a trigger using the item key.

Two possibilities exist for specifying data source for calculations:

1. Specify items in the formula.
2. Do data aggregation based on a filter that selects items.

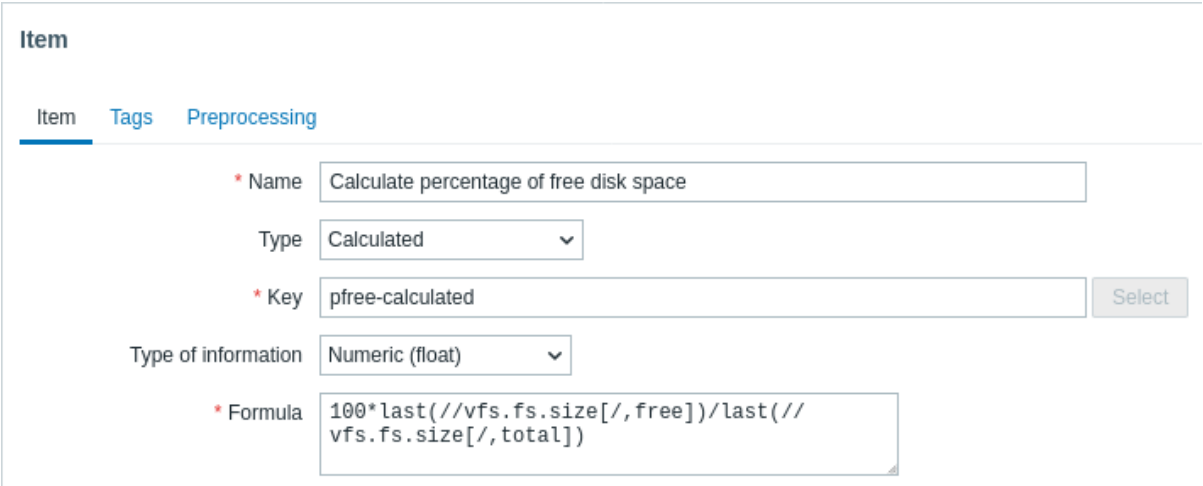
The second option allows you to calculate group averages, totals, etc.

All calculations are done by Zabbix server and handled by the `history poller` processes; proxy or agents are not involved at all. Note that extensive use of calculated items with long time periods may affect the server performance.

Please note that calculated items can access the history of other items, including items that the user creating the calculated item does not have access to. If this behavior is not desired, you can disable calculated items by setting the `StartHistoryPollers` configuration parameter to 0.

Configurable fields

To use calculated items, choose the item type **Calculated**.



The **key** is a unique item identifier (per host). You can create any key name using supported symbols. The key allows you to reference the calculated item in triggers.

The calculation definition must be entered in the **Formula** field.

There is no connection between the formula and the key. The key parameters are not used in the formula in any way.

Unlike trigger expressions, Zabbix processes calculated items according to the item update interval, not upon receiving a new value.

Note that if you select *Numeric (unsigned)* as type of information and the calculation result is a float value, it is trimmed to an integer.

Simple calculation formula

The syntax of a simple formula is:

`function(/host/key,<parameter1>,<parameter2>,...)`

where:

Element	Description
`function`	One of the supported functions: last, min, max, avg, count, etc
`host`	Host of the item that is used for calculation. The current host can be omitted (i.e. as in `function(//key,parameter,...)`).
`key`	Key of the item that is used for calculation. Items returning values with the binary or JSON data type are not supported. If you change the item key of a referenced item, the key must be manually updated in the formula. All items that are referenced by history functions must exist and be collecting data.
`parameter(s)`	Parameters of the function, if required. Time suffixes and memory size suffixes are supported.

A more complex formula may use a combination of functions, operators and brackets.

See examples of calculated item formulas:

- (1) Total traffic on interface
- (2) Percentage of free disk space
- (3) Percentage of incoming traffic
- (4) Average number of processed values

Some variables may be passed to the formula with user macros. Note that user macros:

- will be resolved - if used to reference a function parameter, item filter parameter, or a constant.
- **will not** be resolved - if referencing a function, host name, item key, item key parameter or operator.

Comparison to strings is allowed.

If there is no recent data in the cache and there is no defined querying period in the function, Zabbix will by default go as far back in the past as one week to query the database for historical values.

A calculated item may become unsupported in several cases:

- referenced item
  - is not found
  - is disabled, or belongs to a disabled host
  - is not supported (except with nodata() function and operators with unknown values)
- no data to calculate a function
- division by zero
- incorrect syntax used

Data aggregations

Data aggregations are based on a filter that selects multiple items.

Data aggregations are also a calculated item type, however, you use a slightly different syntax in the formula.

To retrieve aggregates use one of the supported aggregate functions: `avg`, `max`, `min`, `sum`, etc. Then add the **foreach** function as the only parameter and its item filter to select the required items:

```
aggregate_function(function_foreach(/host/key?[group="host group"],timeperiod))
```

A **foreach** function (e.g., *avg_foreach*, *count_foreach*, etc.) returns one aggregate value for each selected item. Items are selected by using the item filter (`/host/key?[group="host group"]`), from item history. For more details, see foreach functions.

If some of the items have no data for the requested period, they are ignored in the calculation. If no items have data, the function will return an error.

Alternatively you may list several items as parameters for aggregation:

```
aggregate_function(function(/host/key,parameter),function(/host2/key2,parameter),...)
```

Note that `function` here must be a history/trend function.

See examples of aggregate calculations:

- (5) Total incoming traffic for host
- (6) Total disk space for group
- (7) Average processor load for group

- (8) Average processor load for group and tag
- (9) Average number of processed values
- (10) Average number of queries for group
- (11) Total number of successful DNS checks
- (12) Total number of unsupported items for group
- (13) Interface stats for group

User macros and low-level discovery macros are supported in:

- item key parameters
- function parameters
- filter conditions (host group name and tag name)
- expression constants

An aggregate calculation may become unsupported if:

- none of the referenced items is found (which may happen if the item key is incorrect, none of the items exists or all included groups are incorrect)
- no data to calculate a function

Examples of calculated items

(1) Total traffic on interface

```
last(//net.if.in[eth0,bytes])+last(//net.if.out[eth0,bytes])
```

Calculate the total bandwidth on eth0.

(2) Percentage of free disk space

```
100*last(//vfs.fs.size[/,free])/last(//vfs.fs.size[/,total])
```

Zabbix will take the latest values for free and total disk spaces (on /) and calculate percentage according to the given formula.

(3) Percentage of incoming traffic

```
100*last(//net.if.in[eth0,bytes])/(last(//net.if.in[eth0,bytes])+last(//net.if.out[eth0,bytes]))
```

Calculate the percentage of incoming traffic from the total traffic.

(4) Average number of processed values

```
avg(/Zabbix server/zabbix[wcache,values],10m)
```

Calculate a 10-minute average of the number of values processed by Zabbix.

(5) Total incoming traffic for host

```
sum(last_foreach(/host/net.if.in[*]))
```

Calculate the total of all items matching `net.if.in[*]` on the host.

(6) Total disk space for group

```
sum(last_foreach(/*/vfs.fs.size[/,total]?[group="MySQL Servers"]))
```

Calculate the total disk space for all `MySQL Servers`.

(7) Average processor load for group

```
avg(last_foreach(/*/system.cpu.load[,avg1]?[group="MySQL Servers"]))
```

Calculate the average processor load for all `MySQL Servers`.

(8) Average processor load for group and tag

```
avg(last_foreach(/*/system.cpu.load?[(group="Servers A" or group="Servers B" or group="Servers C") and (ta
```

Calculate the average CPU load on all hosts in multiple host groups that have specific tags.

(9) Average number of queries for group

```
avg(avg_foreach(/*/mysql.qps?[group="MySQL Servers"],5m))
```

Calculate the 5-minute average of the number of queries per second for `MySQL Servers`.

(10) Total number of successful DNS checks

```
sum(last_foreach(/*/net.dns[*,*,*]))
```

Calculate the total number of successful DNS checks across all hosts. The shown item form is `net.dns[192.0.2.1,example.com,A]` as an example of one possible key.

Note that wildcards must match the number of parameters in the key (here `net.dns` has three parameters: ip, name, type).

(11) Total number of unsupported items for group

```
sum(last_foreach(/*/zabbix[host,,items_unsupported]?[group="Zabbix servers"]))
```

Calculate the total number of unsupported items for `Zabbix servers`.

(12) Interface stats for group

```
sum(last_foreach(/*/net.if.out[eth0,bytes]?[group="video"])) / sum(last_foreach(/*/nginx_stat.sh[active]?[
```

Calculate the interface stats for `video` group.

Examples of correct/incorrect syntax

Expressions (including function calls) cannot be used as history, trend, or foreach function parameters. However, those functions themselves can be used in other (non-historical) function parameters.

Expression	Example
Valid	`avg(last(/host/key1),last(/host/key2)*10,last(/host/key1)*100)`
	`max(avg(avg_foreach(/*/system.cpu.load?[group="Servers A"],5m)),avg(avg_foreach(/*/system.cpu.load?[group="Servers B"],5m)),avg(avg_foreach(/*/system.cpu.load?[group="Servers C"],5m)))`
Invalid	`sum(/host/key,10+2)`
	`sum(/host/key, avg(10,2))`
	`sum(/host/key,last(/host/key2))`

Note that in an expression like:

```
sum(sum_foreach(//resptime[*],5m))/sum(count_foreach(//resptime[*],5m))
```

(where `resptime[*]` is an example of an item key pattern matching the keys `resptime[1]`, `resptime[2]`, etc., used to acquire the response time from several active cluster nodes), it cannot be guaranteed that both parts of the equation will always have the same set of values. While one part of the expression is evaluated, a new value for the requested period may arrive and then the other part of the expression will have a different set of values.

Dependent item

Overview

There are situations when one item gathers multiple metrics at a time or it even makes more sense to collect related metrics simultaneously, for example:

- CPU utilization of individual cores
- Incoming/outgoing/total network traffic

To allow for bulk metric collection and simultaneous use in several related items, Zabbix supports dependent items. Dependent items depend on the master item that collects their data simultaneously, in one query. A new value for the master item automatically populates the values of the dependent items. Dependent items cannot have a different update interval than the master item.

Zabbix preprocessing options can be used to extract the part that is needed for the dependent item from the master item data.

Preprocessing is managed by a `preprocessing manager` process, along with worker threads that perform the preprocessing steps. All values with preprocessing, received from different data gatherers, pass through the preprocessing manager before being added to the history cache. Socket-based IPC communication is used between data gatherers (pollers, trappers, etc) and the preprocessing process.

Zabbix server or Zabbix proxy (if host is monitored by proxy) are performing preprocessing steps and processing dependent items.

Items of any type, even dependent items, can be set as master items. Additional levels of dependent items can be used to extract smaller parts from the value of an existing dependent item.

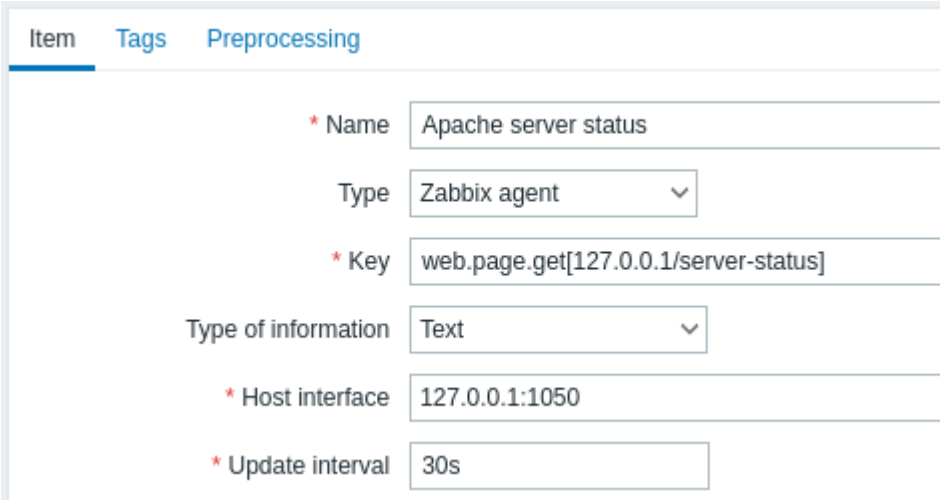
Limitations

- Only same host (template) dependencies are allowed
- An item prototype can depend on another item prototype or regular item from the same host
- Dependent item on a host with master item from template will not be exported to XML

Item configuration

A dependent item depends on its master item for data. That is why the **master item** must be configured (or exist) first:

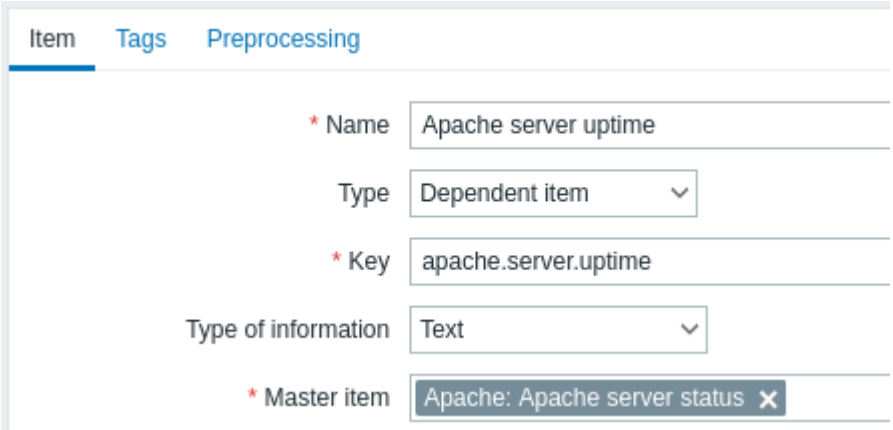
1. Go to *Data collection* > *Hosts*.
2. Click *Items* in the row of the host.
3. Click *Create item*.
4. Enter parameters of the item in the form.



All mandatory input fields are marked with a red asterisk.

Click *Add* to save the master item.

Then you can configure a **dependent item**.

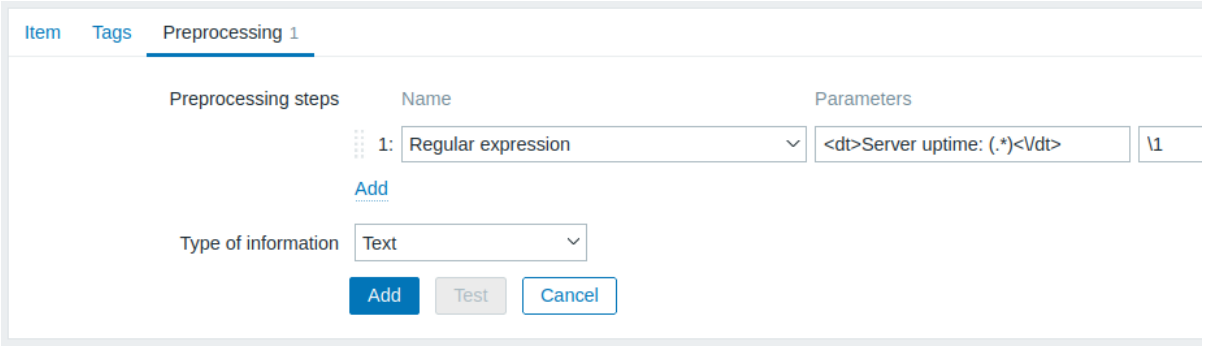


All mandatory input fields are marked with a red asterisk.

The fields that require specific information for dependent items are:

Type	Select **Dependent item** here.
Key	Enter a key that will be used to recognize the item.
Master item	Select the master item. Master item value will be used to populate dependent item value.
Type of information	Select the type of information that will correspond the format of data that will be stored.

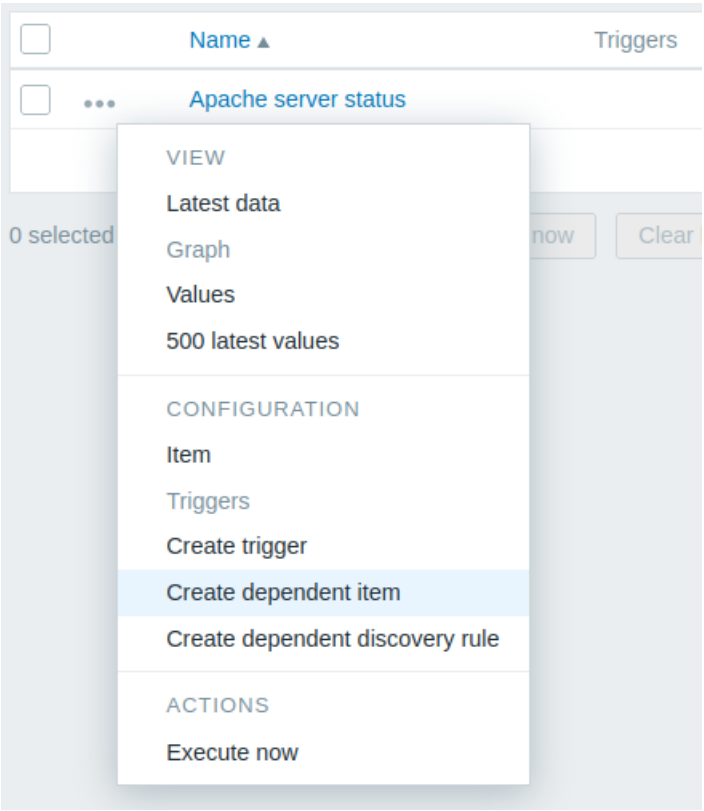
You may use item value preprocessing to extract the required part of the master item value.



Without preprocessing, the dependent item value will be exactly the same as the master item value.

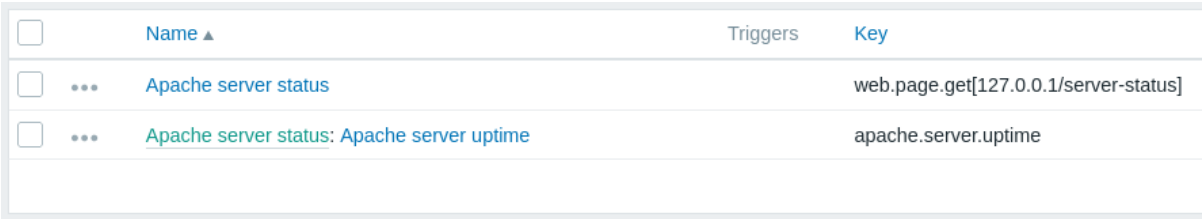
Click *Add* to save the dependent item.

A shortcut to creating a dependent item quicker can be accessed by clicking the ••• button in the item list and selecting *Create dependent item*.



Display

In the item list dependent items are displayed with their master item name as prefix.



If a master item is deleted, so are all its dependent items.

Script item

Overview

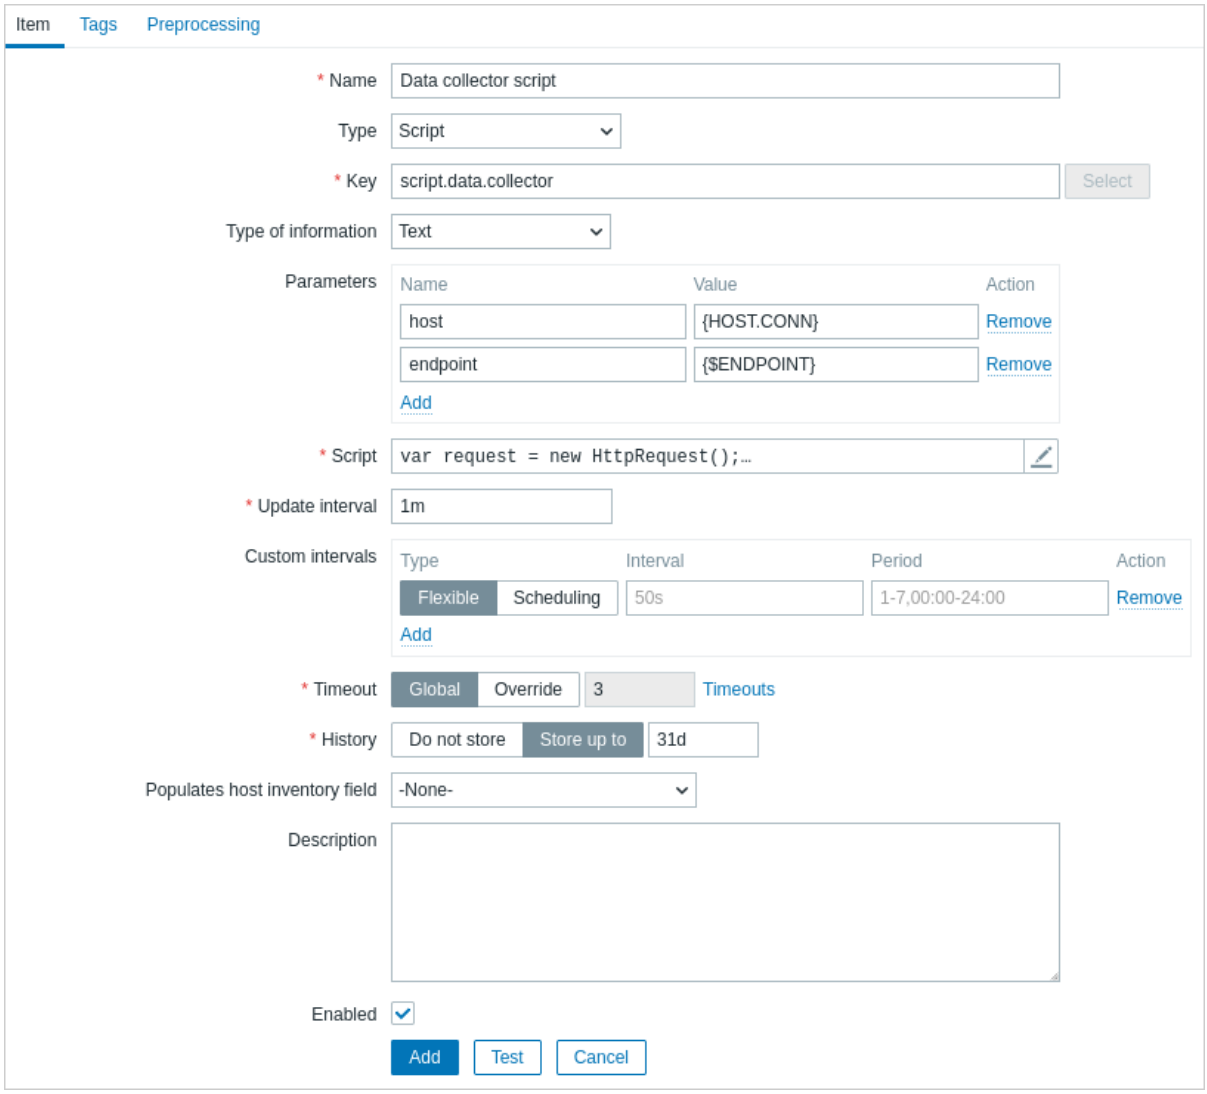
Script items can be used to collect data by executing a user-defined JavaScript code with the ability to retrieve data over HTTP/HTTPS. In addition to the script, an optional list of parameters (pairs of name and value) and timeout can be specified.

This item type may be useful in data collection scenarios that require multiple steps or complex logic. As an example, a Script item can be configured to make an HTTP call, then process the data received in the first step in some way, and pass transformed value to the second HTTP call.

Script items are processed by Zabbix server or proxy pollers.

Configuration

In the *Type* field of item configuration form select Script then fill out required fields.



All mandatory input fields are marked with a red asterisk.

The fields that require specific information for Script items are:

Field	Description
Key	Enter a unique key that will be used to identify the item.
Parameters	Specify the variables to be passed to the script as the attribute and value pairs. User macros are supported. To see which built-in macros are supported, do a search for "Script-type item" in the supported macro table.
Script	Enter JavaScript code in the modal editor that opens when clicking in the parameter field or on the pencil icon next to it. This code must provide the logic for returning the metric value. The code has access to all parameters and all additional JavaScript objects added by Zabbix. See also: JavaScript Guide.
Timeout	JavaScript execution timeout (1-600s; exceeding it will return an error). Note that depending on the script, it might take longer for the timeout to trigger. For more information on the *Timeout* parameter, see general item attributes.

Examples

Simple data collection

Collect the content of *https://www.example.com/release_notes*:

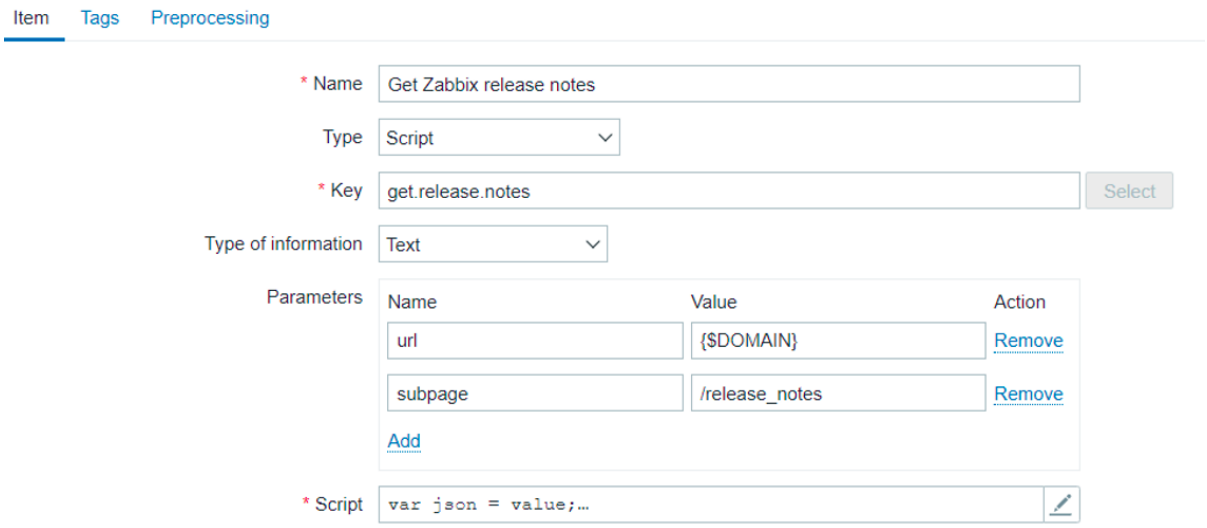
- Create an item with type "Script".
- In the *Script* field, enter:

```
var request = new HttpRequest();
return request.get("https://www.example.com/release_notes");
```

Data collection with parameters

Collect the content of a specific page and make use of parameters:

- Create an item with type "Script" and two parameters:
  - **url : {$DOMAIN}** (the user macro {$DOMAIN} should be defined, preferably on the host level)
  - **subpage : /release_notes**



- In the *Script* field, enter:

```
var obj = JSON.parse(value);
var url = obj.url;
var subpage = obj.subpage;
var request = new HttpRequest();
return request.get(url + subpage);
```

Multiple HTTP requests

Collect the content of both *https://www.example.com* and *https://www.example.com/release_notes*:

- Create an item with type "Script".
- In the *Script* field, enter:

```
var request = new HttpRequest();
return request.get("https://www.example.com") + request.get("https://www.example.com/release_notes");
```

Logging

Add the "Log test" entry to the Zabbix server log and receive the item value "1" in return:

- Create an item with type "Script".
- In the *Script* field, enter:

```
Zabbix.log(3, 'Log test');
return 1;
```

Browser item

Overview

Browser items allow monitoring complex websites and web applications using a browser.

Attention:
The support of Browser items is currently experimental.

Browser items collect data by executing a user-defined JavaScript code and retrieving data over HTTP/HTTPS. This item can simulate such browser-related actions as clicking, entering text, navigating through web pages, and other user interactions with websites or web applications.

In addition to the script, an optional list of parameters (pairs of name and value) and timeout can be specified.

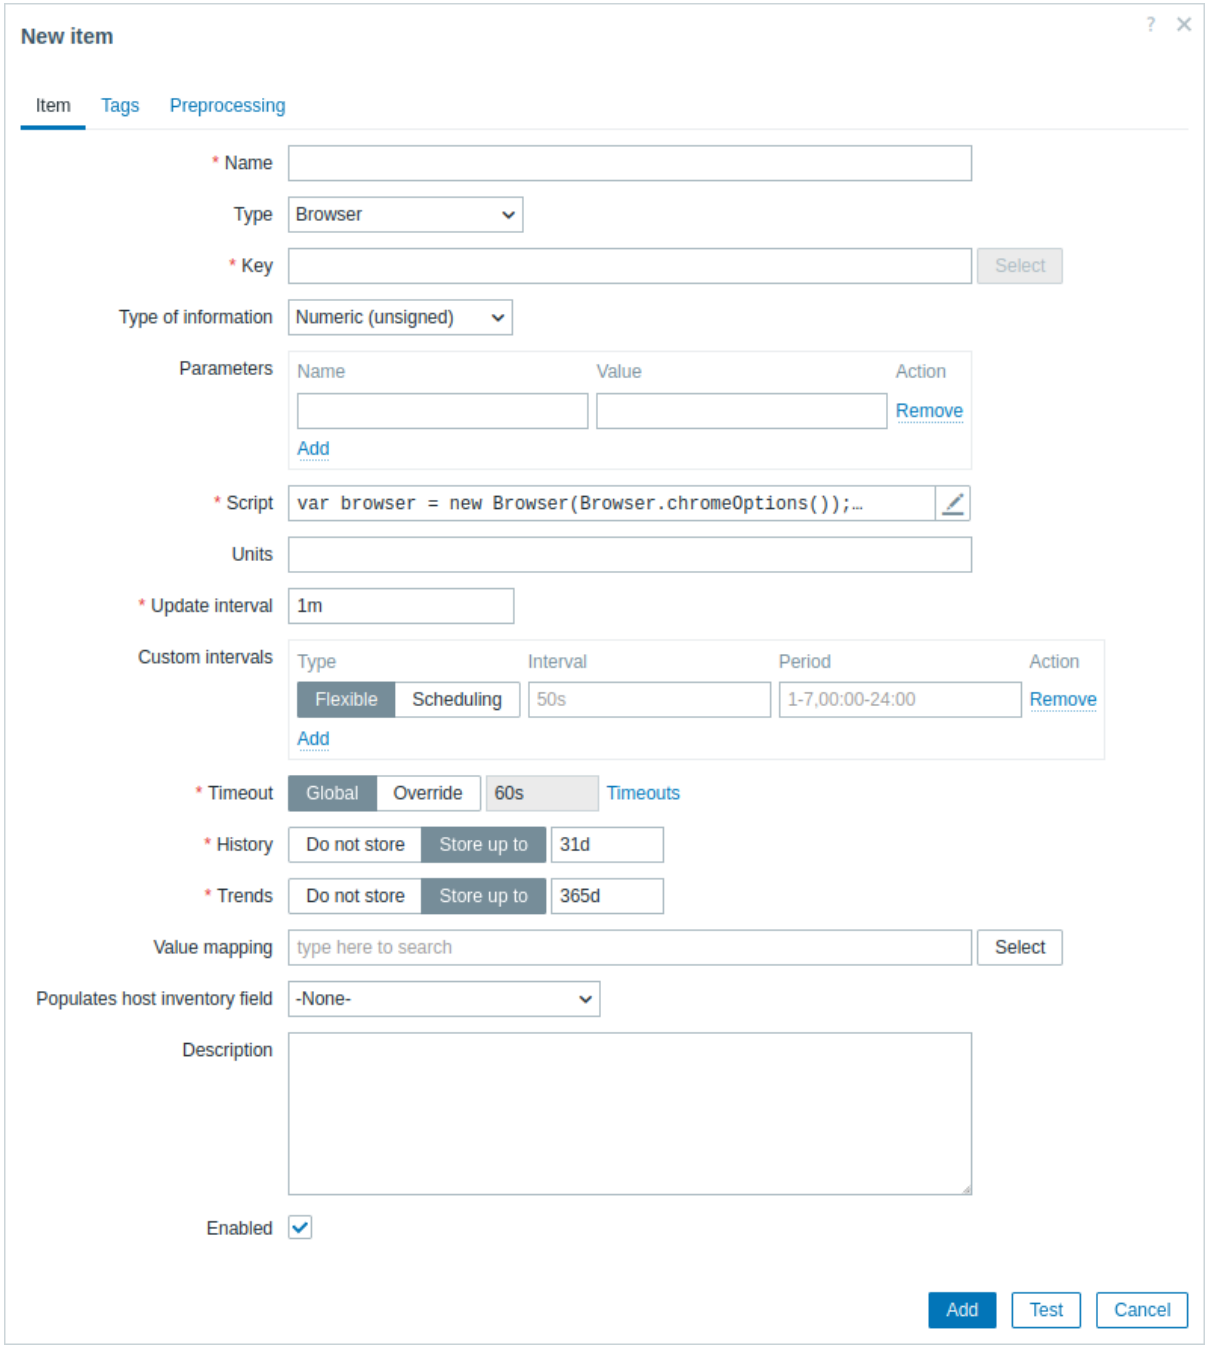
Attention:
The item partially implements the W3C WebDriver standard with either Selenium Server or a plain WebDriver (for example, ChromeDriver) as a web testing endpoint. For the item to work, set the endpoint in the Zabbix server/proxy configuration file `WebDriverURL` parameter (if using ChromeDriver, see Security Considerations). For better performance, consider using a dedicated server for the web testing environment.

Browser item checks are executed and processed by Zabbix server or proxy browser pollers. If necessary, you can adjust the number of pre-forked instances of browser pollers in the Zabbix server/proxy configuration file `StartBrowserPollers` parameter.

For monitoring complex websites and web applications, the Website by Browser template is available as an out-of-the-box template.

Configuration

In the *Type* field of item configuration form, select Browser and then fill out the required fields.



All mandatory input fields are marked with a red asterisk.

The fields that require specific information for Browser items are:

Field	Description
Key	Enter a unique key that will be used to identify the item.
Parameters	Specify the variables to be passed to the script as the attribute and value pairs. User macros are supported. To see which built-in macros are supported, do a search for "Browser-type item" in the supported macro table.
Script	Enter JavaScript code in the modal editor that opens when clicking in the parameter field or on the pencil icon next to it. This code must provide the logic for returning the metric value. The code has access to all parameters, all additional JavaScript objects and Browser item JavaScript objects added by Zabbix. See also: JavaScript Guide.
Timeout	JavaScript execution timeout (1-600s; exceeding it will return an error). Note that depending on the script, it might take longer for the timeout to trigger. For more information on the *Timeout* parameter, see general item attributes.

Examples

For an example of how to set up Zabbix for monitoring websites using the *Website by Browser* template, see Monitor websites with Browser items.

Default script

The following script:

1. Initializes a browser session.
2. Navigates to a specified URL.
3. Collects performance entries and session statistics, and returns them as a JSON string.

In the *Script* field, enter:

```
var browser = new Browser(Browser.chromeOptions());

try {
    browser.navigate("http://example.com");
    browser.collectPerfEntries();
}
finally {
    return JSON.stringify(browser.getResult());
}
```

Initialize browser with custom capabilities

The following script:

1. Initializes a browser session for the available browser based on the first matching browser in the order specified within the script.
2. Defines browser capabilities, including page load strategy and options specific to each browser, such as the headless mode for Chrome, Firefox, and Microsoft Edge browsers.

In the *Script* field, enter:

```
var browser = new Browser({
    "capabilities":{
        "firstMatch":[
            {
                "browserName":"chrome",
                "pageLoadStrategy":"normal",
                "goog:chromeOptions":{
                    "args":[
                        "--headless=new"
                    ]
                }
            },
            {
                "browserName":"firefox",
                "pageLoadStrategy":"normal",
                "moz:firefoxOptions":{
                    "args":[
                        "--headless"
                    ]
                }
            },
            {
                "browserName":"MicrosoftEdge",
                "pageLoadStrategy":"normal",
                "ms:edgeOptions":{
                    "args":[
                        "--headless=new"
                    ]
                }
            },
            {
                "browserName":"safari",
                "pageLoadStrategy":"normal"
            }
```

```
        ]
    }
});
```

Initialize browser with GUI

By default, browser sessions (excluding Safari) are initialized in headless mode, meaning the browser's graphical user interface (GUI) is not displayed.

The following script initializes a browser session with the GUI enabled.

Note that if the WebDriver cannot locate the browser binary, you can specify the path manually.

```
var opts = Browser.chromeOptions();
opts.capabilities.alwaysMatch['goog:chromeOptions'].args = [];

// To initialize a Firefox session with GUI, uncomment the following lines:
// var opts = Browser.firefoxOptions();
// opts.capabilities.alwaysMatch['moz:firefoxOptions'].binary = 'usr/bin/firefox';
// opts.capabilities.alwaysMatch['moz:firefoxOptions'].args = [];

// To initialize a Microsoft Edge session with GUI, uncomment the following lines:
// var opts = Browser.edgeOptions();
// opts.capabilities.alwaysMatch['ms:edgeOptions'].binary = 'usr/bin/microsoft-edge';
// opts.capabilities.alwaysMatch['ms:edgeOptions'].args = [];

var browser = new Browser(opts);
```

Note:
If your tests are running on a remote server or in a container, you can use a Virtual Network Computing (VNC) client to connect to the machine's VNC server. This lets you view and interact with the browser's GUI remotely.

Take screenshots

The following script:

1. Initializes a browser session.
2. Sets the browser's viewport size to determine the screenshot size (specified as parameters, see below).
3. Navigates to a URL (specified as a parameter, see below).
4. Collects session statistics, captures a screenshot, and adds it to the collected statistics.
5. Handles errors by capturing error messages and a screenshot.
6. Returns the collected results as a JSON string.

The script also uses parameters from the item configuration form:

- webURL - http://example.com
- width - 1920
- height - 1080

In the *Script* field, enter:

```
var browser, result;

var browser = new Browser(Browser.chromeOptions());

try {
    var params = JSON.parse(value); // Parse the JSON string containing parameters passed from Zabbix.

    browser.setScreenSize(Number(params.width), Number(params.height))

    browser.navigate(params.webURL);

    result = browser.getResult();
    result.screenshot = browser.getScreenshot();
}
catch (err) {
    if (!(err instanceof BrowserError)) {
```

```
        browser.setError(err.message);
    }
    result = browser.getResult();
    result.error.screenshot = browser.getScreenshot();
}
finally {
    return JSON.stringify(result);
}
```

Check Zabbix login

The following script:

1. Initializes a browser session.
2. Navigates to a page (specified as a parameter, see below).
3. Enters the username and password (specified as parameters, see below).
4. Finds and clicks the login button.
5. Finds and clicks the logout button.
6. Collects performance data before and after login, as well as after logout.
7. Handles errors by capturing error messages and a screenshot.
8. Returns the collected results as a JSON string.

The script also uses parameters from the item configuration form:

- webURL - http://{HOST.CONN}/index.php
- username - {$USERNAME}
- password - {$PASSWORD}

In the *Script* field, enter:

```
var browser, result;

browser = new Browser(Browser.chromeOptions());

try {
    var params = JSON.parse(value); // Parse the JSON string containing parameters passed from Zabbix.

    browser.navigate(params.webURL);
    browser.collectPerfEntries("open page");

    var el = browser.findElement("xpath", "//input[@id='name']");
    if (el === null) {
        throw Error("cannot find name input field");
    }
    el.sendKeys(params.username);

    el = browser.findElement("xpath", "//input[@id='password']");
    if (el === null) {
        throw Error("cannot find password input field");
    }
    el.sendKeys(params.password);

    el = browser.findElement("xpath", "//button[@id='enter']");
    if (el === null) {
        throw Error("cannot find login button");
    }
    el.click();

    browser.collectPerfEntries("login");

    el = browser.findElement("link text", "Sign out");
    if (el === null) {
        throw Error("cannot find logout button");
    }
    el.click();
```

```
    browser.collectPerfEntries("logout");

    result = browser.getResult();
}
catch (err) {
    if (!(err instanceof BrowserError)) {
        browser.setError(err.message);
    }
    result = browser.getResult();
    result.error.screenshot = browser.getScreenshot();
}
finally {
    return JSON.stringify(result);
}
```

Find links

The following script:

1. Initializes a browser session.
2. Defines a function for removing duplicate elements from an array (see step 5).
3. Navigates to a page (specified as parameters, see below).
4. Finds links on the page.
5. Removes duplicate links to ensure they are unique.
6. Extracts only the links that start with "http".
7. Formats the extracted links to a specific structure.
8. Handles errors by capturing error messages and a screenshot.
9. Returns the collected results as a JSON string.

The script also uses parameters from the item configuration form:

- scheme - {$WEBSITE.SCHEME}
- domain - {$WEBSITE.DOMAIN}
- path - {$WEBSITE.PATH}

In the *Script* field, enter:

```
var browser, result;

browser = new Browser(Browser.chromeOptions());

try {
    var params = JSON.parse(value);  // Parse the JSON string containing parameters passed from Zabbix.

    function uniq(a) {
        return a.sort().filter(function (item, pos, ary) {
            return !pos || item != ary[pos - 1];
        });
    }

    browser.navigate(params.scheme + '://' + params.domain + params.path);

    var el = browser.findElements("link text", "");
    var links = [];
    for (var n = 0; n < el.length; n++) {
        links.push(el[n].getAttribute('href'));
    }

    links = uniq(links);

    result = [];
    for (i = 0; i < links.length; i++) {
        if (links[i].match(/^http.*/)) {
            var row = {};
            row["{#URL}"] = links[i];
            result.push(row);
```

```
        }
    }
}
catch (err) {
    if (!(err instanceof BrowserError)) {
        browser.setError(err.message);
    }
    result = browser.getResult();
    result.error.screenshot = browser.getScreenshot();
}
finally {
    return JSON.stringify(result);
}
```

Telemetry query

This item type lets you retrieve or aggregate telemetry data. Telemetry data is used for application performance monitoring (APM), which helps you see how requests flow through your applications and services, measure their performance, and find the cause of slow or failed operations.

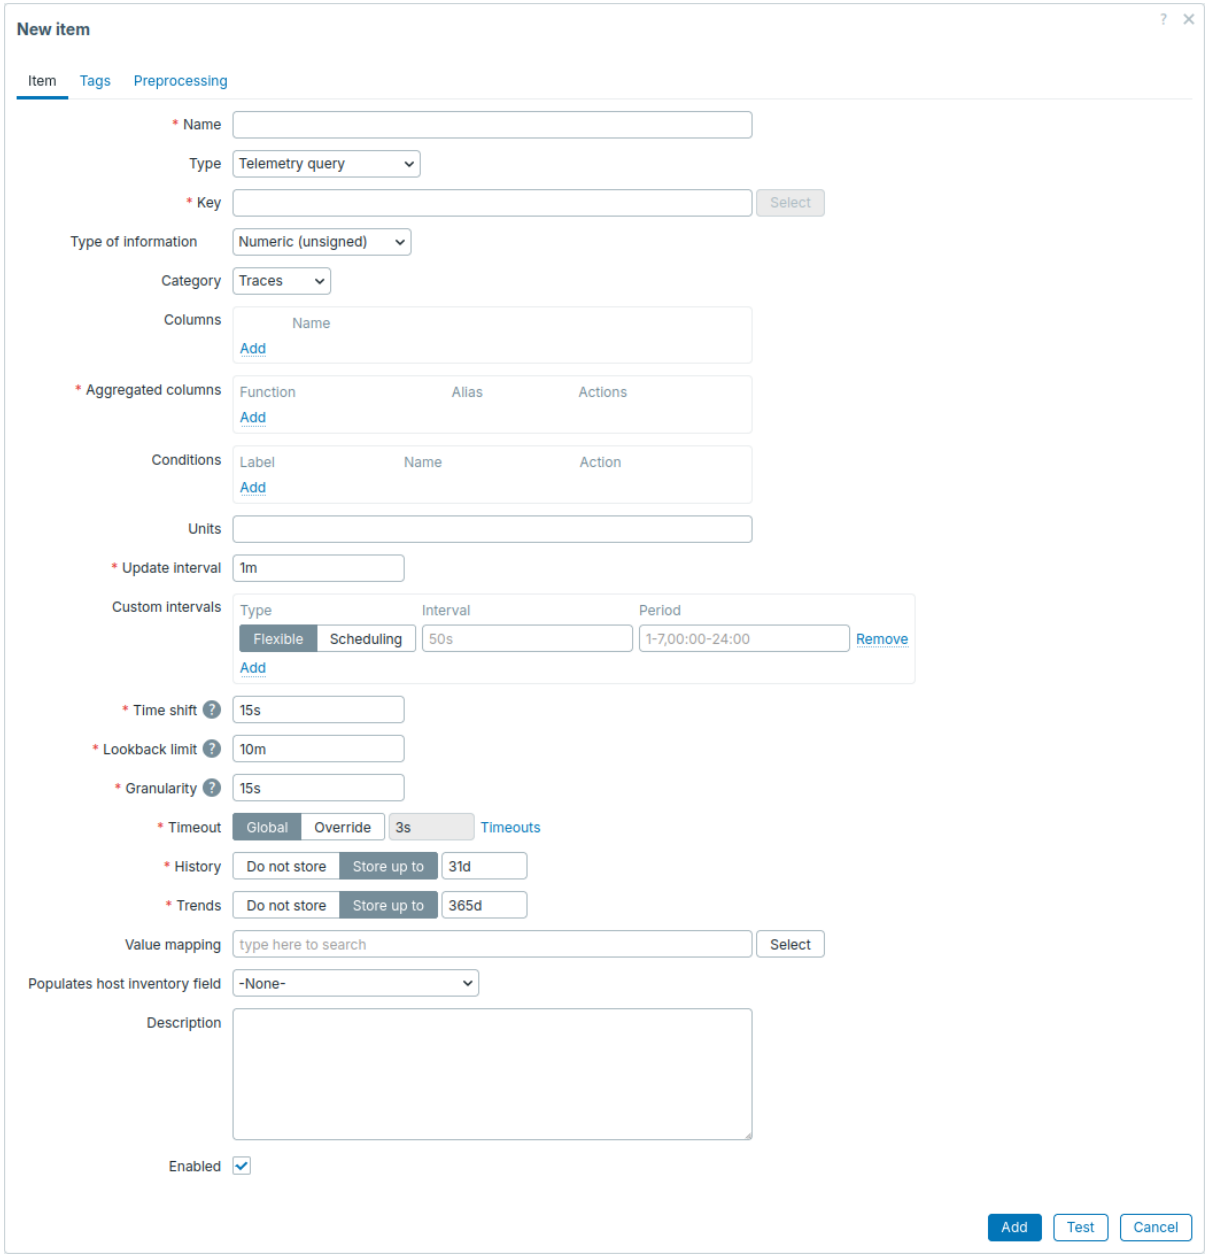
Telemetry data comes in three types: traces (the path that a request takes through your applications and services), metrics (measurements, such as response time), and logs (text records of events). By using the telemetry query item, you can, for example, find all traces where a service returned an error, get the average response time from your metrics, or search your logs for a specific error message.

When you configure a telemetry query item, you define a set of instructions (a query) that tells Zabbix which telemetry data to ask for, how to filter it, and how to aggregate it. When the item check is executed, Zabbix sends this query to the telemetry data source and retrieves the result based on the query.

To use this item, specify the telemetry data source and set the `StartTelemetryQueryPollers` parameter in the Zabbix server or Zabbix proxy configuration, depending on which of them monitors the item's host.

Configuration

When creating an item, set its type to *Telemetry query*:



Parameter	Description
Name	See parameters common to all items.
Type	Select *Telemetry query*.
Key	See parameters common to all items.
Type of information	
Category	Select the signal type that you want to query: *Traces*, *Metrics*, or *Logs*.
Metric points	Select the metric point type that you want to query: *Sum*, *Gauge*, *Histogram*, or *Exponential histogram*. This parameter is available if *Category* is set to *Metrics*.
Columns	Add telemetry data columns that you want to query. For columns that hold key-value pairs, also enter the key name. The records from these columns will be included in the item value as is.
Aggregated columns	Add telemetry data columns that you want to query, and how to calculate a value from their records: **Function** - select the function to apply to the column records: *min*, *max*, *avg*, *count*, *sum*, or *percentile*. **Percentage** - enter the percentile to calculate from the column records (if *Function* is set to *percentile*). **Column** - select the column to calculate from (if *Function* is not set to *count*). **Alias** - enter a unique name for the calculated value, used to identify it in the item value.

Parameter	Description
Type of calculation	Select how to evaluate the conditions set in the *Conditions* parameter below: **And/Or** - all conditions must be met; conditions with the same *Type* and *Attribute key* are combined with *Or*. **And** - all conditions must be met. **Or** - at least one condition must be met. **Custom expression** - conditions must be met as specified in the formula, for example, `A and (B or C)`. This parameter is available if you have set two or more conditions.
Conditions	Add conditions for filtering telemetry data column records: **Type** - select the telemetry data column that the condition will apply to. **Attribute key** - enter the key that the condition will apply to (if *Type* is set to a column with key-value pairs). **Operator** - select the comparison to use (see below) between the record's value and *Attribute value*. **Attribute value** - enter the value to compare the record's value against (if *Operator* is not set to *Exists*). If *Type* is set to a single-value column, the condition will apply to the column's record; if the column holds an array of values, the condition will match the record if at least one value in the array meets the condition. If *Type* is set to a column that holds key-value pairs, the condition will apply to the value of the key specified in *Attribute key*; if the column holds an array of key-value pairs, the condition will match the record if at least one pair with the specified *Attribute key* meets the condition. For single-value columns, you can use these operators: **Equals** - match records that equal *Attribute value* (case-sensitive). **Does not equal** - match records that do not equal *Attribute value* (case-sensitive). **Contains** - match records that contain *Attribute value* (case-sensitive). **Does not contain** - match records that do not contain *Attribute value* (case-sensitive). For columns that hold key-value pairs, you can use these operators: **Equals** - match records where *Attribute key* exists and its value equals *Attribute value* (case-sensitive). **Does not equal** - match records where *Attribute key* exists and its value does not equal *Attribute value* (case-sensitive). **Exists** - match records where *Attribute key* exists (case-sensitive).
Units *Update interval* *Custom intervals*	See parameters common to all items.
Time shift	Enter how far back (before the current time) to end the time range that the item queries. Supports a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., 30s, 1m, 2h, 1d), or user macros. Each query requests data for a time range. The range usually starts where the previous query's time range ended, and ends at the point set by *Time shift*. If the time range of the previous query ended later than the current time minus *Granularity*, then the current query's time range starts at the current time minus *Granularity*, thus overlapping with the previous query's time range. If there is no previous query to continue from, the range starts at the point set by *Lookback limit*. If this item's host becomes monitored by a different entity (e.g., a different proxy from a proxy group), the item may no longer know when the previous query ended. This can cause duplicate values.
Lookback limit	Enter how far back (before the current time) the item queries for data when there is no previous query to continue from (e.g., on the item's first run), or when the previous query is older than *Update interval* (e.g., because Zabbix server was down). Supports a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., 30s, 1m, 2h, 1d), or user macros. If the telemetry data source holds a large amount of data (e.g., over a million records per hour), a large *Lookback limit* may significantly reduce performance when the item uses it.

Parameter	Description
Granularity	Enter the length of a piece that the query's full time range is split into. Supports a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., 30s, 1m, 2h, 1d), or user macros.
Timeout	See parameters common to all items.
History	
Trends	
Value mapping	
Populates host inventory field	
Description	
Enabled	

Supported telemetry data columns

This section lists the columns supported in these item parameters: *Columns*, *Aggregated columns*, and *Conditions*. These are the columns used by ClickHouse to store telemetry data, following the OpenTelemetry ClickHouse exporter schema.

Select the *Category*:

- Traces
- Metrics: Sum
- Metrics: Gauge
- Metrics: Histogram
- Metrics: Exponential histogram
- Logs

Traces

Column	*Columns*	*Aggregated columns*	*Conditions*
Duration	✓	✓	–
Events.Attributes	–	–	✓
Events.Name	–	–	✓
ParentSpanId	✓	–	✓
ResourceAttributes	✓	–	✓
ScopeName	✓	–	✓
ScopeVersion	✓	–	✓
ServiceName	✓	–	✓
SpanAttributes	✓	–	✓
SpanId	✓	–	✓
SpanKind	✓	–	✓
SpanName	✓	–	✓
StatusCode	✓	–	✓
StatusMessage	✓	–	✓
Timestamp	✓	✓	–
TraceId	✓	–	✓
TraceState	✓	–	✓

Metrics: Sum

Column	*Columns*	*Aggregated columns*	*Conditions*
AggregationTemporality	✓	–	–
Attributes	✓	–	✓
Exemplars.FilteredAttributes	–	–	✓
Flags	✓	–	–
IsMonotonic	✓	–	–
MetricDescription	✓	–	✓
MetricName	✓	–	✓
MetricUnit	✓	–	✓
ResourceAttributes	✓	–	✓
ResourceSchemaUrl	✓	–	✓
ScopeAttributes	✓	–	✓
ScopeDroppedAttrCount	✓	✓	–

Column	*Columns*	*Aggregated columns*	*Conditions*
ScopeName	✓	–	✓
ScopeSchemaUrl	✓	–	✓
ScopeVersion	✓	–	✓
ServiceName	✓	–	✓
StartTimeUnix	✓	✓	–
TimeUnix	✓	✓	–
Value	✓	✓	–

Metrics: Gauge

Column	*Columns*	*Aggregated columns*	*Conditions*
Attributes	✓	–	✓
Exemplars.FilteredAttributes	–	–	✓
Flags	✓	–	–
MetricDescription	✓	–	✓
MetricName	✓	–	✓
MetricUnit	✓	–	✓
ResourceAttributes	✓	–	✓
ResourceSchemaUrl	✓	–	✓
ScopeAttributes	✓	–	✓
ScopeDroppedAttrCount	✓	✓	–
ScopeName	✓	–	✓
ScopeSchemaUrl	✓	–	✓
ScopeVersion	✓	–	✓
ServiceName	✓	–	✓
StartTimeUnix	✓	✓	–
TimeUnix	✓	✓	–
Value	✓	✓	–

Metrics: Histogram

Column	*Columns*	*Aggregated columns*	*Conditions*
AggregationTemporality	✓	–	–
Attributes	✓	–	✓
Count	✓	✓	–
Exemplars.FilteredAttributes	–	–	✓
Flags	✓	–	–
Max	✓	✓	–
MetricDescription	✓	–	✓
MetricName	✓	–	✓
MetricUnit	✓	–	✓
Min	✓	✓	–
ResourceAttributes	✓	–	✓
ResourceSchemaUrl	✓	–	✓
ScopeAttributes	✓	–	✓
ScopeDroppedAttrCount	✓	✓	–
ScopeName	✓	–	✓
ScopeSchemaUrl	✓	–	✓
ScopeVersion	✓	–	✓
ServiceName	✓	–	✓
StartTimeUnix	✓	✓	–
Sum	✓	✓	–
TimeUnix	✓	✓	–

Metrics: Exponential histogram

Column	*Columns*	*Aggregated columns*	*Conditions*
AggregationTemporality	✓	–	–
Attributes	✓	–	✓
Count	✓	✓	–
Exemplars.FilteredAttributes	–	–	✓
Flags	✓	–	–
Max	✓	✓	–
MetricDescription	✓	–	✓
MetricName	✓	–	✓
MetricUnit	✓	–	✓
Min	✓	✓	–
NegativeOffset	✓	✓	–
PositiveOffset	✓	✓	–
ResourceAttributes	✓	–	✓
ResourceSchemaUrl	✓	–	✓
Scale	✓	✓	–
ScopeAttributes	✓	–	✓
ScopeDroppedAttrCount	✓	✓	–
ScopeName	✓	–	✓
ScopeSchemaUrl	✓	–	✓
ScopeVersion	✓	–	✓
ServiceName	✓	–	✓
StartTimeUnix	✓	✓	–
Sum	✓	✓	–
TimeUnix	✓	✓	–
ZeroCount	✓	✓	–

Logs

Column	*Columns*	*Aggregated columns*	*Conditions*
Body	✓	–	✓
EventName	✓	–	✓
LogAttributes	✓	–	✓
ResourceAttributes	✓	–	✓
ResourceSchemaUrl	✓	–	✓
ScopeAttributes	✓	–	✓
ScopeName	✓	–	✓
ScopeSchemaUrl	✓	–	✓
ScopeVersion	✓	–	✓
SpanId	✓	–	✓
ServiceName	✓	–	✓
SeverityNumber	✓	✓	–
SeverityText	✓	–	✓
Timestamp	✓	✓	–
TraceFlags	✓	–	–
TraceId	✓	–	✓

History and trends

Overview

History and trends are the two ways of storing collected data in Zabbix.

Whereas history keeps each collected value, trends keep averaged information on hourly basis and therefore are less resource-hungry.

Keeping history

You can set for how many days history will be kept:

- in the item properties form
- when mass-updating items

- when setting up housekeeper tasks

Any older data will be removed by the housekeeper.

The general strong advice is to keep history for the smallest possible number of days and that way not to overload the database with lots of historical values.

Instead of keeping a long history, you can keep longer data of trends. For example, you could keep history for 14 days and trends for 5 years.

You can get a good idea of how much space is required by history versus trends data by referring to the database sizing page.

While keeping shorter history, you will still be able to review older data in graphs, as graphs will use trend values for displaying older data.

Attention:
If history is set to '0', the item will update only dependent items and inventory. No trigger functions will be evaluated because trigger evaluation is based on history data only.

Note:
As an alternative way to preserve history consider to use history export functionality of loadable modules.

Keeping trends

Trends is a built-in historical data reduction mechanism which stores minimum, maximum, average and the total number of values per every hour for numeric data types.

You can set for how many days trends will be kept:

- in the item properties form
- when mass-updating items
- when setting up Housekeeper tasks

Trends usually can be kept for much longer than history. Any older data will be removed by the housekeeper.

Zabbix server accumulates trend data in runtime in the trend cache, as the data flows in. Server flushes **previous hour** trends of every item into the database (where frontend can find them) in these situations:

- server receives the first current hour value of the item
- 5 or less minutes of the current hour left and still no current hour values of the item
- server stops

To see trends on a graph you need to wait at least to the beginning of the next hour (if item is updated frequently) and at most to the end of the next hour (if item is updated rarely), which is 2 hours maximum.

When server flushes trend cache and there are already trends in the database for this hour (for example, server has been restarted mid-hour), server needs to use update statements instead of simple inserts. Therefore on a bigger installation if restart is needed it is desirable to stop server in the end of one hour and start in the beginning of the next hour to avoid trend data overlap.

History tables do not participate in trend generation in any way.

Attention:
If trends are set to '0', Zabbix server does not calculate or store trends at all.

Note:
The trends are calculated and stored with the same data type as the original values. As a result the average value calculations of unsigned data type values are rounded and the less the value interval is the less precise the result will be. For example, if item has values 0 and 1, the average value will be 0, not 0.5.
Also restarting server might result in the precision loss of unsigned data type average value calculations for the current hour.

User parameters

Overview

A user parameter lets you define a custom item key and the command used to retrieve its value when no built-in Zabbix agent item provides the data you need. A user parameter is a user-defined command referenced by a user-defined item key and executed by the Zabbix agent when the corresponding item is checked.

Note:
On UNIX systems, the agent executes user parameter commands using `/bin/sh`, under the operating system account running the agent (usually `zabbix`). That account must have sufficient permissions to execute the command and access any files, sockets, or credentials it uses.

There are two types of user parameters:

- **Simple user parameters** do not accept arguments; they define a fixed item key and a command to execute.
- **Flexible user parameters** accept arguments (specified as parameters in the item key) and pass them to the command, allowing a single definition to serve multiple items.

User parameter configuration

A user parameter is defined using the `UserParameter` configuration parameter.

A user parameter has the following syntax:

`UserParameter=<key>,<shell command>`

where:

- `<key>` - a unique item key that identifies the user parameter.
 When configuring an item, specify this key to reference the command defined by the user parameter.
 For a simple user parameter, the key is fixed and does not accept arguments. For a flexible user parameter, append `[*]` to the key (for example, `mysql.status[*]`) to indicate that it accepts arguments. Specify these arguments as parameters in the item key when configuring the item. The command can refer to the parameters using positional references $1-$9.
- `<shell command>` - a shell command or a path to a custom script that the agent executes to obtain the item value.
 For the flexible user parameter, the command can use positional references $1-$9 to refer to the corresponding parameters in the item key (for example, `UserParameter=mysql.status[*],mysqladmin status --user=$1 --password=$2`). Zabbix parses the parameters enclosed in `[]` in the item key (i.e., `mysql.status[zabbix,{$MYSQL.PASSWORD}]`) and substitutes the positional references in the command accordingly.
 $0 is substituted with the original command before positional references are expanded. Positional references are interpreted regardless of whether they are enclosed in double (") or single (') quotes. To use positional references without substitution, specify a double dollar sign. For example, `awk '{print $$2}'` passes $2 to the command unchanged.
 Note that positional references prefixed with $ are searched for and replaced by the Zabbix agent only in flexible user parameters. For simple user parameters, positional references are not processed, so escaping $ is unnecessary.

Attention:
Certain characters cannot be passed as arguments by default. For the complete list of characters, see UnsafeUserParameters.

`UserParameter` must be configured on each Zabbix agent that uses it: - Directly in the Zabbix agent configuration file (`zabbix_agentd.conf` or `zabbix_agent2.conf`). - In a separate `*.conf` file, provided that the file is included by the agent configuration.

More than one `UserParameter` can be defined per agent, as multiple files can be included with different sets of parameters. Each user parameter must have a unique item key within the monitored host.

The `UserParameterDir` configuration parameter lets you specify the working directory for executing `UserParameter` commands. If configured, the agent changes its working directory to the specified directory before executing a command. Commands can be specified using a relative path with the `./` prefix, for example, `UserParameter=mykey,./myscript.sh`.

To configure a user parameter:

1. Write a custom script or command to retrieve the required data.
2. Add the command and its unique item key to the `UserParameter` configuration parameter in the Zabbix agent configuration file (`zabbix_agentd.conf`).
3. Reload user parameters from the configuration file by running the runtime control command:

   zabbix_agentd -R userparameter_reload

   Alternatively, restart the agent to apply the configuration changes.
4. Create a new item on the monitored host with the item key defined in `UserParameter`.
 Set the item's type to *Zabbix Agent* or *Zabbix Agent (active)*.
 Make sure that the item's *Type of information* is configured correctly; otherwise, Zabbix will reject the returned value.

See Extending Zabbix agents for an example of configuring and using user parameters.

User parameters are subject to the Zabbix agent check timeout. If the timeout is reached, the process executing the user parameter is terminated. See Command execution for more information.

Note that long-running scripts may consume system resources and delay other checks.

Command result

The command's return value consists of standard output and standard error produced by the command.

> **Note:**
> User parameters that return text (the *Character*, *Log*, or *Text* type of information) can also return a whitespace. If the command produces an invalid result, the item becomes unsupported. However, an item with one of these information types will not become unsupported solely because the command produces output on standard error.

The command's return value is limited to 16 MB, including trailing whitespace, which is truncated. This limit applies before item value preprocessing steps. Database limits also apply.

Examples of user parameters

This section provides examples of simple and flexible user parameters.

> **Attention:**
> The examples below are for demonstration purposes only and should not be used in production without appropriate safeguards. Flexible user parameters defined using `[*]` accept arguments that can introduce command injection risks if not handled safely.

Examples of simple user parameters

Example 1: Constant value

```
UserParameter=ping,echo 1
```

An item with the key `ping` always returns 1.

Example 2: MySQL server availability

```
UserParameter=mysql.ping,mysqladmin ping 2>/dev/null | grep -c alive
```

Returns 1 if the MySQL server responds and 0 otherwise. The `2>/dev/null` redirection discards standard error output from `mysqladmin`. Without it, an error message could be returned together with 0, causing an item with the *Numeric* information type to become unsupported (see Command result).
 This example assumes that MySQL connection credentials are configured for the operating system account running the Zabbix agent, for example, in a MySQL option file accessible to that account.

Examples of flexible user parameters

Example 1: Returning a parameter

```
UserParameter=ping[*],echo $1
```

You can define multiple items using this user parameter, with keys in the format `ping[parameter]`:

- `ping[0]` - returns 0.
- `ping[hello]` - returns `hello`.

Example 2: MySQL availability on a specified host

This parameter can be used to monitor the availability of a MySQL server on a specified host.

```
UserParameter=mysql.ping[*],mysqladmin -h "$1" ping 2>/dev/null | grep -c alive
```

The item with the key `mysql.ping[db1.example.com]` returns 1 if the MySQL server on that host responds and 0 otherwise.
 This example assumes that MySQL connection credentials are configured for the operating system account running the Zabbix agent, for example, in a MySQL option file accessible to that account.

Example 3: Counting lines matching a pattern

This parameter can be used to count the number of lines in a file.

```
UserParameter=grep.count[*],grep -c "$2" "$1"
```

Returns the number of lines in the file specified by the first argument that match the regular expression specified by the second argument. For example, `grep.count[/etc/passwd,root]` returns the number of lines in `/etc/passwd` that match the regular expression `root`.

Extending Zabbix agents

This section provides step-by-step instructions for extending the functionality of a Zabbix agent using a user parameter.

The example below retrieves the number of statements executed by a MySQL server. The steps apply to both Zabbix agent and Zabbix agent 2.

> **Note:**
> This example illustrates how user parameters work. For production MySQL monitoring, consider using the predefined MySQL templates, which do not require user parameters.

Step 1: Prepare access to MySQL

The agent runs as the `zabbix` operating system user. Do not use the MySQL `root` account, as it may use socket-based authentication, which can prevent the agent from connecting. Instead, create a dedicated MySQL user:

```
CREATE USER 'zbx_monitor'@'localhost' IDENTIFIED BY '<password>';
```

Store the credentials in a file readable only by the `zabbix` user, for example `/etc/zabbix/mysql_monitor.cnf` (the path depends on the operating system and Zabbix installation method):

```
[client]
user=zbx_monitor
password=<strong password>
```

Set the file's ownership and permissions:

```
sudo chown zabbix:zabbix /etc/zabbix/mysql_monitor.cnf
sudo chmod 600 /etc/zabbix/mysql_monitor.cnf
```

Pass this file explicitly with `--defaults-extra-file` (which must be the first option passed to the MySQL client) rather than relying on `~/.my.cnf`, because the daemon's `HOME` environment variable may be unset or different from what you expect. If MySQL uses a non-default socket path, add a `socket=` line to the same file.

Step 2: Write a command

Write a custom script or command to retrieve the required parameter, and make sure it works under the same conditions as the agent daemon.

The following command returns the number of statements executed by the MySQL server:

```
/usr/bin/mysql --defaults-extra-file=/etc/zabbix/mysql_monitor.cnf -N -B -e "SHOW GLOBAL STATUS LIKE 'Ques
```

You can test the command as the `zabbix` user; it should print a single number:

```
sudo -u zabbix sh -c "/usr/bin/mysql --defaults-extra-file=/etc/zabbix/mysql_monitor.cnf -N -B -e \"SHOW G
```

Step 3: Add the command to the Zabbix agent configuration

Rather than editing the main agent configuration file (`zabbix_agentd.conf` or `zabbix_agent2.conf`), create a separate file in the directory included by the Include parameter. Separate files are easier to manage and are less likely to be affected by package upgrades.

For example, create `/etc/zabbix/zabbix_agentd.d/mysql.conf` (`/etc/zabbix/zabbix_agent2.d/mysql.conf` for Zabbix agent 2) with the following content:

```
UserParameter=mysql.questions,/usr/bin/mysql --defaults-extra-file=/etc/zabbix/mysql_monitor.cnf -N -B -e
```

`mysql.questions` is the unique item key. It must be unique within the agent configuration.

You can test the parameter by running the Zabbix agent with `-t` flag. Run the test as the `zabbix` user, and specify the configuration file if it is not in the default location:

For Zabbix agent:

```
sudo -u zabbix zabbix_agentd -c /etc/zabbix/zabbix_agentd.conf -t mysql.questions
```

For Zabbix agent 2:

```
sudo -u zabbix zabbix_agent2 -c /etc/zabbix/zabbix_agent2.conf -t mysql.questions
```

The expected output contains a number, for example:

```
mysql.questions                                [t|4567]
```

Step 4: Reload and test the user parameter

Reload user parameters from the configuration file by running the following command:

For Zabbix agent:

```
sudo zabbix_agentd -c /etc/zabbix/zabbix_agentd.conf -R userparameter_reload
```

For Zabbix agent 2:

```
sudo zabbix_agent2 -c /etc/zabbix/zabbix_agent2.conf -R userparameter_reload
```

Run the command as `root` or as the operating system user running the agent.

Alternatively, restart the agent to apply the configuration changes.

You can test the parameter against the running agent by using the `zabbix_get` utility, for example, from the Zabbix server:

```
zabbix_get -s <IP address> -k mysql.questions
```

Here, `<IP address>` is the IP address of the host running the agent that `zabbix_get` connects to. The IP address of the host running `zabbix_get` must be listed in the agent's `Server` parameter.

For this test to work, the agent must accept passive checks (for the Zabbix agent, `StartAgents` must not be set to 0). If the agent is configured for active checks only, skip this test and verify the item in *Monitoring* > *Latest* data after completing the final step.

Step 5: Create an item

Create a new item on the monitored host with the `mysql.questions` item key.

Select the item type: either *Zabbix Agent* or *Zabbix Agent (active)*.

For Zabbix agent (passive), the Zabbix server or proxy address must be listed in the agent's `Server` parameter. For Zabbix agent (active), the `ServerActive` parameter must specify the Zabbix server or proxy, and the agent's `Hostname` parameter must match the host name configured in the frontend exactly.

Select the item's *Type of information*. The agent always returns the value as text; *Type of information* defines how Zabbix interprets and stores the value after any preprocessing steps.

For the raw cumulative counter, select *Numeric (unsigned)*. To measure the number of statements per second, select *Numeric (float)* and add the *Change per second* preprocessing step.

Windows performance counters

Overview

You can effectively monitor Windows performance counters using the `perf_counter[]` key.

For example:

```
perf_counter["\Processor(0)\Interrupts/sec"]
```

or

```
perf_counter["\Processor(0)\Interrupts/sec", 10]
```

For more information on using this key or its English-only equivalent `perf_counter_en`, see Windows-specific item keys.

In order to get a full list of performance counters available for monitoring, you may run:

```
typeperf -qx
```

You may also use low-level discovery to discover multiple object instances of Windows performance counters and automate the creation of `perf_counter` items for multiple instance objects.

Numeric representation

Windows maintains numeric representations (indexes) for object and performance counter names. Zabbix supports these numeric representations as parameters to the `perf_counter`, `perf_counter_en` item keys and in `PerfCounter`, `PerfCounterEn` configuration parameters.

However, it's not recommended to use them unless you can guarantee your numeric indexes map to correct strings on specific hosts. If you need to create portable items that work across different hosts with various localized Windows versions, you can

use the `perf_counter_en` key or `PerfCounterEn` configuration parameter which allow you to use English names regardless of system locale.

To find out the numeric equivalents, run `regedit`, then locate the *Counter* in `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib\009`.

The registry entry contains information like this:

```
1
1847
2
System
4
Memory
6
% Processor Time
10
File Read Operations/sec
12
File Write Operations/sec
14
File Control Operations/sec
16
File Read Bytes/sec
18
File Write Bytes/sec
....
```

Here you can find the corresponding numbers for each string part of the performance counter, like in `\System\% Processor Time`:

```
System → 2
% Processor Time → 6
```

Then you can use these numbers to represent the path in numbers:

```
\2\6
```

Performance counter parameters

You can deploy some PerfCounter parameters for the monitoring of Windows performance counters.

For example, you can add these to the Zabbix agent configuration file:

```
PerfCounter=UserPerfCounter1,"\Memory\Page Reads/sec",30
or
PerfCounter=UserPerfCounter2,"\4\24",30
```

With such parameters in place, you can then simply use `UserPerfCounter1` or `UserPerfCounter2` as the keys for creating the respective items.

Remember to restart Zabbix agent after making changes to the configuration file.

Mass update

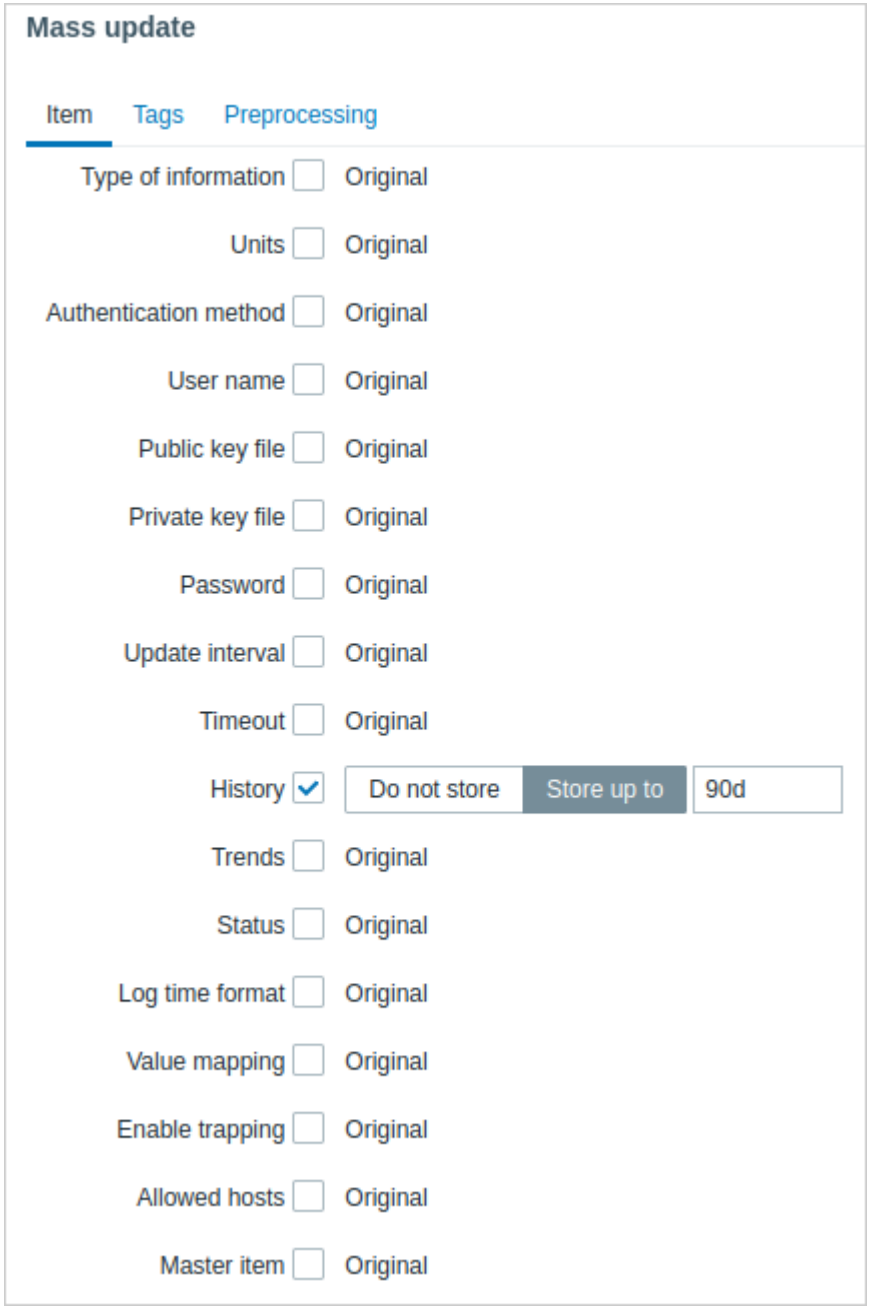
Overview

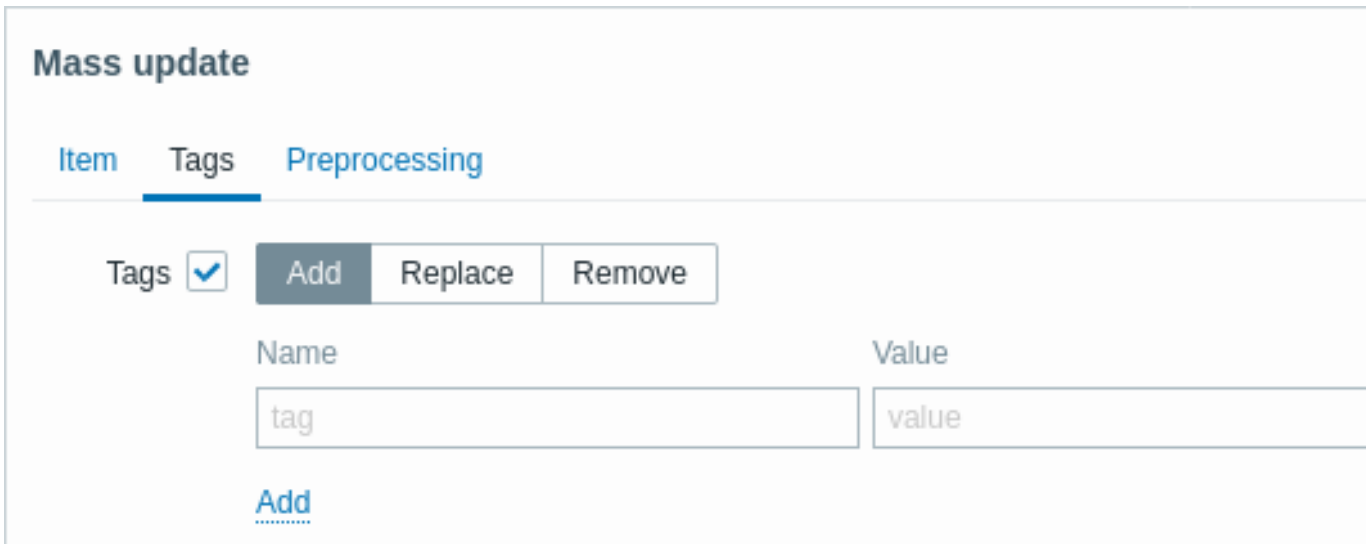
Sometimes you may want to change some attribute for a number of items at once. Instead of opening each individual item for editing, you may use the mass update function for that.

Using mass update

To mass-update some items, do the following:

- Mark the checkboxes of the items to update in the list
- Click *Mass update* below the list
- Navigate to the tab with required attributes (*Item*, *Tags* or *Preprocessing*)
- Mark the checkboxes of the attributes to update
- Enter new values for the attributes



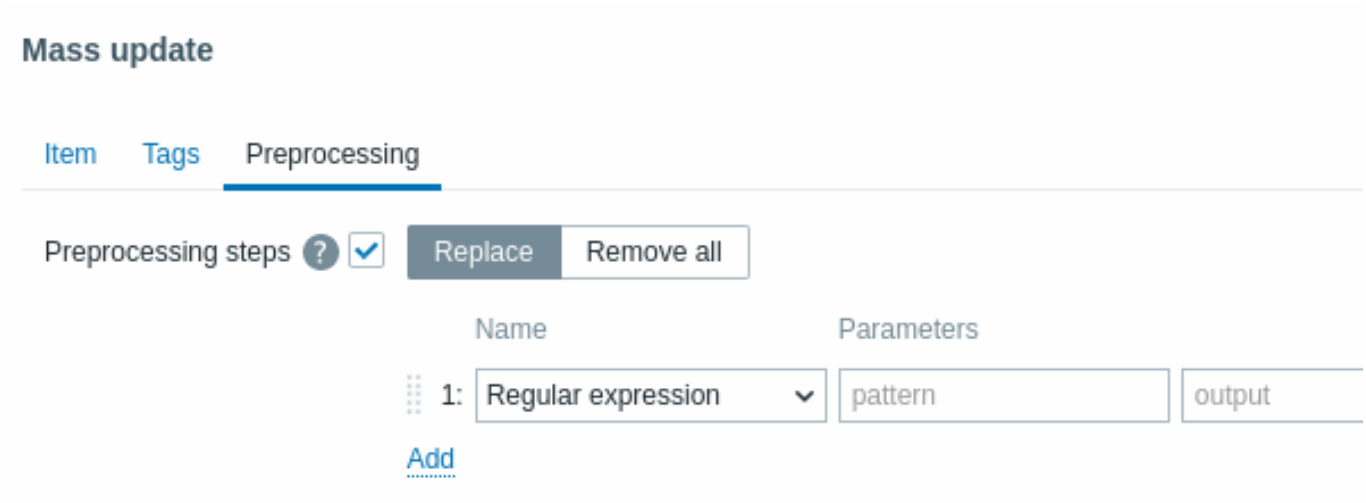


The *Tags* option allows you to:

- *Add* - add specified tags to the items (tags with the same name, but different values are not considered 'duplicates' and can be added to the same item).
- *Replace* - remove the specified tags and add tags with new values

- *Remove* - remove specified tags from the items

User macros, {INVENTORY.*} macros, {HOST.HOST}, {HOST.NAME}, {HOST.CONN}, {HOST.DNS}, {HOST.IP}, {HOST.PORT} and {HOST.ID} macros are supported in tags.



Mark the checkbox next to *Preprocessing steps* to activate mass update for preprocessing steps.

The *Preprocessing* mass update allows you to:

- *Replace* - replace the existing preprocessing steps on the items with the preprocessing steps specified below
- *Remove all* - remove all preprocessing steps from the items

When done, click *Update*.

Value mapping

Overview

Value mapping allows configuring a more user-friendly representation of received values using mappings between numeric/string values and string representations.

For example, when an item's value is "0" or "1," value mappings can be used to represent these values in a more user-friendly manner:

- 0 → Not Available
- 1 → Available

Value mappings for data backup types could be configured as follows:

- F → Full
- D → Differential
- I → Incremental

Value ranges for voltage could be configured as follows:

- <=209 → Low
- 210-230 → OK
- >=231 → High

Value mapping is used in Zabbix frontend and notifications sent by media types.

Attention:
Substitution of the received value with the configured representation is performed both in Zabbix frontend and server; however, the server handles substitution only in the following cases:

- when populating host inventory;
- when expanding supported macros {ITEM.VALUE}, {ITEM.LASTVALUE}, {EVENT.OPDATA}, and {EVENT.CAUSE.OPDATA}.

Value mappings are set up on templates or hosts. Once configured, they are available for all items within the respective template or host. When configuring items, specify the name of a previously configured value mapping in the *Value mapping* parameter.

Note:
There is no value map inheritance - hosts and templates do not inherit value mappings from linked templates. Template items on a host will continue to use the value mappings configured on the template.

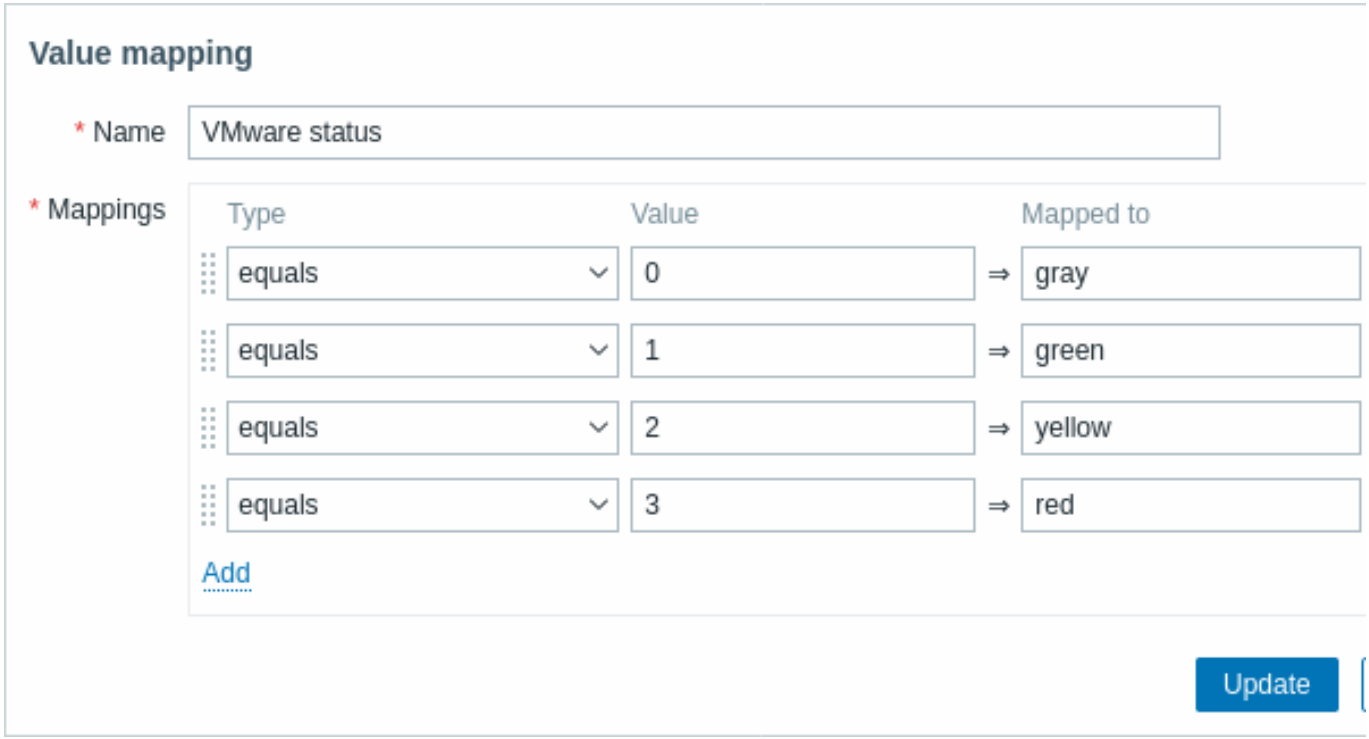
Note:
Value mappings can be used with items having *Numeric (unsigned)*, *Numeric (float)*, and *Character* types of information.

Value mappings are exported/imported with the respective template or host. They can also be mass-updated using the host and template mass update forms.

Configuration

To configure a value mapping, follow these steps:

1. Open the host or template configuration form.
2. In the *Value mapping* tab, click *Add* to add a new value mapping, or click the name of an existing mapping to edit it.



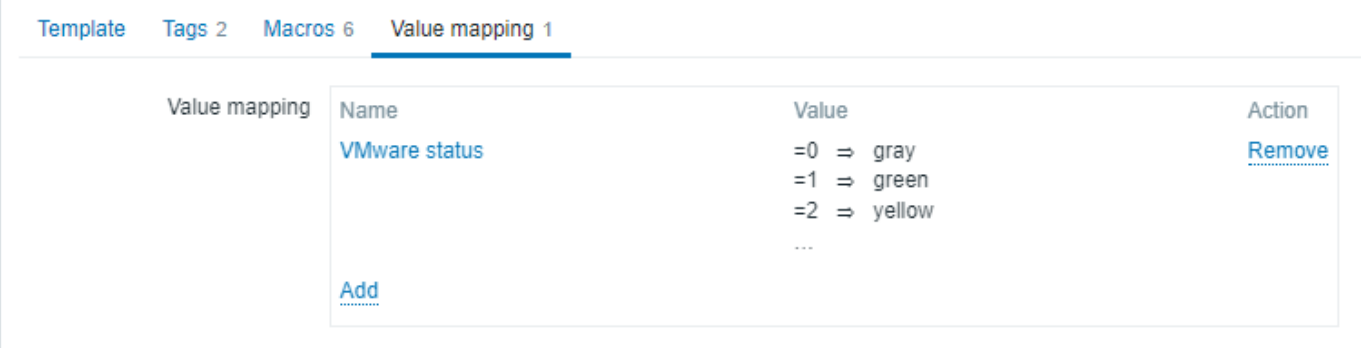
Parameters of a value mapping:

Parameter	Description
Name	Unique name for the set of value mappings.
Mappings	Individual rules for mapping numeric/string values to string representations. Mapping is applied in the order of the rules that can be reordered by dragging.
Type	Mapping type: **equals** - equal values will be mapped; **is greater than or equals** - equal or greater values will be mapped; **is less than or equals** - equal or smaller values will be mapped; **in range** - values in range will be mapped; the range is expressed as <number1>-<number2> or <number>; multiple ranges are supported (for example, 1-10,101-110,201); **regexp** - values corresponding to the regular expression will be mapped (global regular expressions are not supported); **default** - all outstanding values will be mapped, other than those with specific mappings. For mapping ranges, only numeric value types (*is greater than or equals*, *is less than or equals*, *in range*) are supported.
Value	Incoming value (may contain a range or regular expression, depending on the mapping type).

Parameter	Description
Mapped to	String representation (up to 64 characters) for the incoming value.

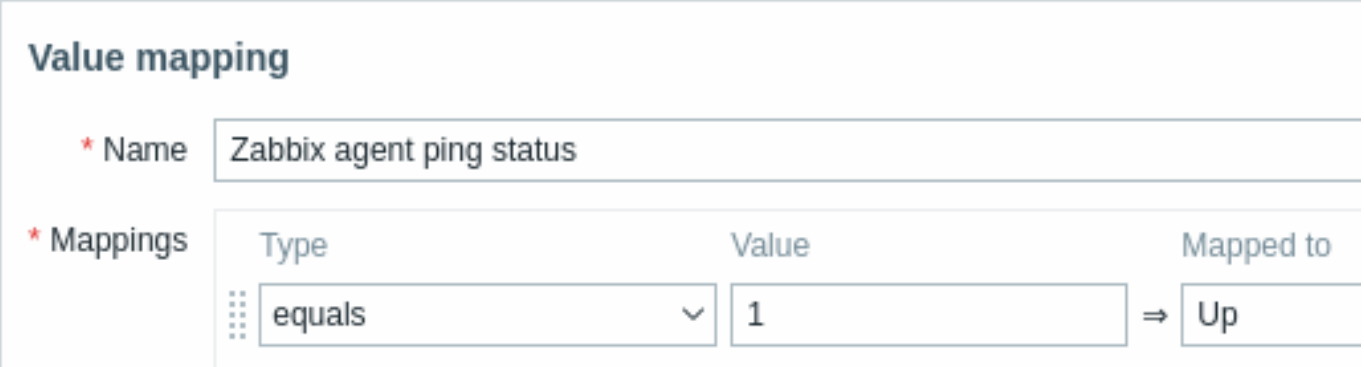
All mandatory input fields are marked with a red asterisk.

When viewing the value mapping in the list, only the first three mappings are visible, with three dots indicating that more mappings exist.



Value mapping example

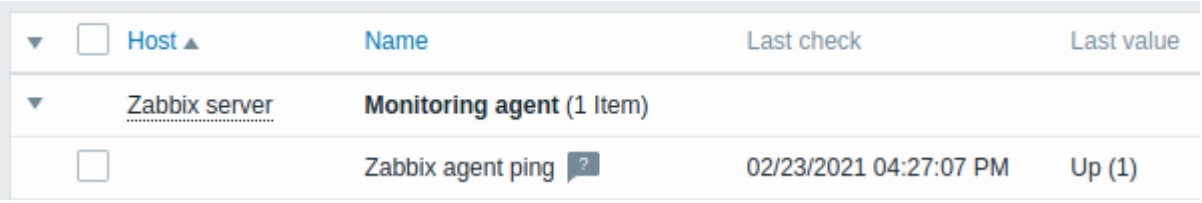
One of the predefined agent items *Zabbix agent ping* uses a template-level value mapping "Zabbix agent ping status" to display its values.



In the item configuration form, you can find a reference to this value mapping in the *Value mapping* field:

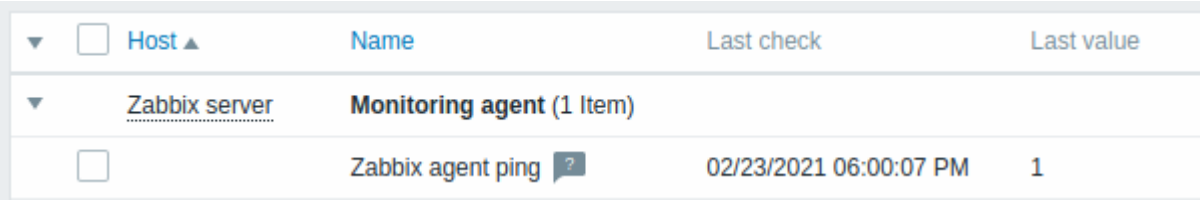


This mapping is used in the *Monitoring* > *Latest data* section to display "Up" (with the raw value in parentheses).



Note:
In the *Latest data* section, displayed values are shortened to 20 symbols. If value mapping is used, this shortening is not applied to the mapped value but only to the raw value (displayed in parentheses).

Without a predefined value mapping, you would only see "1", which might be challenging to understand.



Queue

Overview

The queue displays items that are waiting for a refresh. The queue is just a **logical** representation of data. There is no IPC queue or any other queue mechanism in Zabbix.

Items monitored by proxies are also included in the queue - they will be counted as queued for the proxy history data update period.

Only items with scheduled refresh times are displayed in the queue. This means that the following item types are excluded from the queue:

- log, logrt and event log active Zabbix agent items
- SNMP trap items
- trapper items
- web monitoring items
- dependent items

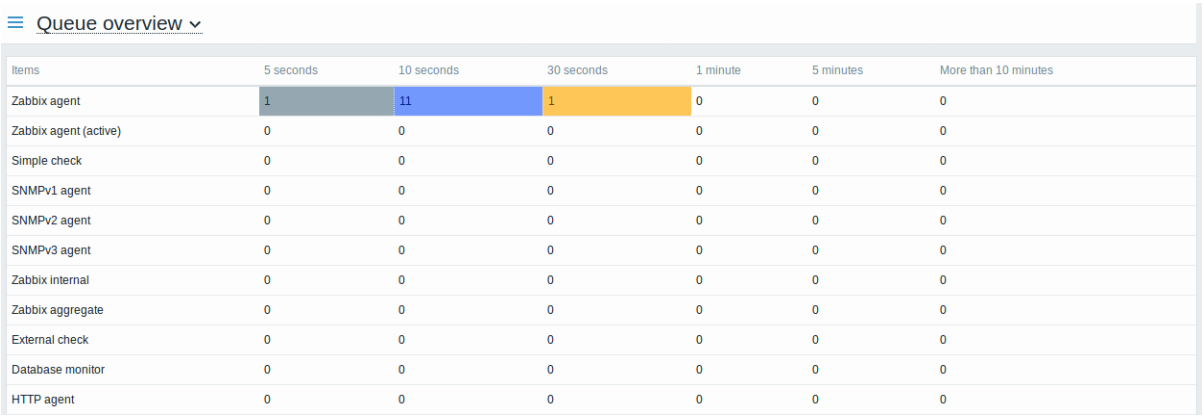
Statistics shown by the queue is a good indicator of the performance of Zabbix server.

The queue is retrieved directly from Zabbix server using JSON protocol. The information is available only if Zabbix server is running.

Attention:
Items are not counted in the queue if the item interface becomes unavailable due to connection problems or agent not working properly.

Reading the queue

To read the queue, go to *Administration* > *Queue* > *Queue overview*.

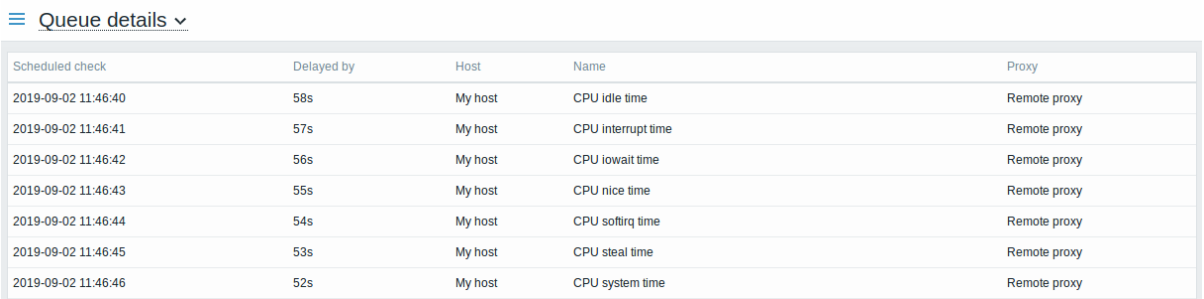

Queue overview

Items	5 seconds	10 seconds	30 seconds	1 minute	5 minutes	More than 10 minutes
Zabbix agent	1	11	1	0	0	0
Zabbix agent (active)	0	0	0	0	0	0
Simple check	0	0	0	0	0	0
SNMPv1 agent	0	0	0	0	0	0
SNMPv2 agent	0	0	0	0	0	0
SNMPv3 agent	0	0	0	0	0	0
Zabbix internal	0	0	0	0	0	0
Zabbix aggregate	0	0	0	0	0	0
External check	0	0	0	0	0	0
Database monitor	0	0	0	0	0	0
HTTP agent	0	0	0	0	0	0

The picture here is generally ”ok” so we may assume that the server is doing fine.

The queue shows some items waiting up to 30 seconds. It would be great to know what items these are.

To do just that, select *Queue details* in the title dropdown. Now you can see a list of those delayed items.

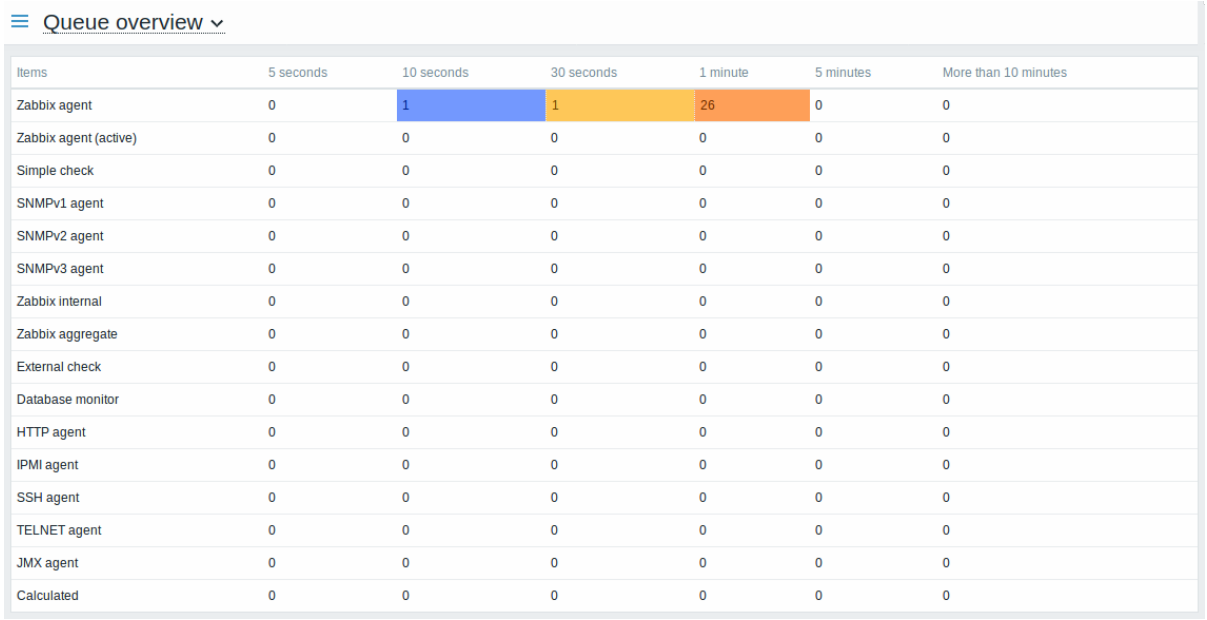

Queue details

Scheduled check	Delayed by	Host	Name	Proxy
2019-09-02 11:46:40	58s	My host	CPU idle time	Remote proxy
2019-09-02 11:46:41	57s	My host	CPU interrupt time	Remote proxy
2019-09-02 11:46:42	56s	My host	CPU iowait time	Remote proxy
2019-09-02 11:46:43	55s	My host	CPU nice time	Remote proxy
2019-09-02 11:46:44	54s	My host	CPU softirq time	Remote proxy
2019-09-02 11:46:45	53s	My host	CPU steal time	Remote proxy
2019-09-02 11:46:46	52s	My host	CPU system time	Remote proxy

With these details provided it may be possible to find out why these items might be delayed.

With one or two delayed items there perhaps is no cause for alarm. They might get updated in a second. However, if you see a bunch of items getting delayed for too long, there might be a more serious problem.

See also: Aligning time zones when using scheduling intervals.



Items	5 seconds	10 seconds	30 seconds	1 minute	5 minutes	More than 10 minutes
Zabbix agent	0	1	1	26	0	0
Zabbix agent (active)	0	0	0	0	0	0
Simple check	0	0	0	0	0	0
SNMPv1 agent	0	0	0	0	0	0
SNMPv2 agent	0	0	0	0	0	0
SNMPv3 agent	0	0	0	0	0	0
Zabbix internal	0	0	0	0	0	0
Zabbix aggregate	0	0	0	0	0	0
External check	0	0	0	0	0	0
Database monitor	0	0	0	0	0	0
HTTP agent	0	0	0	0	0	0
IPMI agent	0	0	0	0	0	0
SSH agent	0	0	0	0	0	0
TELNET agent	0	0	0	0	0	0
JMX agent	0	0	0	0	0	0
Calculated	0	0	0	0	0	0

Queue item

A special internal item `zabbix[queue,<from>,<to>]` can be used to monitor the health of the queue in Zabbix. It will return the number of items delayed by the set amount of time. For more information, see Internal items.

Value cache

Overview

To make the calculation of trigger expressions, calculated items and some macros much faster, a value cache option is supported by the Zabbix server.

This in-memory cache can be used for accessing historical data, instead of making direct SQL calls to the database. If historical values are not present in the cache, the missing values are requested from the database and the cache updated accordingly.

Item values remain in value cache either until:

- the item is deleted (cached values are deleted after the next configuration sync);
- the item value is outside the time or count range specified in the trigger/calculated item expression (cached value is removed when a new value is received);
- the time or count range specified in the trigger/calculated item expression is changed so that less data is required for calculation (unnecessary cached values are removed after 24 hours).

Note:
Value cache status can be observed by using the server runtime control option `diaginfo` (or `diaginfo=valuecache`) and inspecting the section for value cache diagnostic information. This can be useful for determining misconfigured triggers or calculated items.

To enable the value cache functionality, an optional `ValueCacheSize` parameter can be specified in the Zabbix server configuration.

Two internal items are supported for monitoring the value cache: `zabbix[vcache,buffer,<mode>]` and `zabbix[vcache,cache,<param`
See more details with internal items.

Execute now

Overview

The *Execute now* option allows executing passive checks immediately.

Item value collection in Zabbix follows configured update intervals. Some items (e.g., low-level discovery rules) have long intervals, and you may need to fetch a new value immediately—for example, to detect changes in discoverable resources.

The *Execute now* option is supported for the following item types:

- Zabbix agent (passive)

- Simple check
- SNMP agent (v1/v2/v3)
- Zabbix internal
- External check
- Database monitor
- HTTP agent
- IPMI agent
- SSH agent
- TELNET agent
- JMX agent
- Calculated
- Dependent item
- Script
- Browser

If the item is of type ”Dependent item”, then its master item must also be one of the above types.

Configuration

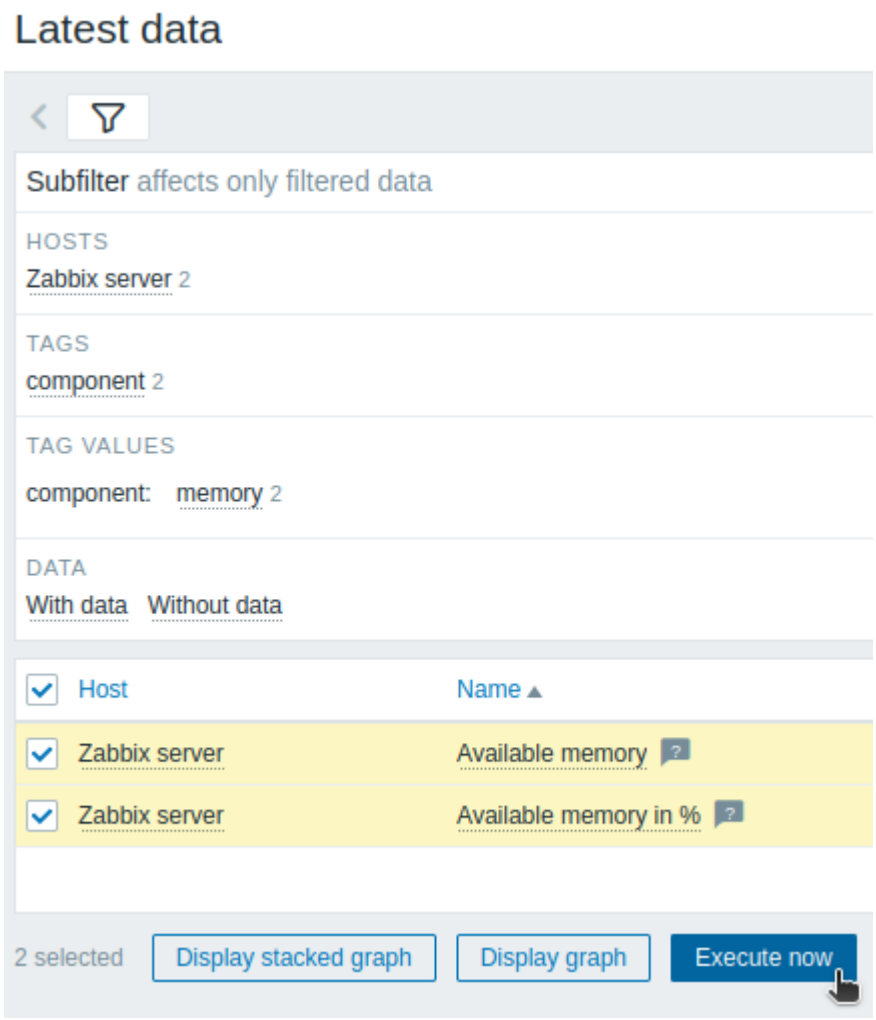
You can execute passive checks from:

- *Monitoring* > *Latest data*
- *Data collection* > *Hosts* > *Items* or *Discovery*
- The item menu
- When editing items or low-level discovery rules (in their configuration forms)

Attention:
The check must exist in the configuration cache in order to get executed; see also CacheUpdateFrequency. The cache is **not updated** before check execution, so recent changes to item or low-level discovery rule configuration will not be picked up. To test items or LLD rules that are being created or newly created, you can use the Test option instead.

For example, to execute a passive check in *Monitoring* > *Latest data*:

1. Select items; you can execute the check for multiple items at once.
2. Click *Execute now*.



Note:
In *Latest data*, users may *Execute now* items on hosts for which they have *read-write* permission, or *read* permission and *Invoke "Execute now" on read-only hosts* action enabled for their role.

Restricting agent checks

Overview

You can control which item keys Zabbix agent or agent 2 is allowed or denied to use when executing item checks, remote commands, or scripts.

To do that, use these agent configuration parameters to define allow/deny rules:

- `AllowKey=<pattern>`
- `DenyKey=<pattern>`
- `AllowKeyRegexp=<pattern>`
- `DenyKeyRegexp=<pattern>`

The <pattern> must contain a single item key.
 For `AllowKey` and `DenyKey`, the <pattern> supports wildcards (`*`). The wildcard matches any number of any characters in its position, and can be used to match item keys or parameters (e.g., `vfs.file.*[*]`).
 For `AllowKeyRegexp` and `DenyKeyRegexp`, the <pattern> supports regular expressions. The regular expression pattern can be used to match item keys and their parameter strings. See also examples of correct escaping.
 Please note that Zabbix agent and Zabbix agent 2 may support different regular expression syntax.

Attention:
To improve security, it is recommended to use exact item keys instead of wildcards for `AllowKey` and `DenyKey`. It is also recommended to use the most specific regular expression patterns for `AllowKeyRegexp` and `DenyKeyRegexp`. For details, see Securing allow/deny rules.

In contrast to other agent configuration parameters, you can specify an unlimited number of `AllowKey`, `DenyKey`, `AllowKeyRegexp`,

or DenyKeyRegexp parameters.

Important notes

- All system.run items are disabled by default (even when DenyKey and DenyKeyRegexp are empty), as if DenyKey=system.run[*] or DenyKeyRegexp=^system\.run\[.*\]$ was set as the last rule. Because of this, you can allow specific system.run items without explicitly denying other system.run items.
- If possible, use AllowKey/AllowKeyRegexp to allow only required items and deny everything else. Some keys can be abused to read unintended files via path traversal (e.g., vfs.file.contents["../../../../etc/passwd"]), and new Zabbix agent versions may introduce keys not covered by your DenyKey/DenyKeyRegexp rules.
- AllowKey, DenyKey, AllowKeyRegexp, and DenyKeyRegexp configuration does not affect HostnameItem, HostMetadataItem, or HostInterfaceItem agent parameters.
- Denied items become unsupported without any hints or error messages; for example:
  - Zabbix agent --print (-p) command-line parameter will not show denied item keys.
  - Zabbix agent --test (-t) command-line parameter will return "Unsupported item key." for denied item keys.
  - When logging is activated (LogRemoteCommands=1), Zabbix agent log file will not log denied remote commands.

Allow/deny rule order

You can specify an unlimited number of AllowKey, DenyKey, AllowKeyRegexp, or DenyKeyRegexp rules, though their order matters.

- Rules are evaluated one by one, from top to bottom.
- When an item key matches a rule, it is either allowed or denied, and rule evaluation stops.

For example, when evaluating vfs.file.contents[/etc/passwd], the rules are processed as follows:

```
AllowKey=vfs.file.contents[/tmp/app.log]    # Item key pattern does not match, agent proceeds to the next
AllowKey=vfs.file.contents[/etc/passwd]     # Item key pattern matches; agent allows the item check and st
DenyKey=vfs.file.*[*]                       # Agent ignores the rule, as the evaluation has stopped.
```

The following rule order will deny the item check:

```
DenyKey=vfs.file.*[*]                       # Item key pattern matches; agent denies the item check and st
AllowKey=vfs.file.contents[/etc/passwd]     # Agent ignores the rule, as the evaluation has stopped.
AllowKey=vfs.file.contents[/tmp/app.log]    # Agent ignores the rule, as the evaluation has stopped.
```

Examples

The following examples show common configuration patterns for AllowKey, DenyKey, AllowKeyRegexp, and DenyKeyRegexp.

Allowing specific checks and commands

Allow only two vfs.file item checks and two system.run commands:

using AllowKey and DenyKey:

```
AllowKey=vfs.file.contents[/tmp/app.log]
AllowKey=vfs.file.size[/tmp/app.log]
AllowKey=system.run[/usr/bin/uptime]
AllowKey=system.run[/usr/bin/df -h /]
DenyKey=vfs.file.*[*]
```

using AllowKeyRegexp and DenyKeyRegexp:

```
AllowKeyRegexp=^vfs\.file\.(contents|size)\[/tmp/app\.log\]$
AllowKeyRegexp=^system\.run\[/usr/bin/(uptime|df -h /)\]$
DenyKeyRegexp=^vfs\.file\..*\[.*\]$
```

Note:
Setting DenyKey=system.run[*] or DenyKeyRegexp=^system\.run\[.*\]$ is unnecessary, because all other system.run commands are denied by default.

Allowing scripts

Allow Zabbix agent to execute scripts on hosts via all available methods:

- Global scripts that can be executed in the frontend or via API (this method always uses the system.run[myscript.sh] key)

- Remote commands from action operations (this method always uses the system.run[myscript.sh,nowait] key)
- system.run Zabbix agent items with the script, for example:
  - system.run[myscript.sh]
  - system.run[myscript.sh,wait]
  - system.run[myscript.sh,nowait]

```
AllowKey=system.run[myscript.sh,*]
```

or

```
AllowKeyRegexp=^system\.run\[myscript\.sh,.*\]$
```

To control the wait/nowait parameter, you must set a different rule. For example, you can allow only system.run[myscript.sh,wait] items, thus excluding other methods:

```
AllowKey=system.run[myscript.sh,wait]
```

or

```
AllowKeyRegexp=^system\.run\[myscript\.sh,wait\]$
```

Securing allow/deny rules

These examples show how to secure overly permissive AllowKey/AllowKeyRegexp or DenyKey/DenyKeyRegexp rules.

Consider the following rules:

using AllowKey and DenyKey:

```
AllowKey=system.run["C:\Program^ Files\Zabbix^ Agent^ 2\scripts\test.bat*"]
DenyKey=vfs.file.*
DenyKey=system.cpu.load[*]
```

using AllowKeyRegexp and DenyKeyRegexp:

```
AllowKeyRegexp=^system\.run\["C:\\Program\sFiles\\Zabbix\sAgent\s2\\scripts\\test\.bat.*"\]$
DenyKeyRegexp=^vfs\.file\..*$
DenyKeyRegexp=^system\.cpu\.load\[.*\]$
```

Note:
On Windows, you must escape spaces in the path using a caret (^) for AllowKey or DenyKey and \s for AllowKeyRegexp or DenyKeyRegexp.

These rules contain a wildcard (*) (or .* for regular expression-based patterns), which can be misused:

- The test.bat script can be executed with any arguments, including unintended ones.
- The vfs.file.* (^vfs\.file\..*$) pattern matches item keys both with and without parameters; however, all vfs.file items require parameters.
- The system.cpu.load[*] (^system\.cpu\.load\[.*\]$) pattern matches only item keys with parameters; however system.cpu.load items do not require parameters.

To secure these rules, explicitly allow executing test.bat only with specific arguments, and deny correct item key patterns; for example:

using AllowKey and DenyKey:

```
AllowKey=system.run["C:\Program^ Files\Zabbix^ Agent^ 2\scripts\test.bat status"]
AllowKey=system.run["C:\Program^ Files\Zabbix^ Agent^ 2\scripts\test.bat version"]
DenyKey=vfs.file.*[*]
DenyKey=system.cpu.load
DenyKey=system.cpu.load[*]
```

using AllowKeyRegexp and DenyKeyRegexp:

```
AllowKeyRegexp=^system\.run\["C:\\Program\sFiles\\Zabbix\sAgent\s2\\scripts\\test\.bat (status|version)"\]
DenyKeyRegexp=^vfs\.file\..+\[.*\]$
DenyKeyRegexp=^system\.cpu\.load(\[.*\])?$
```

You can test the rules by running the following commands, which will return ZBX_NOTSUPPORTED.

```
cd "C:\Program Files\Zabbix Agent 2"
zabbix_agent2.exe -t system.run["C:\Program^ Files\Zabbix^ Agent^ 2\scripts\test.bat debug"]
zabbix_agent2.exe -t vfs.file.size["C:\ProgramData\MyApp\config.ini"]
zabbix_agent2.exe -t vfs.file.contents["C:\Windows\System32\drivers\etc\hosts"]
zabbix_agent2.exe -t system.cpu.load
zabbix_agent2.exe -t system.cpu.load[all,avg1]
```

Pattern examples

The following tables show how item key patterns are matched:

- A key matches the pattern only if it meets **all** conditions in the *Matches* column.
- Parameters must be fully enclosed in square brackets (e.g., `vfs.file.contents[*` and `vfs.file.contents*utf8]` are invalid patterns).
- The examples illustrate pattern matching only; actual items require parameters.

For `AllowKey` and `DenyKey`:

Pattern	Matches	Examples
`*`	Any key with or without parameters	

`vfs.file.*`	Key starts with `vfs.file.`
	With or without parameters
		Matches:
		`vfs.file.size`
		`vfs.file.contents`
		`vfs.file.contents[]`
		`vfs.file.size[/var/log/app.log]`

`vfs.*.contents`	Key starts with `vfs.`
	Key ends with `.contents`
	No parameters
		Matches:
		`vfs..contents`
		`vfs.mount.point.file.contents`
		
		Does not match:
		`vfs.contents`
		`vfs.file.contents[]`

`vfs.file.*[*]`	Key starts with `vfs.file.`
	Any or empty parameters
		Matches
		`vfs.file.get.custom[]`
		`vfs.file.size[/var/log/app.log, utf8]`
		
		Does not match:
		`vfs.file.get.custom`

`vfs.file.contents`	Key is `vfs.file.contents`
	No parameters
		Matches:
		`vfs.file.contents`
		
		Does not match:
		`vfs.file.contents[/etc/passwd]`

`vfs.file.contents[]`	Key is `vfs.file.contents[]`
	Empty parameters
		Matches:
		`vfs.file.contents[]`
		
		Does not match:
		`vfs.file.contents`

`vfs.file.contents[*]`	Key is `vfs.file.contents`
	Any or empty parameters
		Matches:
		`vfs.file.contents[/path/to/file]`
		
		Does not match:
		`vfs.file.contents`

`vfs.file.contents[/etc/passwd,*]`	Key is `vfs.file.contents`
	First parameter is `/etc/passwd`
	Any or empty second parameter
		Matches:
		`vfs.file.contents[/etc/passwd,]`
		`vfs.file.contents[/etc/passwd,utf8]`
		
		Does not match:
		`vfs.file.contents[]`
		`vfs.file.contents[/etc/passwd]`

Pattern	Matches	Examples
vfs.file.contents[*passwd*]	Key is vfs.file.contents Any parameters, at least one containing passwd	Matches: vfs.file.contents[/etc/passwd] vfs.file.contents[/etc/passwd,] vfs.file.contents[/etc/passwd,utf8]
vfs.file.contents[*passwd*,*]	Key is vfs.file.contents First parameter contains passwd Any or empty second parameter	Matches: vfs.file.contents[/etc/passwd,] vfs.file.contents[/etc/passwd,utf8] Does not match: vfs.file.contents[/etc/passwd] vfs.file.contents[/tmp/test]
vfs.file.contents[/etc/passwd,utf8]	Key is vfs.file.contents First parameter is /etc/passwd Second parameter is utf8	Matches: vfs.file.contents[/etc/passwd,utf8] Does not match: vfs.file.contents[/etc/passwd,] vfs.file.contents[/etc/passwd,utf16]

For AllowKeyRegexp and DenyKeyRegexp:

Pattern	Matches	Examples
^.*$	Any key with or without parameters	
^vfs\..*$	Key starts with vfs.file. With or without parameters	Matches: vfs.file.size vfs.file.contents vfs.file.contents[] vfs.file.size[/var/log/app.log]
^vfs\..*\.contents$	Key starts with vfs. Key ends with .contents No parameters	Matches: vfs..contents vfs.mount.point.file.contents Does not match: vfs.contents vfs.file.contents[]
^vfs\.file\..*\[.*\]$	Key starts with vfs.file. Any or empty parameters	Matches vfs.file.get.custom[] vfs.file.size[/var/log/app.log, utf8] Does not match: vfs.file.get.custom
^vfs\.file\.contents$	Key is vfs.file.contents No parameters	Matches: vfs.file.contents Does not match: vfs.file.contents[/etc/passwd]
^vfs\.file\.contents\[\]$	Key is vfs.file.contents[] Empty parameters	Matches: vfs.file.contents[] Does not match: vfs.file.contents
^vfs\.file\.contents\[.*\]$	Key is vfs.file.contents Any or empty parameters	Matches: vfs.file.contents[/path/to/file] Does not match: vfs.file.contents

Pattern	Matches	Examples
^vfs\.file\.contents\[/etc/passwd,.*\]$	Key is vfs.file.contents First parameter is /etc/passwd Any or empty second parameter	Matches: vfs.file.contents[/etc/passwd,] vfs.file.contents[/etc/passwd,utf8] Does not match: vfs.file.contents[] vfs.file.contents[/etc/passwd]
^vfs\.file\.contents\[.*passwd.*\]$	Key is vfs.file.contents Any parameters, at least one containing passwd	Matches: vfs.file.contents[/etc/passwd] vfs.file.contents[/etc/passwd,] vfs.file.contents[/etc/passwd,utf8]
^vfs\.file\.contents\[.*passwd.*,.*\]$	Key is vfs.file.contents First parameter contains passwd Any or empty second parameter	Matches: vfs.file.contents[/etc/passwd,] vfs.file.contents[/etc/passwd,utf8] Does not match: vfs.file.contents[/etc/passwd] vfs.file.contents[/tmp/test]
^vfs\.file\.contents\[/etc/passwd,(utf8|windows-1252)\]$	Key is vfs.file.contents First parameter is /etc/passwd Second parameter is utf8 or windows-1252	Matches: vfs.file.contents[/etc/passwd,utf8] vfs.file.contents[/etc/passwd,windows-1252] Does not match: vfs.file.contents[/etc/passwd,] vfs.file.contents[/etc/passwd,utf16]

Problem detection with triggers

Problem detection in Zabbix is based on **triggers**.

A trigger describes the problem condition, allowing Zabbix to react if the problem happens.

The full workflow of problem detection is:

1. Make sure you have an item collecting data.
2. Define the trigger for this item. Without a trigger, Zabbix just collects data, but never reacts on it.
3. Observe detected problems:
    - In *Monitoring* > *Problems*
    - In the *Problems* widget
4. Receive notifications, execute remote commands (if defined).

Note that official Zabbix templates contain pre-defined triggers. See guide for applying template to monitor Linux.

What is a "trigger"?

A trigger describes the *problem condition*, by means of the trigger expression.

The problem condition, when it occurs in real life, is when the trigger "fires" (a problem is created in Zabbix).

The problem condition can be based on the latest value, the average value, a discovered string, and many other functions. Triggers cannot be created for items returning values with the binary or JSON data type.

A simple trigger expression calculates the function (for example, `max()` for the maximum value) for the data specified in parentheses (typically data item and time window), and then compares this calculation to a threshold.

For example:

```
max(/host/vfs.fs.size[/,free],5m)<10G
```

This trigger fires if free disk space, measured over 5 minutes, has constantly been under 10 GB.

A trigger may have the following states:

Trigger state	Description	In Zabbix
OK	Trigger expression has evaluated to "0" (or FALSE)	No problem created Problem is resolved (if existed)
Problem	Trigger expression has evaluated to "1" (or TRUE)	Problem is created

Sometimes the trigger expression evaluation is not possible for some reason. See Unknown expression state.

Other resources:

- Trigger examples
- Configuring a trigger
- Best practices

Calculation frequency

A trigger is recalculated every time Zabbix server receives a **new value** that is part of the expression. When a new value is received, each function that is included in the expression is recalculated (not just the one that received the new value).

In addition, a trigger is recalculated every 30 seconds (by the history syncer) if the expression contains either:

- Date and time functions
- nodata() function

Triggers containing only trend functions are evaluated once per the smallest time period in the expression. While many trigger functions use item value history data, trend functions use trend data.

Trigger examples

This page is a collection of trigger examples.

Examples are sorted by use case:

- Unavailable host
- Unavailable proxy
- Unavailable SMTP cluster
- Unreachable host
- Unexpected reboot
- Changes in important files
- Changes in DNS query
- Non-matching software on different hosts
- Unsynchronized clocks
- Outdated agent
- Flapping interface
- High incoming traffic
- Low disk space
- Low disk space (dynamic threshold)
- High CPU load
- CPU load comparison
- Long-term CPU load comparison
- String containing special characters

Detect unavailable host

```
max(/host/zabbix[host,agent,available],5m)=0
```

This trigger fires if Zabbix agent on the host has been unavailable for 5 minutes.
 Function: `max`
 Data item: `zabbix[host,agent,available]`

Alternative:

```
nodata(/host/agent.ping,5m)=1
```

This trigger fires if no data has been received from Zabbix agent for 5 minutes.
 Function: `nodata`
 Data item: `agent.ping`

Detect unavailable proxy

```
fuzzytime(/host/zabbix[proxy,{$PROXY_NAME},lastaccess],1m)=0
```

This trigger fires if Zabbix proxy data lags behind Zabbix server time by 1 minute.
 Function: fuzzytime
 Data item: zabbix[proxy,{$PROXY_NAME},lastaccess]

Detect unavailable SMTP cluster

```
last(/smtp1.example.com/net.tcp.service[smtp])=0 and last(/smtp2.example.com/net.tcp.service[smtp])=0
```

This trigger fires if both SMTP servers are down.
 Function: last
 Data item: net.tcp.service

Detect unreachable host

```
count(/host/icmpping,30m,,"0")>5
```

This trigger fires if the host is unreachable by ping more than 5 times in the last 30 minutes.
 Function: count
 Data item: icmpping

Detect unexpected reboot

```
change(/host/system.uptime)<0
```

This trigger fires if a negative change in the system uptime value is encountered (indicating a reboot).
 Function: change
 Data item: system.uptime

Detect changes in important files

```
last(/host/vfs.file.cksum[/etc/passwd],#1)<>last(/host/vfs.file.cksum[/etc/passwd],#2)
```

This trigger fires if /etc/passwd has been changed. The expression is true when the previous /etc/passwd checksum differs from the most recent one. Similar expressions could be useful to monitor changes in important files, such as /etc/passwd, /etc/inetd.conf, /kernel, etc.
 Function: last
 Data item: vfs.file.cksum

Detect changes in DNS query

```
last(/Zabbix server/net.dns.record[192.0.2.1,{$WEBSITE_NAME},{$DNS_RESOURCE_RECORD_TYPE},2,1])<>"{$WEBSITE
```

Notice the quotes around the second operand.

This trigger fires if the query result is not equal to what it normally returns:

```
example.com             MX       0 mail.example.com
```

Function: last
 Data item: net.dns.record[192.0.2.1,{$WEBSITE_NAME},{$DNS_RESOURCE_RECORD_TYPE},2,1], with macros defined as:

```
{$WEBSITE_NAME} = example.com
{$DNS_RESOURCE_RECORD_TYPE} = MX
```

Detect non-matching software on different hosts

```
last(/host/vfs.file.contents[/etc/os-release])<>last(/host2/vfs.file.contents[/etc/os-release])
```

This trigger fires if the Ubuntu version is different on different hosts. Note how operands here are functions that return strings.
 Function: last
 Data item: vfs.file.contents

Detect unsynchronized clocks

```
fuzzytime(/host/system.localtime,10s)=0
```

The trigger fires if the client local time and Zabbix server time differs by more than 10 seconds.
 Function: fuzzytime
 Data item: system.localtime

Note that system.localtime must be configured as a passive check for Zabbix agent; on Zabbix agent 2 it may be configured as an active check.

Detect outdated agent

```
find(/host/agent.version,,"like","beta")=1
```

This trigger fires if Zabbix agent has beta version. Zabbix agent needs to be upgraded.
 Function: find
 Data item: agent.version

Detect flapping interface

```
changecount(/host/vfs.file.contents["/sys/class/net/eth0/operstate"],1h)>5
```

This trigger fires if the operational state (up/down/unknown) of eth0 has changed more than 5 times in an hour.
 Function: changecount
 Data item: vfs.file.contents

Detect high incoming traffic

```
min(/host/net.if.in[eth0,bytes],5m)>100K
```

This trigger fires if the number of received bytes on eth0 within the last five minutes was always over 100 kilobytes. Someone is probably downloading a large file.
 Function: min
 Data item: net.if.in[eth0,bytes]

Detect low disk space

```
max(/host/vfs.fs.size[/,free],5m)<10G
```

The trigger fires if free disk space is consistently (5 minutes) below 10 GB.

You can also define a recovery expression:

```
min(/host/vfs.fs.size[/,free],10m)>40G
```

The problem gets **resolved** only when free disk space is consistently (10 minutes) above 40 GB.
 Function: min
 Data item: vfs.fs.size

Detect low disk space (dynamic threshold)

```
last(/template/hrStorageFree[{#SNMPVALUE}])<last(/template/hrStorageSize[{#SNMPVALUE}])*0.1
```

The trigger fires if free storage (in allocation units) drops below 10 percent. Notice the value of another item being used to get an adaptive trigger threshold, applicable to discovered storage of various size. Function: last

Detect high CPU load

```
last(/host/system.cpu.load[all,avg1])>5
```

The trigger fires when the average processor load has been above 5 for one minute.

Variations:

```
min(/host/system.cpu.load[all,avg1],5m)>2 and time()<060000
```

```
min(/host/system.cpu.load[all,avg1],5m)>2 and not (dayofweek()=7 and time()>230000) and not (dayofweek()=1
```

Such triggers analyze 5 minutes of data and trigger only if CPU load is never under 2. Additionally, these triggers fire:

- only at night time (00:00-06:00)
- at any time except for 2 hours on week change (Sunday, 23:00 - Monday, 01:00)

```
(last(/host/system.cpu.load[all,avg1])>5) + (last(/host2/system.cpu.load[all,avg1])>5) + (last(/host3/syst
```

This trigger fires if the processor load is too high on at least two of the three hosts.
 Functions: last, min, dayofweek, time
 Data item: system.cpu.load

Compare CPU loads

```
avg(/Zabbix server/system.cpu.load,1h)/avg(/Zabbix server/system.cpu.load,1h:now-1d)>2
```

The trigger fires if the average load today tops the average load of the same hour yesterday (using time shift as now-1d) more than two times.
 Function: avg
 Data item: system.cpu.load

Compare long-term CPU loads

```
trendavg(/host/system.cpu.load,1M:now/M)>1.1*trendavg(/host/system.cpu.load,1M:now/M-1M)
```

This trigger fires if the CPU load on the host increased by more than 10% last month
 Function: trendavg
 Data item: system.cpu.load

You may also use the Event name field in trigger configuration to build a meaningful alert message, for example, to receive something like

```
"Load of Exchange server increased by 24% in July (0.69) comparing to June (0.56)"
```

the event name must be defined as:

```
Load of {HOST.HOST} server increased by {{?100*trendavg(//system.cpu.load,1M:now/M)/trendavg(//system.cpu.
```

It is also useful to allow manual closing in trigger configuration for this kind of problem.

Detect string containing special characters

```
last(/host/vfs.file.contents[/tmp/hello])={$HELLO_MACRO}
```

The trigger fires if `/tmp/hello` content is equal to the string defined in `{$HELLO_MACRO}`:

```
{$HELLO_MACRO} = \" //hello ?\"
```

Alternatively, you can compare to the string directly:

```
last(/Zabbix server/vfs.file.contents[/tmp/hello])="\\\" //hello ?\\\""
```

Notice how the special characters (\\ and ") characters are escaped when the string gets compared directly.

Function: `last`
 Data item: `vfs.file.contents`

Configuring a trigger

This page describes how to configure a trigger in Zabbix frontend.

Note that official Zabbix templates contain pre-defined triggers. See guide for applying template to monitor Linux. It is also possible to mass update configured triggers.

Other resources:

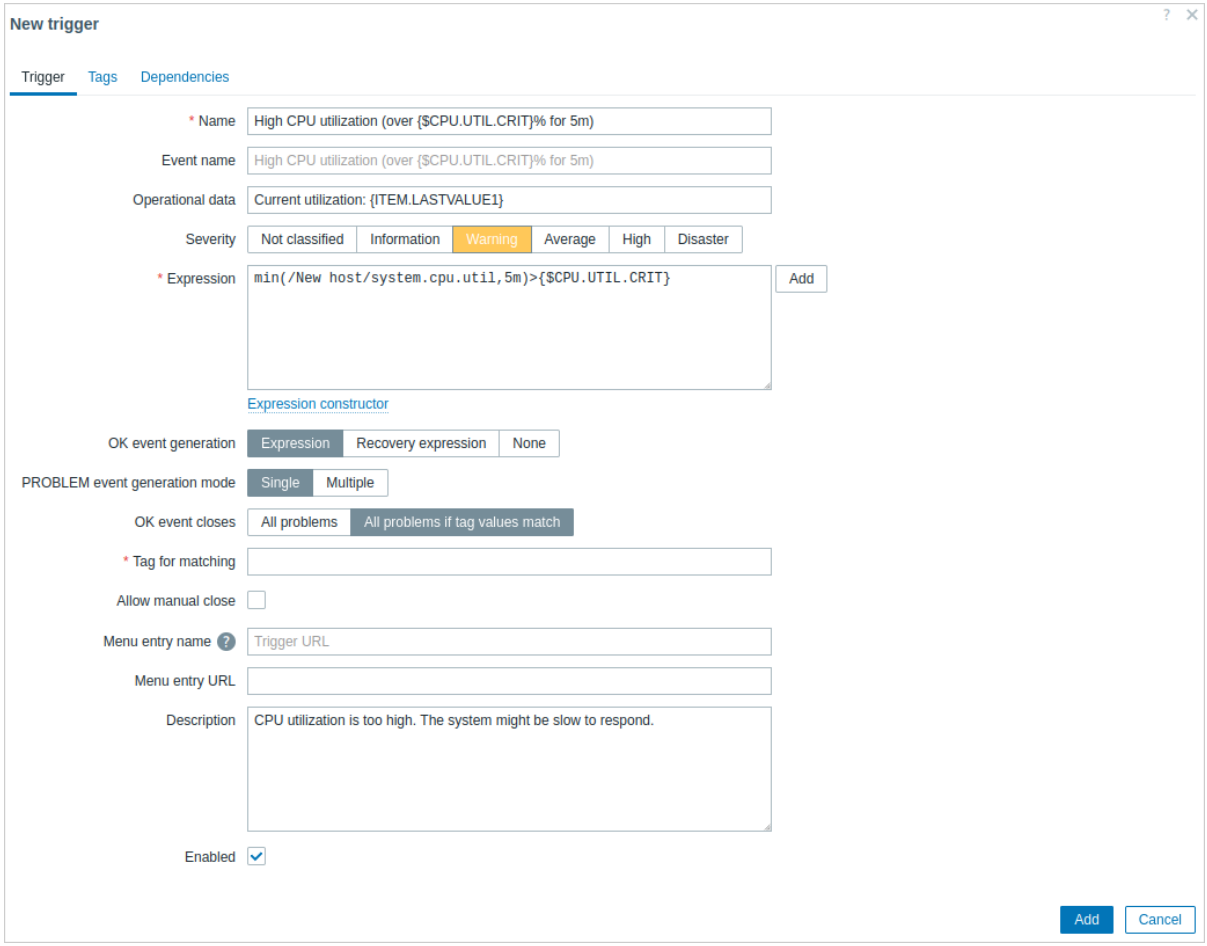
- Problem detection with triggers (general introduction)
- Trigger examples (by use case)
- Trigger expression and functions (syntax details)
- Best practices

Configuration

To configure a trigger, do the following:

1. Go to *Data collection* > *Hosts*.
2. Click *Triggers* in the row of the host.
3. Click *Create trigger* to the right (or the trigger name to edit an existing trigger).
4. Enter parameters of the trigger in the form.

The **Trigger** tab contains all the essential trigger attributes.

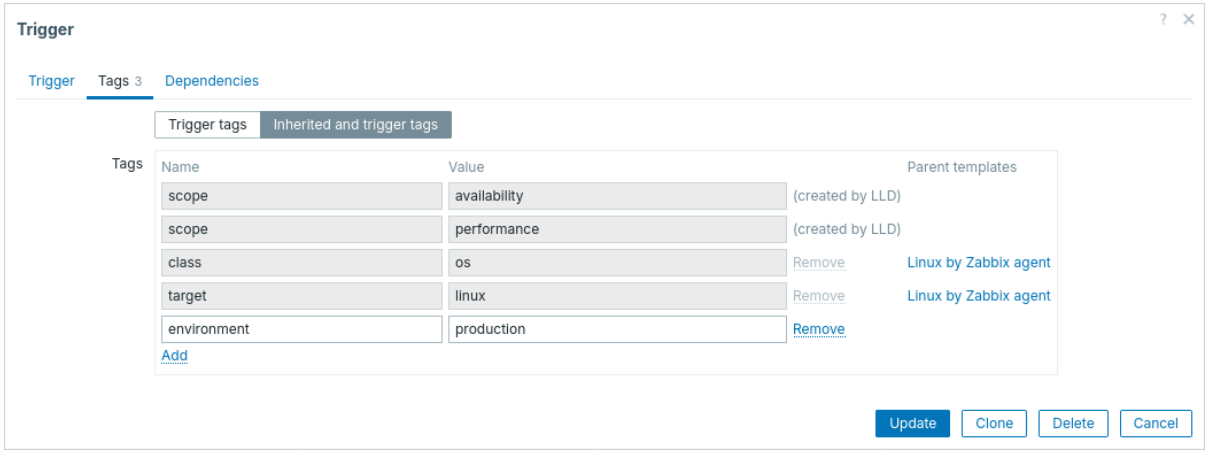


All mandatory input fields are marked with a red asterisk.

Parameter	Description
Name	Trigger name. Supported macros are: {HOST.HOST}, {HOST.NAME}, {HOST.PORT}, {HOST.CONN}, {HOST.DNS}, {HOST.IP}, {ITEM.VALUE}, {ITEM.VALUE.AGE}, {ITEM.VALUE.DATE}, {ITEM.VALUE.TIME}, {ITEM.VALUE.TIMESTAMP}, {ITEM.LASTVALUE}, {ITEM.LASTVALUE.AGE}, {ITEM.LASTVALUE.DATE}, {ITEM.LASTVALUE.TIME}, {ITEM.LASTVALUE.TIMESTAMP}, {ITEM.LOG.*}, and {$MACRO} user macros. **$1, $2...$9** macros can be used to refer to the first, second...ninth constant of the expression. *Note*: $1-$9 macros will resolve correctly if referring to constants in relatively simple, straightforward expressions. For example, the name "Processor load above $1 on {HOST.NAME}" will automatically change to "Processor load above 5 on New host" if the expression is last(/New host/system.cpu.load[percpu,avg1])>5.
Event name	When defined, this name will be used to create the problem event name. By default, the event name is the same as the trigger name. The event name may be used to build meaningful alerts containing problem data (see example). The same set of macros is supported as in the trigger name, plus {TIME}, {TIMESTAMP}, and {?EXPRESSION} expression macros.

Parameter	Description
Operational data	Enter some string (up to 2048 characters) with macros to display dynamic, real-time data in *Monitoring* > *Problems*. Alternatively, leave this field empty to display the latest values of all items from the trigger expression. The same set of macros is supported as in the trigger name, with the ability to resolve dynamically. For example: {ITEM.VALUE<1-9>} resolves to the item values at the moment the trigger state is changed (problem created, resolved, closed manually, or closed by correlation). {ITEM.LASTVALUE<1-9>} resolves to the latest item values. Note that closing a problem manually does not produce a new value, so both macros will still show the value from the problem time. Also note that both macros resolve to *UNKNOWN* if the latest value is older than *Max history display period* (see *Administration > General*). The jsonpath macro function is not supported in the frontend, so it resolves to *UNKNOWN* in *Monitoring* > *Problems*; it is resolved only by Zabbix server, for example, when {EVENT.OPDATA} is used in notifications.
Severity	Set the required trigger severity by clicking the buttons.
Expression	Logical expression used to define the conditions of a problem. Time suffixes and memory size suffixes are supported. A problem is created after all the conditions included in the expression are met, i.e. the expression evaluates to TRUE. The problem will be resolved as soon as the expression evaluates to FALSE, unless additional recovery conditions are specified in *Recovery expression*.
OK event generation	OK event generation options: **Expression** - OK events are generated based on the same expression as problem events; **Recovery expression** - OK events are generated if the problem expression evaluates to FALSE and the recovery expression evaluates to TRUE; **None** - in this case the trigger will never return to an OK state on its own.
Recovery expression	Logical expression (optional) defining additional conditions that have to be met before the problem is resolved, after the original problem expression has already been evaluated as FALSE. Recovery expression is useful for trigger hysteresis. It is **not** possible to resolve a problem by recovery expression alone if the problem expression is still TRUE. This field is only available if 'Recovery expression' is selected for *OK event generation*.
PROBLEM event generation mode	Mode for generating problem events: **Single** - a single event is generated when a trigger goes into the 'Problem' state for the first time; **Multiple** - an event is generated upon *every* 'Problem' evaluation of the trigger.
OK event closes	Select if OK event closes: **All problems** - all problems of this trigger; **All problems if tag values match** - only those trigger problems with matching event tag values.
Tag for matching	Enter event tag name to use for event correlation. This field is displayed if 'All problems if tag values match' is selected for the *OK event closes* property and is mandatory in this case.
Allow manual close	Check to allow manual closing of problem events generated by this trigger. Manual closing is possible when acknowledging problem events.
Menu entry name	If not empty, the name entered here (up to 64 characters) is used in several frontend locations as a label for the trigger URL specified in the *Menu entry URL* parameter. If empty, the default name *Trigger URL* is used. The same set of macros is supported as in the trigger name, plus {EVENT.ID}, {HOST.ID}, and {TRIGGER.ID}.
Menu entry URL	If not empty, the URL entered here (up to 2048 characters) is available as a link in the event menu in several frontend locations, for example, when clicking the problem name in *Monitoring* > *Problems* or *Problems* dashboard widget. The same set of macros is supported as in the trigger name, plus {EVENT.ID}, {HOST.ID}, and {TRIGGER.ID}. Note: user macros with secret values will not be resolved in the URL.
Description	Text field used to provide more information about this trigger. May contain instructions for fixing specific problem, contact detail of responsible staff, etc. The same set of macros is supported as in the trigger name.
Enabled	Unchecking this box will disable the trigger if required. Problems of a disabled trigger are no longer displayed in the frontend, but are not deleted.

The **Tags** tab allows you to define trigger-level tags. All problems of this trigger will be tagged with the values entered here.



In addition, the *Inherited and trigger tags* option allows you to view tags defined on the template level if the trigger comes from that template. If there are multiple templates with the same tag, these tags are displayed once and template names are separated by commas. A trigger does not "inherit" and display host-level tags.

Parameter	Description
Name/Value	Set custom tags to mark trigger events. Tags are a pair of tag name and value. You can use only the name or pair it with a value. A trigger may have several tags with the same name, but different values. User macros, user macros with context, low-level discovery macros, and macro functions with `{{ITEM.VALUE}}`, `{{ITEM.LASTVALUE}}` are supported in event tags. Low-level discovery macros can be used inside macro context. {TRIGGER.ID} macro is supported in trigger tag values. It may be useful for identifying triggers created from trigger prototypes and, for example, suppressing problems from these triggers during maintenance. If the total length of expanded value exceeds 255, it will be cut to 255 characters. See all macros supported for event tags. Event tags can be used for event correlation, in action conditions and will also be seen in *Monitoring* > *Problems* or the *Problems* widget.

The **Dependencies** tab contains all the dependencies of the trigger.

Click *Add* to add a new dependency.

Note:
You can also configure a trigger by opening an existing one, clicking the *Clone* button and then saving under a different name.

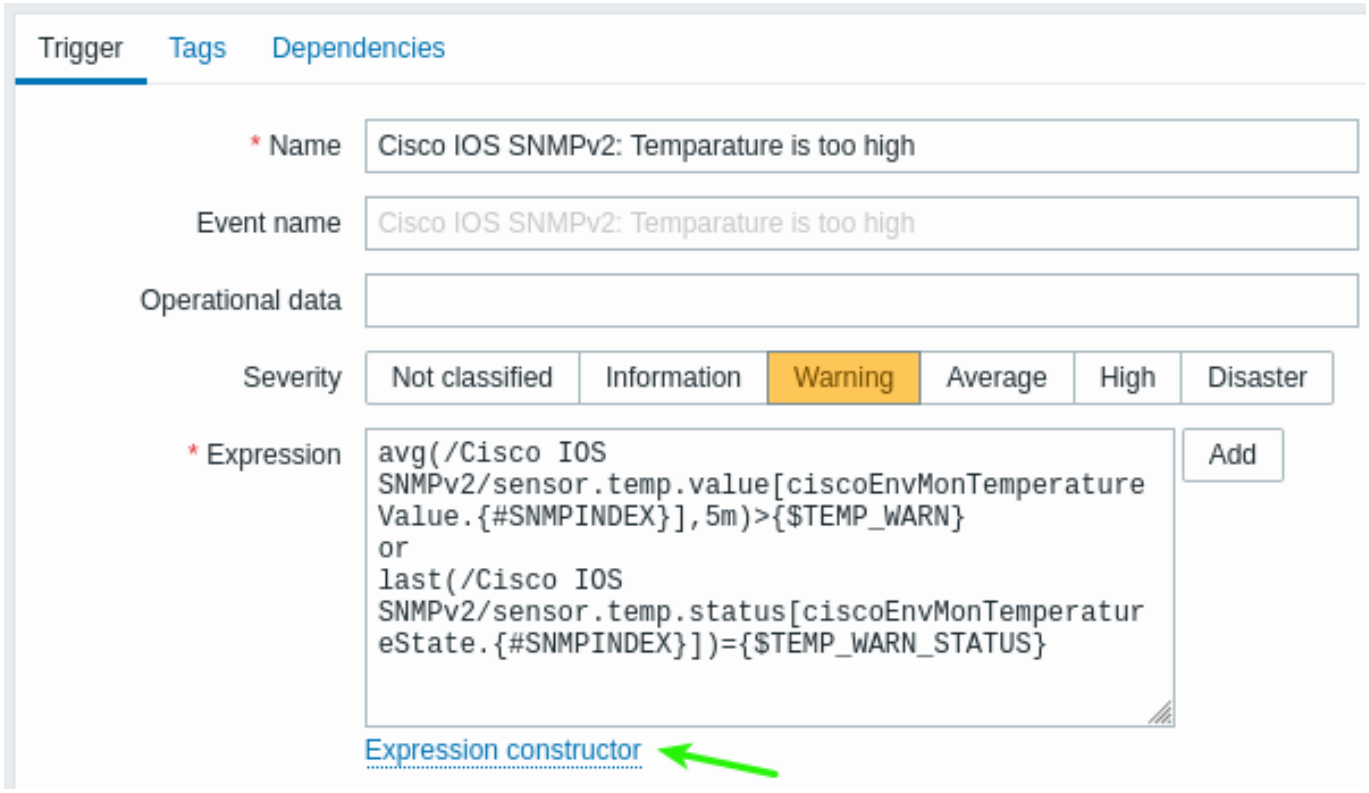
Trigger expression testing

It is possible to test what the expression result would be depending on the received value.

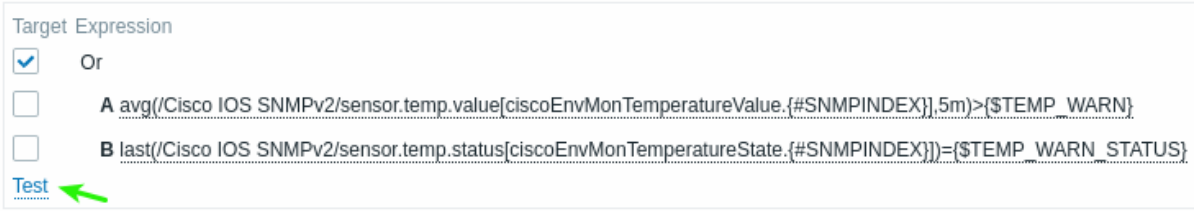
The following expression from an official template is taken as an example:

```
avg(/Cisco IOS SNMPv2/sensor.temp.value[ciscoEnvMonTemperatureValue.{#SNMPINDEX}],5m)>{$TEMP_WARN}
or
last(/Cisco IOS SNMPv2/sensor.temp.status[ciscoEnvMonTemperatureState.{#SNMPINDEX}])={$TEMP_WARN_STATUS}
```

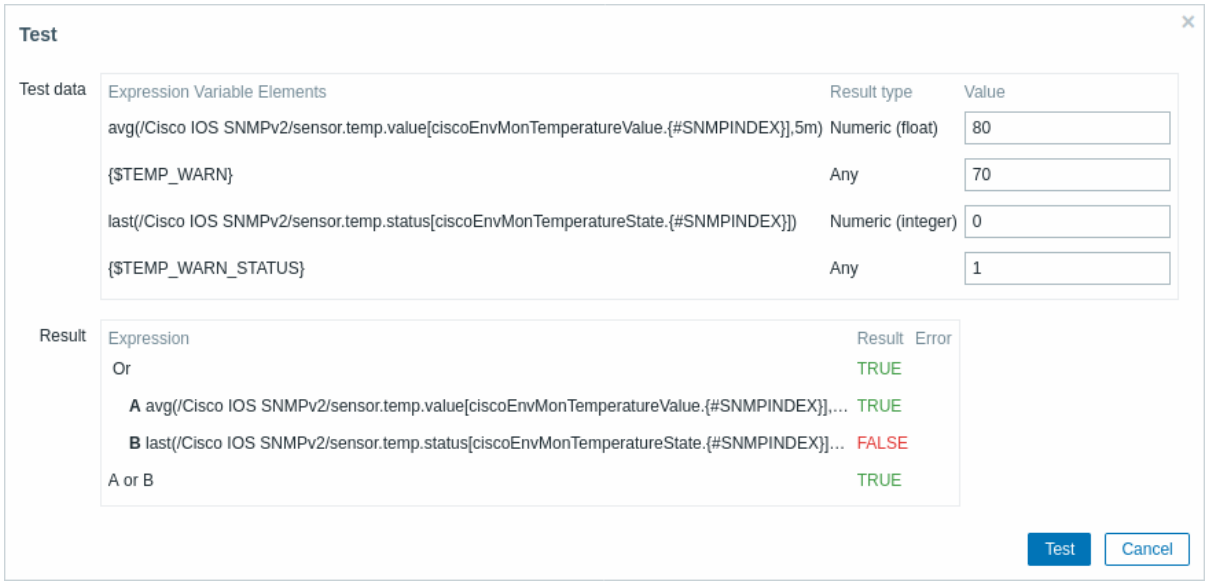
To test the expression, click *Expression constructor* under the expression field.



In the Expression constructor, all individual expressions are listed. To open the testing window, click *Test* below the expression list.



In the testing window you can enter sample values ('80', '70', '0', '1' in this example) and then see the expression result, by clicking the *Test* button.



The result of the individual expressions as well as the whole expression can be seen.

"TRUE" means that the specified expression is correct. In this particular case A, "80" is greater than the {$TEMP_WARN} specified value, "70" in this example. As expected, a "TRUE" result appears.

"FALSE" means that the specified expression is incorrect. In this particular case B, {$TEMP_WARN_STATUS} "1" needs to be equal with specified value, "0" in this example. As expected, a "FALSE" result appears.

The chosen expression type is "OR". If at least one of the specified conditions (A or B in this case) is TRUE, the overall result will

be TRUE as well. Meaning that the current value exceeds the warning value and a problem has occurred.

Trigger expression and functions

Overview

This page describes the trigger expression syntax and evaluation details.

The syntax of a simple expression is:

`function(/host/item,time_period)<operator><constant>`

In this expression, the first operand (to the left of the operator) is a function with its parameters in parentheses (typically the data item and time period).

The function is used to analyze received data during the set time period, resulting in a calculated value.

This value is then compared to the second operand, using the operator. In this example the second operand is a constant, but it can also be another function.

For example:

`min(/Zabbix server/net.if.in[eth0,bytes],5m)>100K`

This trigger fires if the number of received bytes on `eth0` during the last five minutes was always over 100 kilobytes. In this case the expression is true and a problem is created.

Trigger expressions are extremely flexible. In more complex expressions, a number of functions, operators and constants can be combined.

See also:

- Trigger examples
- Problem detection with triggers (general introduction)

Functions

Functions allow the analysis of collected values, for example, calculating the average or finding a specific string.

Click the respective function group to see more details.

Function group	Functions
Aggregate functions	avg, bucket_percentile, count, histogram_quantile, item_count, kurtosis, mad, max, min, skewness, stddevpop, stddevsamp, sum, sumofsquares, varpop, varsamp
Foreach functions	avg_foreach, bucket_rate_foreach, count_foreach, exists_foreach, last_foreach, max_foreach, min_foreach, sum_foreach
Bitwise functions	bitand, bitlshift, bitnot, bitor, bitrshift, bitxor
Date and time functions	date, dayofmonth, dayofweek, now, time
History functions	change, changecount, count, countunique, find, first, firstclock, fuzzytime, jsonpath, last, lastclock, logeventid, logseverity, logsource, logtimestamp, monodec, monoinc, nodata, percentile, rate, xmlxpath
Trend functions	baselinedev, baselinewma, trendavg, trendcount, trendmax, trendmin, trendstl, trendsum
Mathematical functions	abs, acos, asin, atan, atan2, avg, cbrt, ceil, cos, cosh, cot, degrees, e, exp, expm1, floor, log, log10, max, min, mod, pi, power, radians, rand, round, signum, sin, sinh, sqrt, sum, tan, truncate
Operator functions	between, in

Function group	Functions
Predictive functions	forecast, timeleft
String functions	ascii, bitlength, bytelength, char, concat, contains, insert, jsonpath, left, length, ltrim, mid, repeat, replace, right, rtrim, substring, trim, xmlxpath

Except where stated otherwise, these functions are supported in:

- Trigger expressions
- Calculated items
- Expression macros

Foreach functions are supported only for aggregate calculations.

Typically functions return numeric values for comparison. When returning strings, comparison is possible with the **=** and **<>** operators (see Detect non-matching software on different hosts example).

Function parameters

Function parameters allow you to specify:

- item key (as `/host/key`) for functions referencing a host item history
- time period (and other function-specific parameters)
- other expressions

Item key

The referenced item must be in a supported state (except for **nodata()** function, which is calculated for unsupported items as well).

Omission of the host name in the first parameter (i.e. as in `function(//key,parameter,...)`) is only supported in certain contexts:

- In the formula of calculated items
- In expression macros, which can be used in:
  - The *Event name* field
  - Graph name
  - The label of "Host" and "Trigger" map elements

In these contexts, you may also use the `{HOST.HOST}` macro. `{HOST.HOST<1-9>}` could be used in the case of the *Event name* field and the "Trigger" map element to refer to a specific item in the trigger expression. When the host name is omitted or replaced by `{HOST.HOST}` in these contexts, the reference points to the first item in the trigger expression or to the first item in the graph. Outside of these supported contexts, omitting the host name in trigger expressions will result in an error. See Compare long-term CPU loads example for an illustration of double-slash usage in Event name macros.

Time period

Function-specific parameters are placed after the item key and are separated from the item key by a comma.

Most of numeric functions accept **time period** as a parameter. It allows you to specify the interval we are interested in. It can be specified as time period (30s, 10m, 1h) or as a value range (#5 - for five latest values).

You may use seconds or time suffixes to indicate the time period. Preceded by a hash mark, the parameter has a different meaning:

Expression	Description
sum(/host/key,**10m)**	Sum of values in the last 10 minutes.
sum(/host/key,**#10)**	Sum of the last ten values.

Parameters with a hash mark have a different meaning with the function **last** - they denote the Nth previous value, so given the values 30, 70, 20, 60, 50 (from the most recent to the least recent):

- `last(/host/key,#2)` would return '70'
- `last(/host/key,#5)` would return '50'

The time period is measured up to "now" - where "now" is the latest recalculation time of the trigger (see Calculation frequency); "now" is not the "now" time of the server.

The time period specifies either:

- All values between "now - time period" and "now" (or, with time shift applied, between "now - time shift - time period" and "now - time shift").
- The specified number of values from the past, up to "now".

If no data exists for the specified evaluation period, the trigger or calculated item using this function becomes unsupported.

Note that:

- If only a single function (referencing data history) is used in the trigger, "now" is always the latest received value. For example, if the last value was received an hour ago, the time period will be regarded as up to the latest value an hour ago.
- A new trigger is calculated as soon as the first value is received (history functions); it will be calculated within 30 seconds for date and time and nodata() functions. Thus the trigger will be calculated even though perhaps the set time period (for example, one hour) has not yet passed since the trigger was created. The trigger will also be calculated after the first value, even though the time range was set, for example, to ten latest values.

Time shift

An optional time shift is supported with time or value count as the function parameter. This parameter allows you to reference data from a period of time in the past.

Time shift starts with `now` - specifying the current time, and is followed by `+N<time unit>` or `-N<time unit>` - to add or subtract N time units.

For example, `avg(/host/key,1h:now-1d)` will return the average value for an hour one day ago.

> **Attention:**
> Time shift specified in months (M) and years (y) is only supported for trend functions. Other functions support seconds (s), minutes (m), hours (h), days (d), and weeks (w).

Time shift with absolute time periods

Absolute time periods are supported in the time shift parameter, for example, midnight to midnight for a day, Monday-Sunday for a week, first day-last day of the month for a month.

Time shift for absolute time periods starts with `now` - specifying the current time, and is followed by any number of time operations: `/<time unit>` - defines the beginning and end of the time unit, for example, midnight to midnight for a day, `+N<time unit>` or `-N<time unit>` - to add or subtract N time units.

Please note that the value of time shift can be greater or equal to 0, while the time period minimum value is 1.

Parameter	Description
1d:now/d	Yesterday
1d:now/d+1d	Today
2d:now/d+1d	Last 2 days
1w:now/w	Last week
1w:now/w+1w	This week

Other expressions

Function parameters may contain other expressions, as in the following syntax:

```
min(min(/host/key,1h),min(/host2/key2,1h)*10)
```

Note that other expressions may not be used, if the function references item history. For example, the following syntax is **not** allowed:

~~`min(/host/key,#5*10)`~~

While other trigger expressions as function parameters are limited to non-history functions in triggers, this limitation does not apply in calculated items.

Operators

The following operators are supported for triggers **(in descending priority of execution)**:

Priority	Operator	Definition	**Notes for unknown values**	Force cast operand to float [1]
1	**-**	Unary minus	**-**Unknown → Unknown	Yes
2	**not**	Logical NOT	**not** Unknown → Unknown	Yes
3	*	Multiplication	0 * Unknown → Unknown (yes, Unknown, not 0 - to not lose Unknown in arithmetic operations) 1.2 * Unknown → Unknown	Yes
	/	Division	Unknown / 0 → error Unknown / 1.2 → Unknown 0.0 / Unknown → Unknown	Yes
4	**+**	Arithmetical plus	1.2 **+** Unknown → Unknown	Yes
	-	Arithmetical minus	1.2 **-** Unknown → Unknown	Yes
5	**<**	Less than. The operator is defined as: A<B ⇔ (A<B-0.000001)	1.2 **<** Unknown → Unknown	Yes
	<=	Less than or equal to. The operator is defined as: A<=B ⇔ (A≤B+0.000001)	Unknown **<=** Unknown → Unknown	Yes
	>	More than. The operator is defined as: A>B ⇔ (A>B+0.000001)		Yes
	>=	More than or equal to. The operator is defined as: A>=B ⇔ (A≥B-0.000001)		Yes
6	**=**	Is equal. The operator is defined as: A=B ⇔ (A≥B-0.000001) and (A≤B+0.000001)		No [1]
	<>	Not equal. The operator is defined as: A<>B ⇔ (A<B-0.000001) or (A>B+0.000001)		No [1]
7	**and**	Logical AND	0 **and** Unknown → 0 1 **and** Unknown → Unknown Unknown **and** Unknown → Unknown	Yes

Priority	Operator	Definition	Notes for unknown values	Force cast operand to float [1]
8	**or**	Logical OR	1 **or** Unknown → 1 0 **or** Unknown → Unknown Unknown **or** Unknown → Unknown	Yes

[1] String operand is still cast to numeric if:

- another operand is numeric
- operator other than **=** or **<>** is used on an operand

(If the cast fails - numeric operand is cast to a string operand and both operands get compared as strings.)

not, **and** and **or** operators are case-sensitive and must be in lowercase. They also must be surrounded by spaces or parentheses.

All operators, except unary **-** and **not**, have left-to-right associativity. Unary **-** and **not** are non-associative (meaning **-(-1)** and **not (not 1)** should be used instead of **--1** and **not not 1**).

Evaluation result:

- **<**, **<=**, **>**, **>=**, **=**, **<>** operators shall yield '1' in the trigger expression if the specified relation is true and '0' if it is false. If at least one operand is Unknown the result is Unknown;
- **and** for known operands shall yield '1' if both of its operands compare unequal to '0'; otherwise, it yields '0'; for unknown operands **and** yields '0' only if one operand compares equal to '0'; otherwise, it yields 'Unknown';
- **or** for known operands shall yield '1' if either of its operands compare unequal to '0'; otherwise, it yields '0'; for unknown operands **or** yields '1' only if one operand compares unequal to '0'; otherwise, it yields 'Unknown';
- The result of the logical negation operator **not** for a known operand is '0' if the value of its operand compares unequal to '0'; '1' if the value of its operand compares equal to '0'. For unknown operand **not** yields 'Unknown'.

Unknown expression state

It is possible that an unknown operand appears in a trigger expression when:

- unsupported item is used
- function evaluation, for a supported item, results in error

In this case the trigger expression generally evaluates to `Unknown` (as it cannot be evaluated)

It is possible to get notified on unknown triggers.

Exceptions

Despite an unknown operand, trigger expressions can be evaluated to known result (Problem/OK) in some cases:

- The `nodata()` function is evaluated regardless of whether the referenced item is supported or not.
- Expressions with AND/OR can be evaluated to known result in two cases:
  - **Case 1**: "`1 or some_function(unsupported_item1) or some_function(unsupported_item2) or ...`" evaluates to known result ('1' or "Problem"),
  - **Case 2**: "`0 and some_function(unsupported_item1) and some_function(unsupported_item2) and ...`" evaluates to known result ('0' or "OK").
- If the function evaluation for a supported item results in error, the function value becomes `Unknown` and it takes part as unknown operand in further expression evaluation.

Unknown operands may "disappear" only in logical expressions as described above. In arithmetic expressions unknown operands always lead to `Unknown` (except division by 0).

Attention:
Unknown expression state does not change the trigger state (Problem/OK). So, if trigger state was "Problem" (see Case 1), it stays in the problem state even if the known part is resolved ('1' becomes '0'), because the expression is now evaluated to `Unknown` and that does not change the trigger state.

If a trigger expression with several unsupported items evaluates to `Unknown`, the error message in the frontend refers to the last unsupported item evaluated.

Value caching

Values required for trigger evaluation are cached by Zabbix server. Because of this trigger evaluation causes a higher database load for some time after the server restarts.

The value cache is not cleared when item history values are removed (either manually or by housekeeper), so the server will use the cached values until they are older than the time periods defined in trigger functions or server is restarted.

> **Note:**
> If there is no recent data in the cache and there is no defined querying period in the function, Zabbix will by default go as far back in the past as one week to query the database for historical values.

Aggregate functions

Except where stated otherwise, all functions listed here are supported in:

- Trigger expressions
- Calculated item formulas
- Expression macros

Aggregate functions can work with either:

- history of items, for example, `min(/host/key,1h)`
- foreach functions as the only parameter, for example, `min(last_foreach(/*/key))` (only in calculated item formulas; to use in triggers reference the calculated item by its key)

The functions are listed without additional information. Click the function to see the full details.

Function	Description
avg	The average value of an item within the defined evaluation period.
bucket_percentile	Calculates the percentile from the buckets of a histogram.
count	The count of values in an array returned by a foreach function.
histogram_quantile	Calculates the φ-quantile from the buckets of a histogram.
item_count	The count of existing items in configuration that match the filter criteria.
kurtosis	The "tailedness" of the probability distribution in collected values within the defined evaluation period.
mad	The median absolute deviation in collected values within the defined evaluation period.
max	The highest value of an item within the defined evaluation period.
min	The lowest value of an item within the defined evaluation period.
skewness	The asymmetry of the probability distribution in collected values within the defined evaluation period.
stddevpop	The population standard deviation in collected values within the defined evaluation period.
stddevsamp	The sample standard deviation in collected values within the defined evaluation period.
sum	The sum of collected values within the defined evaluation period.
sumofsquares	The sum of squares in collected values within the defined evaluation period.
varpop	The population variance of collected values within the defined evaluation period.
varsamp	The sample variance of collected values within the defined evaluation period.

Common parameters

- `/host/key` is a common mandatory first parameter for the functions referencing the host item history.
- `(sec|#num)<:time shift>` is a common second parameter for the functions referencing the host item history, where:
  - **sec** - maximum evaluation period in seconds (time suffixes can be used).
  - **#num** - maximum evaluation range in latest collected values (if preceded by a hash mark).
  - **time shift** (optional) allows the evaluation point to be moved back in time. See more details on specifying time shift.

Function details

Some general notes on function parameters:

- Function parameters are separated by a comma
- Optional function parameters (or parameter parts) are indicated by < >
- Function-specific parameters are described with each function
- `/host/key` and `(sec|#num)<:time shift>` parameters must never be quoted

avg(/host/key,(sec|#num)<:time shift>)

The average value of an item within the defined evaluation period.
 Supported value types: *Float*, *Integer*.
 Supported foreach functions: *avg_foreach*, *count_foreach*, *exists_foreach*, *last_foreach*, *max_foreach*, *min_foreach*, *sum_foreach*.

Parameters: see common parameters.

Time shift is useful when there is a need to compare the current average value with the average value some time ago.

Examples:

```
avg(/host/key,1h) #calculate the average value for the last hour until now
avg(/host/key,#5) #calculate the average value of the five latest values
avg(/host/key,1h:now-1d) #calculate the average value for an hour from 25 hours ago to 24 hours ago from n
avg(last_foreach(/*/system.cpu.load[,avg1]?[group="MySQL Servers"])) #calculate the average processor load
avg(/host/proc.num,5m)>300 #trigger if the average number of processes in the last 5 minutes has been abov
```

bucket_percentile(item filter,time period,percentage)

Calculates the percentile from the buckets of a histogram.

Parameters:

- **item filter** - see item filter;
- **time period** - see time period;
- **percentage** - percentage (0-100).

Comments:

- Supported only in calculated items;
- This function is an alias for `histogram_quantile(percentage/100, bucket_rate_foreach(item filter, time period, 1))`.

Examples:

```
bucket_percentile(/*/http_request_duration_seconds_bucket[*],5m,95) #calculate the 95th percentile respons
bucket_percentile(/*/apiserver_request_duration_seconds_bucket[*],10m,99) #calculate the 99th percentile A
```

count(func_foreach(item filter,<time period>),<operator>,<pattern>)

The count of values in an array returned by a foreach function.
 Supported foreach functions: *avg_foreach*, *count_foreach*, *exists_foreach*, *last_foreach*, *max_foreach*, *min_foreach*, *sum_foreach*.

Parameters:

- **func_foreach** - foreach function for which the number of returned values should be counted. See foreach functions for details. Note that count_foreach and bucket_rate_foreach support additional parameters.
- **item filter** - see item filter;
- **time period** - see time period;
- **operator** (must be double-quoted). Supported `operators`:
eq - equal
ne - not equal
gt - greater
ge - greater or equal
lt - less
le - less or equal
like - matches if contains pattern (case-sensitive)
bitand - bitwise AND
regexp - case-sensitive match of the regular expression given in `pattern`
iregexp - case-insensitive match of the regular expression given in `pattern`
- **pattern** - the required pattern (string arguments must be double-quoted); supported if *operator* is specified in the third parameter.

Comments:

- Using **count()** with a history-related foreach function (max_foreach, avg_foreach, etc.) may lead to performance implications, whereas using **exists_foreach()**, which works only with configuration data, will not have such effect.
- Optional parameters *operator* or *pattern* can't be left empty after a comma, only fully omitted.
- With *bitand* as the third parameter, the fourth `pattern` parameter can be specified as two numbers, separated by '/': **number_to_compare_with/mask**. count() calculates "bitwise AND" from the value and the *mask* and compares the result to *number_to_compare_with*. If the result of "bitwise AND" is equal to *number_to_compare_with*, the value is counted.
If *number_to_compare_with* and *mask* are equal, only the *mask* need be specified (without '/').
- With *regexp* or *iregexp* as the third parameter, the fourth `pattern` parameter can be an ordinary or global (starting with '@') regular expression. In case of global regular expressions case sensitivity is inherited from global regular expression settings. For the purpose of regexp matching, float values will always be represented with 4 decimal digits after '.'. Also note that for large numbers difference in decimal (stored in database) and binary (used by Zabbix server) representation may affect the 4th decimal digit.

Examples:

```
count(max_foreach(/*/net.if.in[*],1h)) #calculate the number of net.if.in items that received data in the
count(last_foreach(/*/vfs.fs.size[*,pused]),"gt",95) #calculate the number of file systems with over 95% c
```

histogram_quantile(quantile,bucket1,value1,bucket2,value2,...)

Calculates the φ-quantile from the buckets of a histogram.
 Supported foreach function: *bucket_rate_foreach*.

Parameters:

- **quantile** - 0 ≤ φ ≤ 1;
- **bucketN, valueN** - manually entered pairs (>=2) of parameters or the response of bucket_rate_foreach.

Comments:

- Supported only in calculated items;
- Functionally corresponds to 'histogram_quantile' of PromQL;
- Returns -1 if values of the last 'Infinity' bucket ("+*inf*") are equal to 0.

Examples:

```
histogram_quantile(0.75,1.0,last(/host/rate_bucket[1.0]),"+Inf",last(/host/rate_bucket[Inf]))
histogram_quantile(0.5,bucket_rate_foreach(//item_key,30s))
```

item_count(item filter)

The count of existing items in configuration that match the filter criteria.
 Supported value type: *Integer*.

Parameter:

- **item filter** - criteria for item selection, allows referencing by host group, host, item key, and tags. Wildcards are supported. See item filter for more details.

Comments:

- Supported only in calculated items;
- Works as an alias for the *count(exists_foreach(item_filter))* function.

Examples:

```
item_count(/*/agent.ping?[group="Host group 1"]) #calculate the number of hosts with the agent.ping item i
```

kurtosis(/host/key,(sec|#num)<:time shift>)

The "tailedness" of the probability distribution in collected values within the defined evaluation period. See also: Kurtosis.
 Supported value types: *Float*, *Integer*.
 Supported foreach function: *last_foreach*.

Parameters: see common parameters.

Examples:

```
kurtosis(/host/system.cpu.util[,user],5m)>5 #trigger on sharp spikes of CPU utilization over 5 minutes
```

mad(/host/key,(sec|#num)<:time shift>)

The median absolute deviation in collected values within the defined evaluation period. See also: Median absolute deviation.
 Supported value types: *Float*, *Integer*.
 Supported foreach function: *last_foreach*.

Parameters: see common parameters.

Examples:

```
last(/host/system.cpu.util[,user])>avg(/host/system.cpu.util[,user],1h)+3*mad(/host/system.cpu.util[,user]
```

max(/host/key,(sec|#num)<:time shift>)

The highest value of an item within the defined evaluation period.
 Supported value types: *Float*, *Integer*.
 Supported foreach functions: *avg_foreach*, *count_foreach*, *exists_foreach*, *last_foreach*, *max_foreach*, *min_foreach*, *sum_foreach*.

Parameters: see common parameters.

Examples:

```
max(/host/key,1h) - min(/host/key,1h) #calculate the difference between the maximum and minimum values wit
max(last_foreach(/*/vfs.fs.size[*,pused]?[group="Linux servers"])) #return the highest disk usage across a
```

min(/host/key,(sec|#num)<:time shift>)

The lowest value of an item within the defined evaluation period.
 Supported value types: *Float*, *Integer*.
 Supported foreach functions: *avg_foreach*, *count_foreach*, *exists_foreach*, *last_foreach*, *max_foreach*, *min_foreach*, *sum_foreach*.

Parameters: see common parameters.

Examples:

```
max(/host/key,1h) - min(/host/key,1h) #calculate the difference between the maximum and minimum values wit
min(last_foreach(/*/vfs.file.contents["/sys/class/net/enp0s3/operstate"]?[group="Linux servers"])) #return
```

skewness(/host/key,(sec|#num)<:time shift>)

The asymmetry of the probability distribution in collected values within the defined evaluation period. See also: Skewness.
 Supported value types: *Float*, *Integer*.
 Supported foreach function: *last_foreach*.

Parameters: see common parameters.

Examples:

```
skewness(/host/vm.memory.size[used],5m)<-20 #trigger on sudden negative drops indicating crashes/restarts
```

stddevpop(/host/key,(sec|#num)<:time shift>)

The population standard deviation in collected values within the defined evaluation period. See also: Standard deviation.
 Supported value types: *Float*, *Integer*.
 Supported foreach function: *last_foreach*.

Parameters: see common parameters.

Examples:

```
stddevpop(/host/system.cpu.util[,user],10m)>7 #trigger if CPU usage fluctuates heavily (population standar
```

stddevsamp(/host/key,(sec|#num)<:time shift>)

The sample standard deviation in collected values within the defined evaluation period. See also: Standard deviation.
 Supported value types: *Float*, *Integer*.
 Supported foreach function: *last_foreach*.

Parameters: see common parameters.

At least two data values are required for this function to work.

Examples:

```
stddevsamp(/host/system.cpu.util[,user],10m)>7 #trigger if CPU usage fluctuates heavily (sample standard d
```

sum(/host/key,(sec|#num)<:time shift>)

The sum of collected values within the defined evaluation period.
 Supported value types: *Float*, *Integer*.
 Supported foreach functions: *avg_foreach*, *count_foreach*, *exists_foreach*, *last_foreach*, *max_foreach*, *min_foreach*, *sum_foreach*.

Parameters: see common parameters.

Examples:

```
sum(last_foreach(/*/net.if.in[*]?[group="Linux servers"])) #calculate the total incoming network traffic f
sum(last_foreach(/*/vfs.fs.size[/,total]?[group="MySQL Servers"])) #calculate the total disk space for all
sum(last_foreach(/*/net.dns[*,*,*])) #calculate the total number of successful DNS checks
```

sumofsquares(/host/key,(sec|#num)<:time shift>)

The sum of squares in collected values within the defined evaluation period.
 Supported value types: *Float*, *Integer*.
 Supported foreach function: *last_foreach*.

Parameters: see common parameters.

Examples:

```
sumofsquares(/host/calculated.net.if.rate,1m) #the sum of squares for the network traffic rate (incoming v
sumofsquares(last_foreach(/*/net.if.in[*]?[group="Linux servers"])) #the sum of squares for latest incomin
```

varpop(/host/key,(sec|#num)<:time shift>)

The population variance of collected values within the defined evaluation period. See also: Variance.
 Supported value types: *Float*, *Integer*.
 Supported foreach function: *last_foreach*.

Parameters: see common parameters.

Examples:

```
varpop(/host/system.cpu.util[,user],10m)>50 #trigger if CPU usage fluctuates heavily (variance is above 50
```

varsamp(/host/key,(sec|#num)<:time shift>)

The sample variance of collected values within the defined evaluation period. See also: Variance.
 Supported value types: *Float*, *Integer*.
 Supported foreach function: *last_foreach*.

Parameters: see common parameters.

At least two data values are required for this function to work.

Examples:

```
varsamp(/host/system.cpu.util[,user],10m)>50 #trigger if CPU usage fluctuates heavily (sample variance is
```

See all supported functions.

Foreach functions

Overview

Foreach functions are used in aggregate calculations to return one aggregate value for each item that is selected by the used **item filter**. An array of values is returned.

For example, the `avg_foreach` function will return an array of values, where each value is the *average* history value of the selected item, during the time interval that is specified.

Attention:
Foreach functions are supported only in Calculated item formulas. They cannot be invoked in plain item keys or used directly to build trigger expressions. In trigger expressions, a calculated item can be used by referencing its key.

The item filter is part of the syntax used by foreach functions. The use of wildcards is supported in the item filter, thus the required items can be selected quite flexibly.

Supported functions

Function	Description
avg_foreach	Returns the average value for each item.
bucket_rate_foreach	Returns pairs (bucket upper bound, rate value) suitable for use in the `histogram_quantile()` function, where "bucket upper bound" is the value of item key parameter defined by the <parameter number> parameter.
count_foreach	Returns the number of values for each item.
exists_foreach	Returns '1' for each enabled item.
last_foreach	Returns the last value for each item.
max_foreach	Returns the maximum value for each item.
min_foreach	Returns the minimum value for each item.
sum_foreach	Returns the sum of values for each item.

Function syntax

Foreach functions support two common parameters: `item filter` (see details below) and `time period`:

```
foreach_function(item filter,time period)
```

For example:

```
avg_foreach(/*/mysql.qps?[group="MySQL Servers"],5m)
```

will return the five-minute average of each 'mysql.qps' item in the MySQL server group.

Note that some functions support additional parameters.

Item filter syntax

The item filter:

```
/host/key[parameters]?[conditions]
```

consists of four parts, where:
 `host` - host name;
 `key` - item key (without parameters);
 `parameters` - item key parameters;
 `conditions` - host group and/or item tag based conditions (as expression).

Spaces are allowed only inside the conditions expression.

Wildcard usage

- Wildcard can be used to replace the host name, item key or an individual item key parameter.
- Either the host or item key must be specified without wildcard. So `/host/*` and `/*/key` are valid filters, but `/*/*` is invalid.
- Wildcard cannot be used for a *part* of host name, item key, item key parameter.
- Wildcard does not match more than a single item key parameter. So a wildcard must be specified for each parameter in separation (i.e. `key[abc,*,*]`).

Conditions expression

The conditions expression supports:

- Operands:
 group - host group;
 tag - item tag;
 "<text>" - string constant, with the \ escape character to escape " and \.
- Case-sensitive string comparison operators: =, <>
- Logical operators: `and`, `or`, `not`
- Grouping with parentheses: ( )

Quotation of string constants is mandatory. Only case-sensitive full string comparison is supported.

Warning:
When specifying tags in the filter (i.e. `tag="tagname:value"`), the colon ":" is used as a delimiter. Everything after it is considered the tag value. Thus it is currently not supported to specify a tag name containing ":" in it.

Examples

A complex filter may be used, referencing the item key, host group and tags, as illustrated by the examples:

Syntax example	Description
`/host/key[abc,*]`	Matches similar items on this host.
`/*/key`	Matches the same item of any host.
`/*/key?[group="ABC" and tag="tagname:value"]`	Matches the same item of any host from the ABC group having 'tagname:value' tags.
`/*/key[a,*,c]?[(group="ABC" and tag="Tag1") or (group="DEF" and (tag="Tag2" or tag="Tag3:value"))]`	Matches similar items of any host from the ABC or DEF group with the respective tags.

All referenced items must exist and collect data. Only enabled items on enabled hosts are included in the calculations. Items in the unsupported state are not included.

Attention:
If the item key of a referenced item is changed, the filter must be updated manually.

Specifying a parent host group includes the parent group and all nested host groups with their items.

Time period

The **second** parameter allows you to specify the time period for aggregation. The time period can only be expressed as time, the amount of values (prefixed with **#**) is not supported.

Supported unit symbols can be used in this parameter for convenience, for example, `5m` (five minutes) instead of 300s (300 seconds) or `1d` (one day) instead of 86400 (86400 seconds).

For the `last_foreach` function time period is an optional parameter (supported since Zabbix 7.0), which can be omitted:

```
last_foreach(/*/key?[group="host group"])
```

Time period is not supported with the `exists_foreach` function.

Additional parameters

bucket_rate_foreach

A third optional parameter is supported by the `bucket_rate_foreach` function:

```
bucket_rate_foreach(item filter,time period,<parameter number>)
```

where `<parameter number>` is the position of the "bucket" value in the item key. For example, if the "bucket" value in `myItem[aaa,0.2]` is `0.2`, then its position is 2.

The default value of <parameter number> is 1.

count_foreach

Third and fourth optional parameters are supported by the `count_foreach` function:

```
count_foreach(item filter,time period,<operator>,<pattern>)
```

Where:

- **operator** is the conditional operator for item values (must be double-quoted). Supported `operators`:
eq - equal
ne - not equal
gt - greater
ge - greater or equal
lt - less
le - less or equal
like - matches if contains pattern (case-sensitive)
bitand - bitwise AND
regexp - case-sensitive match of the regular expression given in `pattern`
iregexp - case-insensitive match of the regular expression given in `pattern`.
- **pattern** is the required pattern (string arguments must be double-quoted); supported if *operator* is specified in the third parameter.

Comments:

- Optional parameters *operator* or *pattern* can't be left empty after a comma, only fully omitted.
- With *bitand* as the third parameter, the fourth `pattern` parameter can be specified as two numbers, separated by '/': **number_to_compare_with/mask**. `count_foreach()` calculates "bitwise AND" from the value and the *mask* and compares the result to *number_to_compare_with*. If the result of "bitwise AND" is equal to *number_to_compare_with*, the value is counted.
If *number_to_compare_with* and *mask* are equal, only the *mask* need be specified (without '/').
- With *regexp* or *iregexp* as the third parameter, the fourth `pattern` parameter can be an ordinary or global (starting with '@') regular expression. In case of global regular expressions case sensitivity is inherited from global regular expression settings. For the purpose of regexp matching, float values will always be represented with 4 decimal digits after '.'. Also note that for large numbers difference in decimal (stored in database) and binary (used by Zabbix server) representation may affect the 4th decimal digit.

See aggregate calculations for more details and examples on using foreach functions.

Behavior depending on availability

The following table illustrates how each function behaves in cases of limited availability of host/item and history data.

Function	Disabled host	Unavailable host with data	Unavailable host without data	Disabled item	Unsupported item	Data retrieval error (SQL)
avg_foreach	ignore	return avg	ignore	ignore	ignore	ignore
bucket_rate_foreach	ignore	return bucket rate	ignore	ignore	ignore	ignore
count_foreach	ignore	return count	0	ignore	ignore	ignore
exists_foreach	ignore	1	1	ignore	1	n/a
last_foreach	ignore	return last	ignore	ignore	ignore	ignore
max_foreach	ignore	return max	ignore	ignore	ignore	ignore
min_foreach	ignore	return min	ignore	ignore	ignore	ignore
sum_foreach	ignore	return sum	ignore	ignore	ignore	ignore

If the item is *ignored*, nothing is added to the aggregation.

Bitwise functions

All functions listed here are supported in:

- Trigger expressions
- Calculated item formulas
- Expression macros

The functions are listed without additional information. Click the function to see the full details.

Function	Description
bitand	The value of "bitwise AND" of an item value and mask.
bitlshift	The bitwise shift left of an item value.
bitnot	The value of "bitwise NOT" of an item value.
bitor	The value of "bitwise OR" of an item value and mask.

Function	Description
bitrshift	The bitwise shift right of an item value.
bitxor	The value of ”bitwise exclusive OR” of an item value and mask.

Function details

Some general notes on function parameters:

- Function parameters are separated by a comma
- Expressions are accepted as parameters
- Optional function parameters (or parameter parts) are indicated by < >

bitand(value,mask)

The value of ”bitwise AND” of an item value and mask.
 Supported value types: *Integer*.

Parameters:

- **value** - the value to check;
- **mask** (mandatory) - a 64-bit unsigned integer (0 - 18446744073709551615).

Although the comparison is done in a bitwise manner, all the values must be supplied and are returned in decimal. For example, checking for the 3rd bit is done by comparing to 4, not 100.

Examples:

```
bitand(last(/host/key),12)=8 or bitand(last(/host/key),12)=4 #3rd or 4th bit set, but not both at the same
bitand(last(/host/key),20)=16 #3rd bit not set and 5th bit set
```

bitlshift(value,bits to shift)

The bitwise shift left of an item value.
 Supported value types: *Integer*.

Parameters:

- **value** - the value to check;
- **bits to shift** (mandatory) - the number of bits to shift.

Although the comparison is done in a bitwise manner, all the values must be supplied and are returned in decimal. For example, checking for the 3rd bit is done by comparing to 4, not 100.

bitnot(value)

The value of ”bitwise NOT” of an item value.
 Supported value types: *Integer*.

Parameter:

- **value** - the value to check.

Although the comparison is done in a bitwise manner, all the values must be supplied and are returned in decimal. For example, checking for the 3rd bit is done by comparing to 4, not 100.

bitor(value,mask)

The value of ”bitwise OR” of an item value and mask.
 Supported value types: *Integer*.

Parameters:

- **value** - the value to check;
- **mask** (mandatory) - a 64-bit unsigned integer (0 - 18446744073709551615).

Although the comparison is done in a bitwise manner, all the values must be supplied and are returned in decimal. For example, checking for the 3rd bit is done by comparing to 4, not 100.

bitrshift(value,bits to shift)

The bitwise shift right of an item value.
 Supported value types: *Integer*.

Parameters:

- **value** - the value to check;
- **bits to shift** (mandatory) - the number of bits to shift.

Although the comparison is done in a bitwise manner, all the values must be supplied and are returned in decimal. For example, checking for the 3rd bit is done by comparing to 4, not 100.

bitxor(value,mask)

The value of "bitwise exclusive OR" of an item value and mask.
 Supported value types: *Integer*.

Parameters:

- **value** - the value to check;
- **mask** (mandatory) - a 64-bit unsigned integer (0 - 18446744073709551615).

Although the comparison is done in a bitwise manner, all the values must be supplied and are returned in decimal. For example, checking for the 3rd bit is done by comparing to 4, not 100.

See all supported functions.

Date and time functions

All functions listed here are supported in:

- Trigger expressions
- Calculated item formulas
- Expression macros

Attention:
Date and time functions cannot be used in the expression by themselves; at least one function from another group, referencing the host item, must be included in the expression (except the nodata() function). For detailed information on how date and time functions work within expressions, see Calculation time.

The functions are listed without additional information. Click the function to see the full details.

Function	Description
date	The current date in YYYYMMDD format.
dayofmonth	The day of month in range of 1 to 31.
dayofweek	The day of week in range of 1 to 7.
now	The number of seconds since the Epoch (00:00:00 UTC, January 1, 1970).
time	The current time in HHMMSS format.

Function details

date

The current date in YYYYMMDD format.

Example:

```
date()<20220101
```

dayofmonth

The day of month in range of 1 to 31.

Example:

```
dayofmonth()=1
```

dayofweek

The day of week in range of 1 to 7 (Mon - 1, Sun - 7).

Example (only weekdays):

```
dayofweek()<6
```

Example (only weekend):

```
dayofweek()>5
```

now

The number of seconds since the Epoch (00:00:00 UTC, January 1, 1970).

Example:

`now()<1640998800`

time

The current time in HHMMSS format.

Example (only nighttime, 00:00-06:00):

`time()<060000`

See all supported functions.

History functions

All functions listed here are supported in:

- Trigger expressions
- Calculated item formulas
- Expression macros

The functions are listed without additional information. Click the function to see the full details.

Function	Description
change	The amount of difference between the previous and latest value.
changecount	The number of changes between adjacent values within the defined evaluation period.
count	The number of values within the defined evaluation period.
countunique	The number of unique values within the defined evaluation period.
find	Find a value match within the defined evaluation period.
first	The first (the oldest) value within the defined evaluation period.
firstclock	The timestamp of the first (the oldest) value within the defined evaluation period.
fuzzytime	Check how much the passive agent time differs from the Zabbix server/proxy time.
jsonpath	Return an array of JSONPath query results.
last	The most recent value.
lastclock	The timestamp of the Nth most recent value within the defined evaluation period.
logeventid	Check if the event ID of the last log entry matches a regular expression.
logseverity	The log severity of the last log entry.
logsource	Check if log source of the last log entry matches a regular expression.
logtimestamp	The log message timestamp of the Nth most recent log item value.
monodec	Check if there has been a monotonous decrease in values.
monoinc	Check if there has been a monotonous increase in values.
nodata	Check for no data received.
percentile	The P-th percentile of a period, where P (percentage) is specified by the third parameter.
rate	The per-second average rate of the increase in a monotonically increasing counter within the defined time period.
xmlxpath	Return an array of XML XPath query results.

Common parameters

- `/host/key` is a common mandatory first parameter for the functions referencing the host item history.
- `(sec|#num)<:time shift>` is a common second parameter for the functions referencing the host item history, where:
  - **sec** - maximum evaluation period in seconds (time suffixes can be used), or
  - **#num** - maximum evaluation range in latest collected values (if preceded by a hash mark).
  - **time shift** (optional) allows the evaluation point to be moved back in time. See more details on specifying time shift.

Function details

Some general notes on function parameters:

- Function parameters are separated by a comma
- Optional function parameters (or parameter parts) are indicated by < >
- Function-specific parameters are described with each function
- `/host/key` and `(sec|#num)<:time shift>` parameters must never be quoted

change(/host/key)

The amount of difference between the previous and latest value.
 Supported value types: *Float*, *Integer*, *String*, *Text*, *Log*.
 For strings returns: 0 - values are equal; 1 - values differ.

Parameters: see common parameters.

Comments:

- Numeric difference will be calculated, as seen with these incoming example values ('previous' and 'latest' value = difference):
'1' and '5' = +4
'3' and '1' = −2
'0' and '-2.5' = −2.5.
- In triggers, useful for catching sudden spikes (or drops), counter resets, numeric state transitions.
- See also: abs for comparison.

Examples:

```
change(/host/system.uptime)<0 #system uptime change has been negative since the last value (indicating a r
change(/host/system.cpu.load[all,avg1])>2 #CPU load (for one minute) has jumped by more than 2 since the l
change(/host/vfs.fs.size[/,free])<-1G #free disk space has dropped by more than 1 GB between checks
```

changecount(/host/key,(sec|#num)<:time shift>,<mode>)

The number of changes between adjacent values within the defined evaluation period.
 Supported value types: *Float*, *Integer*, *String*, *Text*, *Log*.

Parameters:

- See common parameters
- **mode** (must be double-quoted) - possible values: *all* - count all changes (default); *dec* - count decreases; *inc* - count increases

Comments:

- For non-numeric value types, the *mode* parameter is ignored.
- In triggers, useful for catching frequent state changes (indicating instability).

Examples:

```
changecount(/host/icmpping,10m)>5 #ping status has changed more than 5 times in 10 minutes
changecount(/host/vfs.file.contents["/sys/class/net/eth0/operstate"],1h)>5 #operational state of eth0 has
changecount(/host/proc.num[httpd],15m)>10 #the number of httpd processes has changed more than 10 times in
changecount(/host/key,#10,"inc") #the number of value increases (relative to the adjacent value) among the
changecount(/host/key,24h,"dec") #the number of value decreases (relative to the adjacent value) for the l
```

count(/host/key,(sec|#num)<:time shift>,<operator>,<pattern>)

The number of values within the defined evaluation period.
 Supported value types: *Float*, *Integer*, *String*, *Text*, *Log*.

Parameters:

- See common parameters
- **operator** (must be double-quoted) Supported `operators`:
eq - equal (default for integer, float)
ne - not equal
gt - greater
ge - greater or equal
lt - less
le - less or equal
like (default for string, text, log) - matches if contains pattern (case-sensitive)
bitand - bitwise AND
regexp - case-sensitive match of the regular expression given in `pattern`
iregexp - case-insensitive match of the regular expression given in `pattern`
- **pattern** - the required pattern (string arguments must be double-quoted)

Comments:

- Float items match with the precision of 2.22e-16.
- *like* is not supported as operator for integer values.
- *like* and *bitand* are not supported as operators for float values.
- For string, text, and log values only *eq*, *ne*, *like*, *regexp* and *iregexp* operators are supported.
- With *bitand* as operator, the fourth `pattern` parameter can be specified as two numbers, separated by '/': **number_to_compare_with/mask**. count() calculates "bitwise AND" from the value and the *mask* and compares the result to *number_to_compare_with*. If the result of "bitwise AND" is equal to *number_to_compare_with*, the value is counted.
If *number_to_compare_with* and *mask* are equal, only the *mask* need be specified (without '/').
- With *regexp* or *iregexp* as operator, the fourth `pattern` parameter can be an ordinary or global (starting with '@') regular expression. In case of global regular expressions case sensitivity is inherited from global regular expression settings. For the purpose of regexp matching, float values will always be represented with 4 decimal digits after '.'. Also note that for large numbers difference in decimal (stored in database) and binary (used by Zabbix server) representation may affect the 4th decimal digit.

Examples:

```
count(/host/icmpping,30m,,"0")>5 #ping has failed more than 5 times in 30 minutes
count(/host/key,10m,"like","error") #the number of values for the last 10 minutes until now that contain '
count(/host/key,10m,,12) #the number of values for the last 10 minutes until now that equal '12'
count(/host/key,10m,"gt",12) #the number of values for the last 10 minutes until now that are over '12'
count(/host/key,#10,"gt",12) #the number of values within the last 10 values until now that are over '12'
count(/host/key,10m:now-1d,"gt",12) #the number of values between 24 hours and 10 minutes and 24 hours ago
count(/host/key,10m,"bitand","6/7") #the number of values for the last 10 minutes until now having '110' (
count(/host/key,10m:now-1d) #the number of values between 24 hours and 10 minutes and 24 hours ago from no
```

countunique(/host/key,(sec|#num)<:time shift>,<operator>,<pattern>)

The number of unique values within the defined evaluation period.
 Supported value types: *Float*, *Integer*, *String*, *Text*, *Log*.

Parameters:

- See common parameters
- **operator** (must be double-quoted). Supported `operators`:
eq - equal (default for integer, float)
ne - not equal
gt - greater
ge - greater or equal
lt - less
le - less or equal
like (default for string, text, log) - matches if contains pattern (case-sensitive)
bitand - bitwise AND
regexp - case-sensitive match of the regular expression given in `pattern`
iregexp - case-insensitive match of the regular expression given in `pattern`
- **pattern** - the required pattern (string arguments must be double-quoted)

Comments:

- Float items match with the precision of 2.22e-16.
- *like* is not supported as operator for integer values.
- *like* and *bitand* are not supported as operators for float values.
- For string, text, and log values only *eq*, *ne*, *like*, *regexp* and *iregexp* operators are supported.
- With *bitand* as operator, the fourth `pattern` parameter can be specified as two numbers, separated by '/': **number_to_compare_with/mask**. countunique() calculates "bitwise AND" from the value and the *mask* and compares the result to *number_to_compare_with*. If the result of "bitwise AND" is equal to *number_to_compare_with*, the value is counted.
If *number_to_compare_with* and *mask* are equal, only the *mask* need be specified (without '/').
- With *regexp* or *iregexp* as operator, the fourth `pattern` parameter can be an ordinary or global (starting with '@') regular expression. In case of global regular expressions case sensitivity is inherited from global regular expression settings. For the purpose of regexp matching, float values will always be represented with 4 decimal digits after '.'. Also note that for large numbers difference in decimal (stored in database) and binary (used by Zabbix server) representation may affect the 4th decimal digit.

Examples:

```
countunique(/host/key,10m) #the number of unique values for the last 10 minutes until now
countunique(/host/key,10m,"like","error") #the number of unique values for the last 10 minutes until now t
countunique(/host/key,10m,,12) #the number of unique values for the last 10 minutes until now that equal '
countunique(/host/key,10m,"gt",12) #the number of unique values for the last 10 minutes until now that are
countunique(/host/key,#10,"gt",12) #the number of unique values within the last 10 values until now that a
countunique(/host/key,10m:now-1d,"gt",12) #the number of unique values between 24 hours and 10 minutes and
countunique(/host/key,10m,"bitand","6/7") #the number of unique values for the last 10 minutes until now h
countunique(/host/key,10m:now-1d) #the number of unique values between 24 hours and 10 minutes and 24 hour
```

find(/host/key,(sec|#num)<:time shift>,<operator>,<pattern>)

Find a value match within the defined evaluation period.
 Supported value types: *Float*, *Integer*, *String*, *Text*, *Log*.
 Returns: 1 - found; 0 - otherwise.

Parameters:

- See common parameters
- **sec** or **#num** (optional) - defaults to the latest value if not specified
- **operator** (must be double-quoted) Supported `operators`:
eq - equal (default for integer, float)
ne - not equal
gt - greater
ge - greater or equal
lt - less
le - less or equal
like (default for string, text, log) - matches if contains the string given in `pattern` (case-sensitive)
bitand - bitwise AND
regexp - case-sensitive match of the regular expression given in `pattern`
iregexp - case-insensitive match of the regular expression given in `pattern`
- **pattern** - the required pattern (string arguments must be double-quoted); Perl Compatible Regular Expression (PCRE) regular expression if `operator` is *regexp*, *iregexp*

Comments:

- If more than one value is processed, '1' is returned if there is at least one matching value.

- *like* is not supported as operator for integer values.
- *like* and *bitand* are not supported as operators for float values.
- For string, text, and log values only *eq*, *ne*, *like*, *regexp* and *iregexp* operators are supported.
- With *regexp* or *iregexp* as operator, the fourth `pattern` parameter can be an ordinary or global (starting with '@') regular expression. In case of global regular expressions case sensitivity is inherited from the global regular expression settings.

Examples:

```
find(/host/key,10m,"like","error") #find a value that contains 'error' within the last 10 minutes until no
find(/host/agent.version,,"like","beta")=1 #Zabbix agent has beta version, must be upgraded
find(/host/log[/var/log/nginx/access.log],,"regexp"," 500 ")=1 #internal web server error has been found
```

first(/host/key,sec<:time shift>)

The first (the oldest) value within the defined evaluation period.
 Supported value types: *Float*, *Integer*, *String*, *Text*, *Log*.

Parameters:

- See common parameters

See also last().

Example:

```
first(/host/key,1h) #retrieve the oldest value within the last hour until now
```

firstclock(/host/key,sec<:time shift>)

The timestamp of the oldest value within the defined evaluation period.
 Supported value types: *Float*, *Integer*, *String*, *Text*, *Log*.

Parameters:

- See common parameters.

The function fails with an error if no data has been collected in the given period.

See also lastclock().

Examples:

```
firstclock(/host/key,1h) #retrieve the timestamp of the oldest value within the last hour
firstclock(/host/key,1h:now-24h) #retrieve the timestamp of the oldest value within the last hour a day ag
```

fuzzytime(/host/key,sec)

Check how much the passive agent time differs from the Zabbix server/proxy time.
 Supported value types: *Float*, *Integer*.
 Returns: 1 - difference between the passive item value (as timestamp) and Zabbix server/proxy timestamp (the clock of value collection) is less than or equal to *sec* seconds; 0 - otherwise.

Parameters:

- See common parameters.

Comments:

- Usually used with the 'system.localtime' item to check that local time is in sync with the local time of Zabbix server. *Note* that 'system.localtime' must be configured as a passive check for Zabbix agent; on Zabbix agent 2 it may be configured as an active check.
- Can be used also with the `vfs.file.time[/path/file,modify]` key to check that the file did not get updates for long time.
- This function is not recommended for use in complex trigger expressions (with multiple items involved), because it may cause unexpected results (time difference will be measured with the most recent metric), e.g., in `fuzzytime(/Host/system.localtime,60 or last(/Host/trap)<>0`.

Examples:

```
fuzzytime(/host/system.localtime,5m)=0 #client local time differs from Zabbix server/proxy time by more th
fuzzytime(/host/system.localtime,5m)=0 and nodata(/host/system.localtime,10m)=0 #client local time differs
```

jsonpath(/host/key,(sec|#num)<:time shift>,pattern)

Return an array of JSONPath query results.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- See common parameters
- **pattern** - the required pattern (string arguments must be enclosed in double quotes)

Comments:

- This function is used with aggregate functions.
- A JSONPath query is executed on each selected history value; each non-empty successful query result is added to the array. If a query result for a single history value contains multiple matches, each match is added to the array separately.
- The function fails with an error if:
  - The `pattern` format is invalid.
  - Any of the history values contains invalid JSON or the JSONPath query fails for that value.
- The function returns an empty array if:
  - No history values were selected.
  - All selected history values contain invalid JSON.

Example:

```
max(jsonpath(/host/http.response,5m,"$.response_time")) > 500
```

last(/host/key,<#num<:time shift»)

The most recent value.
 Supported value types: *Float*, *Integer*, *String*, *Text*, *Log*.

Parameters:

- See common parameters
- **#num** (optional) - the Nth most recent value

Comments:

- Take note that a hash-tagged time period (#N) works differently here than with many other functions. For example: `last(/host/key)` is always equal to `last(/host/key,#1)`; `last(/host/key,#3)` - the third most recent value (*not* three latest values).
- Zabbix does not guarantee the exact order of values if more than two values exist within one second in history.
- See also first().

Examples:

```
last(/host/key) #retrieve the last value
last(/host/key,#2) #retrieve the previous value
last(/host/key,#1)<>last(/host/key,#2) #last two values differ
```

lastclock(/host/key,<#num<:time shift»)

The timestamp of the Nth most recent value within the defined evaluation period.
 Supported value types: *Float*, *Integer*, *String*, *Text*, *Log*.

Parameters:

- See common parameters;
- **#num** (optional) - the Nth most recent value.

The function fails with an error if no data has been collected in the given period or no Nth value has been collected.

See also firstclock().

Examples:

```
lastclock(/host/key) #retrieve the timestamp of the latest value
lastclock(/host/key,#2) #retrieve the timestamp of the 2nd latest value
now()-lastclock(/host/heartbeat.item)>300 #more than 5 minutes have passed since the heartbeat item was up
lastclock(/host/system.cpu.load[all,avg1])-lastclock(/host/system.cpu.load[all,avg1],#2)>300 #the update i
((now()-lastclock(/host/system.cpu.load[all,avg1])<120 and last(/host/system.cpu.load[all,avg1])>5)) #CPU
```

logeventid(/host/key,<#num<:time shift»,<pattern>)

Check if the event ID of the last log entry matches a regular expression.
 Supported value types: *Log*.
 Returns: 0 - does not match; 1 - matches.

Parameters:

- See common parameters
- **#num** (optional) - the Nth most recent value

- **pattern** (optional) - the regular expression describing the required pattern, Perl Compatible Regular Expression (PCRE) style (string arguments must be double-quoted)

Examples:

```
logeventid(/host/eventlog[Security],,"^4625$")=1 #a log entry with ID matching "4625" (failed authenticati
logeventid(/host/eventlog[System],,"^(6008|41)$")=1 #a log entry with ID matching "6008" or "41" found
```

logseverity(/host/key,<#num<:time shift»)

Log severity of the last log entry.
 Supported value types: *Log*.
 Returns: 0 - default severity; N - severity (integer, useful for Windows event logs: 1 - Information, 2 - Warning, 4 - Error, 7 - Failure Audit, 8 - Success Audit, 9 - Critical, 10 - Verbose).

Parameters:

- See common parameters
- **#num** (optional) - the Nth most recent value

Zabbix takes log severity from the **Information** field of Windows event log.

Examples:

```
logseverity(/host/log[/var/log/syslog],10m)>3 #a log entry with severity above "3" found
logseverity(/host/eventlog[System],10m)=4 #a log entry with severity equaling "Error" found
```

logsource(/host/key,<#num<:time shift»,<pattern>)

Check if log source of the last log entry matches a regular expression.
 Supported value types: *Log*.
 Returns: 0 - does not match; 1 - matches.

Parameters:

- See common parameters
- **#num** (optional) - the Nth most recent value
- **pattern** (optional) - the regular expression describing the required pattern, Perl Compatible Regular Expression (PCRE) style (string arguments must be double-quoted)

Normally used for Windows event logs.

Examples:

```
logsource(/host/eventlog[Application],,"MSSQLSERVER")=1 #a log entry with source matching "MSSQLSERVER" fo
logsource(/host/eventlog[System],,"Service Control Manager")=1 #a log entry with source matching "Service
logsource(/host/eventlog[System],,"Service Control Manager")=1 and logeventid(/host/eventlog[System],,"^70
```

logtimestamp(/host/key,<#num<:time shift»)

The log message timestamp of the Nth most recent log item value.
 Supported value types: *Log*.

Parameters:

- See common parameters;
- **#num** (optional) - the Nth most recent value.

Comments:

- Calculation of time shift is based on the item value clock, not on the log message timestamp;
- The function fails with an error if:
  - a non-log type item is received;
  - no data has been collected in the given period;
  - no Nth value has been collected;
  - log message does not contain a timestamp value.

Examples:

```
logtimestamp(/host/key) #retrieve the timestamp of the latest log message
logtimestamp(/host/key,#2) #retrieve the timestamp of the 2nd latest log message
logtimestamp(/host/key,#3:now-1d) #retrieve the timestamp of the 3rd latest log message a day ago
```

monodec(/host/key,(sec|#num)<:time shift>,<mode>)

Check if there has been a monotonous decrease in values.
 Supported value types: *Integer*.
 Returns: 1 - if all elements in the time period continuously decrease; 0 - otherwise.

Parameters:

- See common parameters
- **mode** (must be double-quoted) - *weak* (every value is smaller or the same as the previous one; default) or *strict* (every value has decreased)

Examples:

```
monodec(/host/system.swap.size[all,free],60s) + monodec(/host2/system.swap.size[all,free],60s) + monodec(/
monodec(/host/proc.num[nginx],10m,"strict")=1 #the number of nginx processes has monotonously decreased ov
```

monoinc(/host/key,(sec|#num)<:time shift>,<mode>)

Check if there has been a monotonous increase in values.
 Supported value types: *Integer*.
 Returns: 1 - if all elements in the time period continuously increase; 0 - otherwise.

Parameters:

- See common parameters
- **mode** (must be double-quoted) - *weak* (every value is bigger or the same as the previous one; default) or *strict* (every value has increased)

Examples:

```
monoinc(/host/system.localtime,#3,"strict")=0 #the system local time has not been increasing consistently
monoinc(/host/vfs.dir.count[/mnt/data/logs],24h,"weak")=0 #trigger if the file count has stagnated over 24
```

nodata(/host/key,sec,<mode>)

Check for no data received.
 Supported value types: *Integer*, *Float*, *Character*, *Text*, *Log*.
 Returns: 1 - if no data received during the defined period of time; 0 - otherwise.

Parameters:

- See common parameters
- **sec** - the period should not be less than 30 seconds because the history syncer process calculates this function only every 30 seconds; `nodata(/host/key,0)` is disallowed
- **mode** - if set to *strict* (double-quoted), this function will be insensitive to proxy availability (see comments for details)

Comments:

- the 'nodata' triggers monitored by proxy are, by default, sensitive to proxy availability - if proxy becomes unavailable, the 'nodata' triggers will not fire immediately after a restored connection, but will skip the data for the delayed period. Note that for passive proxies suppression is activated if connection is restored more than 15 seconds and no less than 2 seconds later. For active proxies suppression is activated if connection is restored more than 15 seconds later. To turn off sensitiveness to proxy availability, use the third parameter, e.g.: `nodata(/host/key,5m,"strict")`; in this case the function will fire as soon as the evaluation period (five minutes) without data has past.
- This function will display an error if, within the period of the 1st parameter:
- there's no data and Zabbix server was restarted
- there's no data and maintenance was completed
- there's no data and the item was added or re-enabled.
- Errors are displayed in the *Info* column of the trigger list.
- This function may not work properly if there are time differences between Zabbix server, proxy and agent. See also: Time synchronization requirement.

Example:

```
nodata(/host/agent.ping,5m)=1 #trigger if no data has been received from Zabbix agent for 5 minutes
```

percentile(/host/key,(sec|#num)<:time shift>,percentage)

The P-th percentile of a period, where P (percentage) is specified by the third parameter.
 Supported value types: *Float*, *Integer*.

Parameters:

- See common parameters
- **percentage** - a floating-point number between 0 and 100 (inclusive) with up to 4 digits after the decimal point

Examples:

```
percentile(/host/net.if.in[eth0,bytes],1h,95)>1000000 #95th percentile network input (bytes/sec) over 1 ho
percentile(/host/system.cpu.util,5m,95)>80 #95th percentile of CPU utilization user-time percentage has go
percentile(/host/icmppingsec[192.168.0.2],15m,95)>0.15 #most latency measurements are below 150 ms but the
percentile(/host/net.if.in[eth0,bytes],1h,50) #calculate the 50th percentile (the median value) of incomin
(percentile(/host/net.if.in[eth0,bytes],1h,50)+percentile(/host/net.if.in[eth0,bytes],1h,51))/2 #calculate
```

rate(/host/key,sec<:time shift>)

The average per-second increase in a monotonically increasing counter within the defined time period.
 Supported value types: *Float*, *Integer*.

Parameters:

- See common parameters

Functionally corresponds to 'rate' of PromQL.

Examples:

```
rate(/host/key,30s) #if the monotonic increase over 30 seconds is 20, this function returns 0.67.
rate(/host/net.if.in[eth0,bytes],5m)>500000 #the incoming interface traffic rate on eth0 has exceeded 500
rate(/host/app.requests.count,1m)>100 #the request count counter has risen at >100 requests per second in
```

xmlxpath(/host/key,(sec|#num)<:time shift>,pattern)

Return an array of XML XPath query results.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- See common parameters
- **pattern** - the required pattern (string arguments must be enclosed in double quotes)

Comments:

- This function is used with aggregate functions.
- An XML XPath query is executed on each selected history value; each non-empty successful query result is added to the array. If a query result for a single history value contains multiple matches, each match is added to the array separately.
- The function fails with an error if:
  - The `pattern` format is invalid.
  - Zabbix was compiled without XML support.
  - Any of the history values contains invalid XML or the XPath query fails for that value.
- The function returns an empty array if:
  - No history values were selected.
  - All selected history values contain invalid XML.

Example:

```
max(xmlxpath(/host/http.response,5m,"/response/response_time")) > 500
```

See all supported functions.

Trend functions

Trend functions, in contrast to history functions, use trend data for calculations.

Trends store hourly aggregate values. Trend functions use these hourly averages, and thus are useful for long-term analysis.

Trend function results are cached so multiple calls to the same function with the same parameters fetch info from the database only once. The trend function cache is controlled by the TrendFunctionCacheSize server parameter.

Use `zabbix[tcache,cache,paccessed]` to estimate how much of the cache is actually used. Low `paccessed`, together with low `pitems` and low `pmisses`, usually means that the cache size can be reduced.

Triggers that reference trend functions **only** are evaluated once per the smallest time period in the expression. For instance, a trigger like

```
trendavg(/host/key,1d:now/d) > 1 or trendavg(/host/key2,1w:now/w) > 2
```

will be evaluated once per day. If the trigger contains both trend and history (or date and time and/or nodata()) functions, it is calculated in accordance with the usual principles.

All functions listed here are supported in:

- Trigger expressions
- Calculated items
- Expression macros

The functions are listed without additional information. Click the function to see the full details.

Function	Description
baselinedev	Returns the number of deviations (by stddevpop algorithm) between the last data period and the same data periods in preceding seasons.
baselinewma	Calculates the baseline by averaging data from the same timeframe in multiple equal time periods (’seasons’) using the weighted moving average algorithm.
trendavg	The average of trend values within the defined time period.
trendcount	The number of successfully retrieved history values used to calculate the trend value within the defined time period.
trendmax	The maximum in trend values within the defined time period.
trendmin	The minimum in trend values within the defined time period.
trendstl	Returns the rate of anomalies during the detection period - a decimal value between 0 and 1 that is `((the number of anomaly values)/(total number of values))`.
trendsum	The sum of trend values within the defined time period.

Common parameters

- `/host/key` is a common mandatory first parameter
- `time period:time shift` is a common second parameter, where:
  - **time period** - the time period (minimum ’1h’), defined as <N><time unit> where `N` - the number of time units, `time unit` - h (hour), d (day), w (week), M (month) or y (year).
  - **time shift** - the time period offset (see function examples)

Function details

Some general notes on function parameters:

- Function parameters are separated by a comma
- Optional function parameters (or parameter parts) are indicated by < >
- Function-specific parameters are described with each function
- `/host/key` and `time period:time shift` parameters must never be quoted

baselinedev(/host/key,data period:time shift,season unit,num seasons)

Returns the number of deviations (by stddevpop algorithm) between the last data period and the same data periods in preceding seasons.

Parameters:

- See common parameters
- **data period** - the data gathering period within a season, defined as <N><time unit> where:
`N` - the number of time units
`time unit` - h (hour), d (day), w (week), M (month) or y (year), must be equal to or less than season
- **season unit** - the calendar unit that defines one season (h, d, w, M, y); cannot be smaller than data period
- **num seasons** - the number of seasons to evaluate

Examples:

```
baselinedev(/host/key,1d:now/d,"M",6) #calculating the number of standard deviations (population) between
baselinedev(/host/key,1h:now/h,"d",10)  #calculating the number of population standard deviations between
```

baselinewma(/host/key,data period:time shift,season unit,num seasons)

Calculates the baseline by averaging data from the same timeframe in multiple equal time periods (’seasons’) using the weighted moving average algorithm.

Parameters:

- See common parameters
- **data period** - the data gathering period within a season, defined as <N><time unit> where:
`N` - the number of time units
`time unit` - h (hour), d (day), w (week), M (month) or y (year), must be equal to or less than season
Time shift - the time period offset, defines the end of data gathering time frame in seasons (see examples)
- **season unit** - the calendar unit that defines one season (h, d, w, M, y); cannot be smaller than data period
- **num seasons** - the number of seasons to evaluate

Examples:

```
baselinewma(/host/key,1h:now/h,"d",3) #calculating the baseline from the same hour-of-day across the last
baselinewma(/host/key,2h:now/h,"d",3) #calculating the baseline from the same two-hour timeframe across th
baselinewma(/host/key,1d:now/d,"M",4) #calculating the baseline from the same day-of-month as 'yesterday'
```

trendavg(/host/key,time period:time shift)

The average of trend values within the defined time period.

Parameters:

- See common parameters

Examples:

```
trendavg(/host/key,1h:now/h) #the average for the previous hour (e.g., 12:00-13:00)
trendavg(/host/key,1h:now/h-1h) #the average for two hours ago (11:00-12:00)
trendavg(/host/key,1h:now/h-2h) #the average for three hours ago (10:00-11:00)
trendavg(/host/key,1M:now/M-1y) #the average for the previous month a year ago
```

trendcount(/host/key,time period:time shift)

The number of successfully retrieved history values used to calculate the trend value within the defined time period.

Parameters:

- See common parameters

Examples:

```
trendcount(/host/key,1h:now/h) #the value count for the previous hour (e.g., 12:00-13:00)
trendcount(/host/key,1h:now/h-1h) #the value count for two hours ago (11:00-12:00)
trendcount(/host/key,1h:now/h-2h) #the value count for three hours ago (10:00-11:00)
trendcount(/host/key,1M:now/M-1y) #the value count for the previous month a year ago
```

trendmax(/host/key,time period:time shift)

The maximum in trend values within the defined time period.

Parameters:

- See common parameters

Examples:

```
trendmax(/host/key,1h:now/h) #the maximum for the previous hour (e.g., 12:00-13:00)
trendmax(/host/key,1h:now/h) - trendmin(/host/key,1h:now/h) → calculate the difference between the maximum
trendmax(/host/key,1h:now/h-1h) #the maximum for two hours ago (11:00-12:00)
trendmax(/host/key,1h:now/h-2h) #the maximum for three hours ago (10:00-11:00)
trendmax(/host/key,1M:now/M-1y) #the maximum for the previous month a year ago
```

trendmin(/host/key,time period:time shift)

The minimum in trend values within the defined time period.

Parameters:

- See common parameters

Examples:

```
trendmin(/host/key,1h:now/h) #the minimum for the previous hour (e.g., 12:00-13:00)
trendmax(/host/key,1h:now/h) - trendmin(/host/key,1h:now/h) → calculate the difference between the maximum
trendmin(/host/key,1h:now/h-1h) #the minimum for two hours ago (11:00-12:00)
trendmin(/host/key,1h:now/h-2h) #the minimum for three hours ago (10:00-11:00)
trendmin(/host/key,1M:now/M-1y) #the minimum for the previous month a year ago
```

trendstl(/host/key,eval period:time shift,detection period,season,<deviations>,<devalg>,<s window>)

Returns the rate of anomalies during the detection period - a decimal value between 0 and 1 that is `((the number of anomaly values)/(total number of values))`.

Parameters:

- See common parameters
- **eval period** - the time period that must be decomposed (minimum '1h'), defined as <N><time unit> where
`N` - the number of time units
`time unit` - h (hour), d (day), w (week), M (month) or y (year)
- **detection period** - the time period before the end of eval period for which anomalies are calculated (minimum '1h', cannot be longer than eval period), defined as <N><time unit> where
`N` - the number of time units
`time unit` - h (hour), d (day), w (week)

- **season** - the shortest time period where a repeating pattern (”season”) is expected (minimum ’2h’, cannot be longer than eval period, the number of entries in the eval period must be greater than the two times of the resulting frequency (season/h)), defined as <N><time unit> where
N - the number of time units
time unit - h (hour), d (day), w (week)
- **deviations** - the number of deviations (calculated by devalg) to count as anomaly (can be decimal), (must be greater than or equal to 1, default is 3)
- **devalg** (must be double-quoted) - the deviation algorithm, can be *stddevpop*, *stddevsamp* or *mad* (default)
- **s window** - the span (in lags) of the loess window for seasonal extraction (default is 10 * number of entries in eval period + 1)

Examples:

```
trendstl(/host/key,100h:now/h,10h,2h) #analyse the last 100 hours of trend data, find the anomaly rate for
trendstl(/host/key,100h:now/h-10h,100h,2h,2.1,"mad") #analyse the period of 100 hours of trend data, up to
trendstl(/host/key,100d:now/d-1d,10d,1d,4,,10) #analyse 100 days of trend data up to a day ago, find the a
trendstl(/host/key,1M:now/M-1y,1d,2h,,"stddevsamp") #analyse the previous month a year ago, find the anoma
```

trendsum(/host/key,time period:time shift)

The sum of trend values within the defined time period.

Parameters:

- See common parameters

Examples:

```
trendsum(/host/key,1h:now/h) #the sum for the previous hour (e.g., 12:00-13:00)
trendsum(/host/key,1h:now/h-1h) #the sum for two hours ago (11:00-12:00)
trendsum(/host/key,1h:now/h-2h) #the sum for three hours ago (10:00-11:00)
trendsum(/host/key,1M:now/M-1y) #the sum for the previous month a year ago
```

See all supported functions.

Mathematical functions

All functions listed here are supported in:

- Trigger expressions
- Calculated item formulas
- Expression macros

Mathematical functions are supported with float and integer value types, unless stated otherwise.

The functions are listed without additional information. Click the function to see the full details.

Function	Description
abs	The absolute value of a value.
acos	The arccosine of a value as an angle, expressed in radians.
asin	The arcsine of a value as an angle, expressed in radians.
atan	The arctangent of a value as an angle, expressed in radians.
atan2	The arctangent of the ordinate (value) and abscissa coordinates specified as an angle, expressed in radians.
avg	The average value of the referenced item values.
cbrt	The cube root of a value.
ceil	Round the value up to the nearest greater or equal integer.
cos	The cosine of a value, where the value is an angle expressed in radians.
cosh	The hyperbolic cosine of a value.
cot	The cotangent of a value, where the value is an angle expressed in radians.
degrees	Converts a value from radians to degrees.
e	The Euler’s number (2.718281828459045).
exp	The Euler’s number at a power of a value.
expm1	The Euler’s number at a power of a value minus 1.
floor	Round the value down to the nearest smaller or equal integer.
log	The natural logarithm.
log10	The decimal logarithm.
max	The highest value of the referenced item values.
min	The lowest value of the referenced item values.

Function	Description
mod	The division remainder.
pi	The Pi constant (3.14159265358979).
power	The power of a value.
radians	Converts a value from degrees to radians.
rand	Return a random integer value.
round	Round the value to decimal places.
signum	Returns '-1' if a value is negative, '0' if a value is zero, '1' if a value is positive.
sin	The sine of a value, where the value is an angle expressed in radians.
sinh	The hyperbolical sine of a value, where the value is an angle expressed in radians.
sqrt	The square root of a value.
sum	The sum of the referenced item values.
tan	The tangent of a value.
truncate	Truncate the value to decimal places.

Function details

Some general notes on function parameters:

- Function parameters are separated by a comma
- Expressions are accepted as parameters
- Optional function parameters (or parameter parts) are indicated by < >

abs(value)

The absolute value (from 0) of a value.

Parameter:

- **value** - the value to check

For example, the absolute value of either '3' or '-3' will be '3'.

Example:

```
abs(last(/host/key))>10
```

acos(value)

The arccosine of a value as an angle, expressed in radians.

Parameter:

- **value** - the value to check

The value must be between -1 and 1. For example, the arccosine of a value '0.5' will be '2.0943951'.

Example:

```
acos(last(/host/key))
```

asin(value)

The arcsine of a value as an angle, expressed in radians.

Parameter:

- **value** - the value to check

The value must be between -1 and 1. For example, the arcsine of a value '0.5' will be '-0.523598776'.

Example:

```
asin(last(/host/key))
```

atan(value)

The arctangent of a value as an angle, expressed in radians.

Parameter:

- **value** - the value to check

For example, the arctangent of a value '1' will be '0.785398163'.

Example:

```
atan(last(/host/key))
```

atan2(value,abscissa)

The arctangent of the ordinate (value) and abscissa coordinates specified as an angle, expressed in radians.

Parameters:

- **value** - the value to check;
- **abscissa** - the abscissa value.

For example, the arctangent of the ordinate and abscissa coordinates of a value '1' will be '2.21429744'.

Example:

```
atan2(last(/host/key),2)
```

avg(<value1>,<value2>,...)

The average value of the referenced item values.

Parameter:

- **valueX** - the value returned by another function that is working with item history.

Example:

```
avg(avg(/host/key,1h),avg(/host2/key2,1h))
```

cbrt(value)

The cube root of a value.

Parameter:

- **value** - the value to check

For example, the cube root of '64' will be '4', of '63' will be '3.97905721'.

Example:

```
cbrt(last(/host/key))
```

ceil(value)

Round the value up to the nearest greater or equal integer.

Parameter:

- **value** - the value to check

For example, '2.4' will be rounded up to '3'. See also floor().

Example:

```
ceil(last(/host/key))
```

cos(value)

The cosine of a value, where the value is an angle expressed in radians.

Parameter:

- **value** - the value to check

For example, the cosine of a value '1' will be '0.54030230586'.

Example:

```
cos(last(/host/key))
```

cosh(value)

The hyperbolic cosine of a value. Returns the value as a real number, not as scientific notation.

Parameter:

- **value** - the value to check

For example, the hyperbolic cosine of a value '1' will be '1.54308063482'.

Example:

```
cosh(last(/host/key))
```

cot(value)

The cotangent of a value, where the value is an angle expressed in radians.

Parameter:

- **value** - the value to check

For example, the cotangent of a value '1' will be '0.54030230586'.

Example:

```
cot(last(/host/key))
```

degrees(value)

Converts a value from radians to degrees.

Parameter:

- **value** - the value to check

For example, a value '1' converted to degrees will be '57.2957795'.

Example:

```
degrees(last(/host/key))
```

e

The Euler's number (2.718281828459045).

Example:

```
e()
```

exp(value)

The Euler's number at a power of a value.

Parameter:

- **value** - the value to check

For example, Euler's number at a power of a value '2' will be '7.38905609893065'.

Example:

```
exp(last(/host/key))
```

expm1(value)

The Euler's number at a power of a value minus 1.

Parameter:

- **value** - the value to check

For example, Euler's number at a power of a value '2' minus 1 will be '6.38905609893065'.

Example:

```
expm1(last(/host/key))
```

floor(value)

Round the value down to the nearest smaller or equal integer.

Parameter:

- **value** - the value to check

For example, '2.6' will be rounded down to '2'. See also ceil().

Example:

```
floor(last(/host/key))
```

log(value)

The natural logarithm.

Parameter:

- **value** - the value to check

For example, the natural logarithm of a value ’2’ will be ’0.69314718055994529’.

Example:

```
log(last(/host/key))
```

log10(value)

The decimal logarithm.

Parameter:

- **value** - the value to check

For example, the decimal logarithm of a value ’5’ will be ’0.69897000433’.

Example:

```
log10(last(/host/key))
```

max(<value1>,<value2>,...)

The highest value of the referenced item values.

Parameter:

- **valueX** - the value returned by another function that is working with item history.

Example:

```
max(avg(/host/key,1h),avg(/host2/key2,1h))
```

min(<value1>,<value2>,...)

The lowest value of the referenced item values.

Parameter:

- **valueX** - the value returned by another function that is working with item history.

Example:

```
min(avg(/host/key,1h),avg(/host2/key2,1h))
```

mod(value,denominator)

The division remainder.

Parameters:

- **value** - the value to check;
- **denominator** - the division denominator.

For example, division remainder of a value ’5’ with division denominator ’2’ will be ’1’.

Example:

```
mod(last(/host/key),2)
```

pi

The Pi constant (3.14159265358979).

Example:

```
pi()
```

power(value,power value)

The power of a value.

Parameters:

- **value** - the value to check;
- **power value** - the Nth power to use.

For example, the 3rd power of a value ’2’ will be ’8’.

Example:

```
power(last(/host/key),3)
```

radians(value)

Converts a value from degrees to radians.

Parameter:

- **value** - the value to check

For example, a value ’1’ converted to radians will be ’0.0174532925’.

Example:

```
radians(last(/host/key))
```

rand

Return a random integer value. A pseudo-random generated number using time as seed (enough for mathematical purposes, but not cryptography).

Example:

```
rand()
```

round(value,decimal places)

Round the value to decimal places.

Parameters:

- **value** - the value to check;
- **decimal places** - specify decimal places for rounding (0 is also possible).

For example, a value ’2.5482’ rounded to 2 decimal places will be ’2.55’.

Example:

```
round(last(/host/key),2)
```

signum(value)

Returns ’-1’ if a value is negative, ’0’ if a value is zero, ’1’ if a value is positive.

Parameter:

- **value** - the value to check.

Example:

```
signum(last(/host/key))
```

sin(value)

The sine of a value, where the value is an angle expressed in radians.

Parameter:

- **value** - the value to check

For example, the sine of a value ’1’ will be ’0.8414709848’.

Example:

```
sin(last(/host/key))
```

sinh(value)

The hyperbolical sine of a value.

Parameter:

- **value** - the value to check

For example, the hyperbolical sine of a value ’1’ will be ’1.17520119364’.

Example:

```
sinh(last(/host/key))
```

sqrt(value)

The square root of a value.
 This function will fail with a negative value.

Parameter:

- **value** - the value to check

For example, the square root of a value ’3.5’ will be ’1.87082869339’.

Example:

```
sqrt(last(/host/key))
```

sum(<value1>,<value2>,...)

The sum of the referenced item values.

Parameter:

- **valueX** - the value returned by another function that is working with item history.

Example:

```
sum(avg(/host/key,1h),avg(/host2/key2,1h))
```

tan(value)

The tangent of a value.

Parameter:

- **value** - the value to check

For example, the tangent of a value ’1’ will be ’1.55740772465’.

Example:

```
tan(last(/host/key))
```

truncate(value,decimal places)

Truncate the value to decimal places.

Parameters:

- **value** - the value to check;
- **decimal places** - specify decimal places for truncating (0 is also possible).

For example, a value ’2.5482’ truncated to 2 decimal places will be ’2.54’.

Example:

```
truncate(last(/host/key),2)
```

See all supported functions.

Operator functions

All functions listed here are supported in:

- Trigger expressions
- Calculated item formulas
- Expression macros

The functions are listed without additional information. Click the function to see the full details.

Function	Description
between	Check if the value belongs to the given range.
in	Check if the value is equal to at least one of the listed values.

Function details

Some general notes on function parameters:

- Function parameters are separated by a comma
- Expressions are accepted as parameters

between(value,min,max)

Check if the value belongs to the given range.
 Supported value types: *Integer*, *Float*.
 Returns: 1 - in range; 0 - otherwise.

Parameters:

- **value** - the value to check;
- **min** - the minimum value;
- **max** - the maximum value.

Example:

```
between(last(/host/key),1,10)=1 #trigger if the value is between 1 and 10
```

in(value,value1,value2,...valueN)

Check if the value is equal to at least one of the listed values.
 Supported value types: *Integer*, *Float*, *Character*, *Text*, *Log*.
 Returns: 1 - if equal; 0 - otherwise.

Parameters:

- **value** - the value to check;
- **valueX** - listed values (string values must be double-quoted).

The value is compared to the listed values as numbers, if all of these values can be converted to numeric; otherwise, compared as strings.

Example:

```
in(last(/host/key),5,10)=1 #trigger if the last value is equal to 5 or 10
in("text",last(/host/key),last(/host/key,#2))=1 #trigger if "text" is equal to either of the last 2 values
```

See all supported functions.

Predictive functions

All functions listed here are supported in:

- Trigger expressions
- Calculated item formulas
- Expression macros

The functions are listed without additional information. Click the function to see the full details.

Function	Description
forecast	The future value, max, min, delta or avg of the item.
timeleft	The time in seconds needed for an item to reach the specified threshold.

Common parameters

- `/host/key` is a common mandatory first parameter for the functions referencing the host item history.
- `(sec|#num)<:time shift>` is a common second parameter for the functions referencing the host item history, where:
  - **sec** - maximum evaluation period in seconds (time suffixes can be used), or
  - **#num** - maximum evaluation range in latest collected values (if preceded by a hash mark).
  - **time shift** (optional) - allows the evaluation point to be moved back in time. See more details on specifying time shift.

Function details

Some general notes on function parameters:

- Function parameters are separated by a comma
- Optional function parameters (or parameter parts) are indicated by < >
- Function-specific parameters are described with each function
- `/host/key` and `(sec|#num)<:time shift>` parameters must never be quoted

forecast(/host/key,(sec|#num)<:time shift>,time,<fit>,<mode>)

The future value, max, min, delta or avg of the item.
 Supported value types: *Float*, *Integer*.

Parameters:

- See common parameters;

- **time** - the forecasting horizon in seconds (e.g., 1, 60) (time suffixes can be used (e.g., 30s, 1m)); negative values are supported;
- **fit** (optional; must be double-quoted) - the function used to fit historical data. Supported fits:
linear - linear function (default)
polynomialN - polynomial of degree N (1 <= N <= 6)
exponential - exponential function
logarithmic - logarithmic function
power - power function
Note that *polynomial1* is equivalent to *linear*;
- **mode** (optional; must be double-quoted) - the demanded output. Supported modes:
value - value (default)
max - maximum
min - minimum
delta - *max-min*
avg - average
Note that *value* estimates the item value at the moment `now + time`; *max*, *min*, *delta* and *avg* investigate the item value estimate on the interval between `now` and `now + time`.

Comments:

- If the value to return is larger than 1.7976931348623158E+308 or less than -1.7976931348623158E+308, the return value is cropped to 1.7976931348623158E+308 or -1.7976931348623158E+308 correspondingly;
- Becomes unsupported only if misused in the expression (wrong item type, invalid parameters); otherwise, it returns -1 in case of errors;
- See also additional information on predictive trigger functions.

Examples:

```
forecast(/host/key,#10,1h) #forecast the item value in one hour based on the last 10 values
forecast(/host/key,1h,30m) #forecast the item value in 30 minutes based on the last hour data
forecast(/host/key,1h:now-1d,12h) #forecast the item value in 12 hours based on one hour one day ago
forecast(/host/key,1h,10m,"exponential") #forecast the item value in 10 minutes based on the last hour dat
forecast(/host/key,1h,2h,"polynomial3","max") #forecast the maximum value the item can reach in the next t
forecast(/host/key,#2,-20m) #estimate the item value 20 minutes ago based on the last two values (this can
```

timeleft(/host/key,(sec|#num)<:time shift>,threshold,<fit>)

The time in seconds needed for an item to reach the specified threshold.
 Supported value types: *Float*, *Integer*.

Parameters:

- See common parameters;
- **threshold** - the value to reach (unit suffixes can be used);
- **fit** (optional; must be double-quoted) - see forecast().

Comments:

- If the value to return is larger than 1.7976931348623158E+308, the return value is cropped to 1.7976931348623158E+308;
- Returns 1.7976931348623158E+308 if the threshold cannot be reached;
- Becomes unsupported only if misused in the expression (wrong item type, invalid parameters); otherwise,it returns -1 in case of errors;
- See also additional information on predictive trigger functions.

Examples:

```
timeleft(/host/key,#10,0) #the time until the item value reaches zero based on the last 10 values
timeleft(/host/key,1h,100) #the time until the item value reaches 100 based on the last hour data
timeleft(/host/key,1h:now-1d,100) #the time until the item value reaches 100 based on one hour one day ago
timeleft(/host/key,1h,200,"polynomial2") #the time until the item value reaches 200 based on the last hour
```

See all supported functions.

Predictive trigger functions

Overview

Sometimes there are signs of the upcoming problem. These signs can be spotted so that actions may be taken in advance to prevent or at least minimize the impact of the problem.

Zabbix has tools to predict the future behavior of the monitored system based on historic data. These tools are realized through predictive trigger functions.

Functions

Before setting a trigger, it is necessary to define what a problem state is and how much time is needed to take action. Then there are two ways to set up a trigger signaling about a potential unwanted situation. First: the trigger must fire when the system is expected to be in a problem state after ”time to act”. Second: the trigger must fire when the system is going to reach the problem

state in less than "time to act". Corresponding trigger functions to use are **forecast** and **timeleft**. Note that underlying statistical analysis is basically identical for both functions. You may set up a trigger whichever way you prefer with similar results.

Parameters

Both functions use almost the same set of parameters. Use the list of supported functions for reference.

Time interval

First of all, you should specify the historic period Zabbix should analyze to come up with the prediction. You do it in a familiar way by means of the `time period` parameter and optional time shift like you do it with **avg**, **count**, **delta**, **max**, **min** and **sum** functions.

Forecasting horizon

(**forecast** only)
 Parameter `time` specifies how far in the future Zabbix should extrapolate dependencies it finds in historic data. No matter if you use `time_shift` or not, `time` is always counted starting from the current moment.

Threshold to reach

(**timeleft** only)
 Parameter `threshold` specifies a value the analyzed item has to reach, no difference if from above or from below. Once we have determined f(t) (see below), we should solve equation f(t) = `threshold` and return the root which is closer to now and to the right from now or 1.7976931348623158E+308 if there is no such root.

Note:
When item values approach the threshold and then cross it, **timeleft** assumes that intersection is already in the past and therefore switches to the next intersection with `threshold` level, if any. Best practice should be to use predictions as a complement to ordinary problem diagnostics, not as a substitution.[1]

Fit functions

Default `fit` is the *linear* function. But if your monitored system is more complicated you have more options to choose from.

`fit`	x = f(t)
linear	$x = a + b*t$
polynomialN[2]	$x = a_0 + a_1*t + a_2*t^2 + ... + a_n*t^n$
exponential	$x = a*\exp(b*t)$
logarithmic	$x = a + b*\log(t)$
power	$x = a*t^b$

Modes

(**forecast** only)
 Every time a trigger function is evaluated, it gets data from the specified history period and fits a specified function to the data. So, if the data is slightly different, the fitted function will be slightly different. If we simply calculate the value of the fitted function at a specified time in the future, you will know nothing about how the analyzed item is expected to behave between now and that moment in the future. For some `fit` options (like *polynomial*) a simple value from the future may be misleading.

`mode`	**forecast** result
value	f(now + `time`)
max	$\max_{\text{now <= t <= now + time}}$ f(t)
min	$\min_{\text{now <= t <= now + time}}$ f(t)
delta	*max* - *min*
avg	average of f(t) (now <= t <= now + `time`) according to definition

Details

To avoid calculations with huge numbers, we consider the timestamp of the first value in specified period plus 1 ns as a new zero-time (current epoch time is of order 10^9, epoch squared is 10^{18}, double precision is about 10^{-16}). 1 ns is added to provide all positive time values for *logarithmic* and *power* fits which involve calculating log(t). Time shift does not affect *linear*, *polynomial*, *exponential* (apart from easier and more precise calculations) but changes the shape of *logarithmic* and *power* functions.

Potential errors

Functions return -1 in such situations:

- specified evaluation period contains no data;

- result of mathematical operation is not defined[3];
- numerical complications (unfortunately, for some sets of input data range and precision of double-precision floating-point format become insufficient)[4].

Note:
No warnings or errors are flagged if chosen fit poorly describes provided data or there is just too few data for accurate prediction.

Examples and dealing with errors

To get a warning when you are about to run out of free disk space on your host, you may use a trigger expression like this:

```
timeleft(/host/vfs.fs.size[/,free],1h,0)<1h
```

However, error code -1 may come into play and put your trigger in a problem state. Generally it's good because you get a warning that your predictions don't work correctly and you should look at them more thoroughly to find out why. But sometimes it's bad because -1 can simply mean that there was no data about the host free disk space obtained in the last hour. If you are getting too many false positive alerts, consider using more complicated trigger expression [5]:

```
timeleft(/host/vfs.fs.size[/,free],1h,0)<1h and timeleft(/host/vfs.fs.size[/,free],1h,0)<>-1
```

The situation is a bit more difficult with **forecast**. First of all, -1 may or may not put the trigger in a problem state depending on whether you have expression like `forecast(/host/item,(...))<...` or like `forecast(/host/item,(...))>...`

Furthermore, -1 may be a valid forecast if it's normal for the item value to be negative. But probability of this situation in the real world situation is negligible (see how the operator **=** works). So add `... or forecast(/host/item,(...))=-1` or `... and forecast(/host/item,(...))<>-1` if you want or don't want to treat -1 as a problem respectively.

Footnotes

[1] For example, a simple trigger like `timeleft(/host/item,1h,X) < 1h` may go into problem state when the item value approaches X and then suddenly recover once value X is reached. If the problem is item value being below X, use: `last(/host/item) < X or timeleft(/host/item,1h,X) < 1h` If the problem is item value being above X use: `last(/host/item) > X or timeleft(/host/item,1h,X) < 1h`

[2] Polynomial degree can be from 1 to 6, *polynomial1* is equivalent to *linear*. However, use higher degree polynomials with caution. If the evaluation period contains less points than needed to determine polynomial coefficients, polynomial degree will be lowered (e.g., *polynomial5* is requested, but there are only 4 points, therefore *polynomial3* will be fitted).

[3] For example, fitting *exponential* or *power* functions involves calculating log() of item values. If data contains zeros or negative numbers, you will get an error since log() is defined for positive values only.

[4] For *linear*, *exponential*, *logarithmic* and *power* fits all necessary calculations can be written explicitly. For *polynomial* only *value* can be calculated without any additional steps. Calculating *avg* involves computing polynomial antiderivative (analytically). Computing *max*, *min* and *delta* involves computing polynomial derivative (analytically) and finding its roots (numerically). Solving f(t) = 0 involves finding polynomial roots (numerically).

[5] But in this case -1 can cause your trigger to recover from the problem state. To be fully protected use: `timeleft(/host/vfs.fs.size[/,f and ({TRIGGER.VALUE}=0 and timeleft(/host/vfs.fs.size[/,free],1h,0)<>-1 or {TRIGGER.VALUE}=1)`

String functions

All functions listed here are supported in:

- Trigger expressions
- Calculated item formulas
- Expression macros

The functions are listed without additional information. Click the function to see the full details.

Function	Description
ascii	The ASCII code of the leftmost character of the value.
bitlength	The length of value in bits.
bytelength	The length of value in bytes.
char	Return the character by interpreting the value as ASCII code.
concat	The string resulting from concatenating the referenced item values or constant values.
contains	Return 1 if the value contains the substring specified by `search value`, otherwise 0.
insert	Insert specified characters or spaces into the character string beginning at the specified position in the string.

Function	Description
jsonpath	Return the JSONPath result.
left	Return the leftmost characters of the value.
length	The length of value in characters.
ltrim	Remove specified characters from the beginning of string.
mid	Return a substring of N characters beginning at the character position specified by 'start'. This function is an alias for the `substring()` function.
repeat	Repeat a string.
replace	Find the pattern in the value and replace with replacement.
right	Return the rightmost characters of the value.
rtrim	Remove specified characters from the end of string.
substring	Return a substring of N characters beginning at the character position specified by 'start'.
trim	Remove specified characters from the beginning and end of string.
xmlxpath	Return the XML XPath result.

Function details

Some general notes on function parameters:

- Function parameters are separated by a comma
- Expressions are accepted as parameters
- String parameters must be double-quoted; otherwise, they might get misinterpreted
- Optional function parameters (or parameter parts) are indicated by < >

ascii(value)

The ASCII code of the leftmost character of the value.
 Supported value types: *String*, *Text*, *Log*.

Parameter:

- **value** - the value to check

For example, a value like 'Abc' will return '65' (ASCII code for 'A').

Example:

```
ascii(last(/host/key))
```

bitlength(value)

The length of value in bits.
 Supported value types: *String*, *Text*, *Log*, *Integer*.

Parameter:

- **value** - the value to check

Example:

```
bitlength(last(/host/key))
```

bytelength(value)

The length of value in bytes.
 Supported value types: *String*, *Text*, *Log*, *Integer*.

Parameter:

- **value** - the value to check

Example:

```
bytelength(last(/host/key))
```

char(value)

Return the character by interpreting the value as ASCII code.
 Supported value types: *Integer*.

Parameter:

- **value** - the value to check

The value must be in the 0-255 range. For example, a value like '65' (interpreted as ASCII code) will return 'A'.

Example:

```
char(last(/host/key))
```

concat(<value1>,<value2>,...)

The string resulting from concatenating the referenced item values or constant values.
 Supported value types: *String*, *Text*, *Log*, *Float*, *Integer*.

Parameter:

- **valueX** - the value returned by one of the history functions or a constant value (string, integer, or float number). Must contain at least two parameters.

For example, a value like 'Zab' concatenated to 'bix' (the constant string) will return 'Zabbix'.

Examples:

```
concat(last(/host/key),"bix")
concat("1 min: ",last(/host/system.cpu.load[all,avg1]),", 15 min: ",last(/host/system.cpu.load[all,avg15])
```

contains(value, search value)

Return 1 if the value contains the substring specified by `search value`, otherwise 0.| Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **search value** - the value to search.

Comments:

- The search is case-sensitive.
- When the value is valid and `search value` is empty, the function returns 1.

For example, the function returns 1 if the value contains 'Zabbix'.

Example:

```
contains(last(/host/key),"Zabbix")
```

insert(value,start,length,replacement)

Insert specified characters or spaces into the character string beginning at the specified position in the string.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **start** - start position;
- **length** - positions to replace;
- **replacement** - replacement string.

For example, a value like 'Zabbbix' will be replaced by 'Zabbix' if 'bb' (starting position 3, positions to replace 2) is replaced by 'b'.

Example:

```
insert(last(/host/key),3,2,"b")
```

jsonpath(value,path,<default>)

Return the JSONPath result.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **path** - the path (must be quoted);
- **default** - the optional fallback value if the JSONPath query returns no data. Note that on other errors failure is returned (e.g., "unsupported construct").

Example:

```
jsonpath(last(/host/proc.get[zabbix_agentd,,,summary]),"$..size")
```

left(value,count)

Return the leftmost characters of the value.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **count** - the number of characters to return.

For example, you may return 'Zab' from 'Zabbix' by specifying 3 leftmost characters to return. See also right().

Example:

```
left(last(/host/key),3) #return three leftmost characters
```

length(value)

The length of value in characters.
 Supported value types: *String*, *Text*, *Log*.

Parameter:

- **value** - the value to check.

Examples:

```
length(last(/host/key)) #the length of the latest value
length(last(/host/key,#3)) #the length of the third most recent value
length(last(/host/key,#1:now-1d)) #the length of the most recent value one day ago
```

ltrim(value,<chars>)

Remove specified characters from the beginning of string.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **chars** (optional) - specify the characters to remove.

Whitespace is left-trimmed by default (if no optional characters are specified). See also: rtrim(), trim().

Examples:

```
ltrim(last(/host/key)) #remove whitespace from the beginning of string
ltrim(last(/host/key),"Z") #remove any 'Z' from the beginning of string
ltrim(last(/host/key)," Z") #remove any space and 'Z' from the beginning of string
```

mid(value,start,length)

Return a substring of N characters beginning at the character position specified by 'start'.
This function is an alias for the `substring()` function.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **start** - start position of the substring;
- **length** - positions to return in substring.

For example, it is possible return 'abbi' from a value like 'Zabbix' if starting position is 2, and positions to return is 4.

Example:

```
mid(last(/host/key),2,4)="abbi"
```

repeat(value,count)

Repeat a string.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **count** - the number of times to repeat.

Example:

```
repeat(last(/host/key),2) #repeat the value two times
```

replace(value,pattern,replacement)

Find the pattern in the value and replace with replacement. All occurrences of the pattern will be replaced.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **pattern** - the pattern to find;
- **replacement** - the string to replace the pattern with.

Example:

```
replace(last(/host/key),"ibb","abb") #replace all 'ibb' with 'abb'
```

right(value,count)

Return the rightmost characters of the value.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **count** - the number of characters to return.

For example, you may return 'bix' from 'Zabbix' by specifying 3 rightmost characters to return. See also left().

Example:

```
right(last(/host/key),3) #return three rightmost characters
```

rtrim(value,<chars>)

Remove specified characters from the end of string.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **chars** (optional) - specify the characters to remove.

Whitespace is right-trimmed by default (if no optional characters are specified). See also: ltrim(), trim().

Examples:

```
rtrim(last(/host/key)) #remove whitespace from the end of string
rtrim(last(/host/key),"x") #remove any 'x' from the end of string
rtrim(last(/host/key),"x ") #remove any 'x' and space from the end of string
```

substring(value,start,length)

Return a substring of N characters beginning at the character position specified by 'start'.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **start** - start position of the substring;
- **length** - positions to return in substring.

For example, it is possible return 'abbi' from a value like 'Zabbix' if starting position is 2, and positions to return is 4.

Example:

```
substring(last(/host/key),2,4)="abbi"
```

trim(value,<chars>)

Remove specified characters from the beginning and end of string.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **chars** (optional) - specify the characters to remove.

Whitespace is trimmed from both sides by default (if no optional characters are specified). See also: ltrim(), rtrim().

Examples:

```
trim(last(/host/key)) #remove whitespace from the beginning and end of string
trim(last(/host/key),"_") #remove '_' from the beginning and end of string
```

xmlxpath(value,path,<default>)

Return the XML XPath result.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **value** - the value to check;
- **path** - the path (must be quoted);
- **default** - the optional fallback value if the XML XPath query returns an empty nodeset. It will not be returned if the empty result is not a nodeset (i.e., empty string). On other errors failure is returned (e.g., "invalid expression").

Example:

`xmlxpath(last(/host/xml_result),"/response/error/status")`

See all supported functions.

Best practices for triggers

Triggers are a powerful tool, but they can also create undesired alert noise. To see more real signals, and less noise, follow these tips:

1. Desensitize triggers. Instead of alerting on the latest value (too high/too low), analyze the average for an extended period by using functions like **avg**, **min**, and **max**.
2. Consider using the **percentile** function (set it to 95% or 5%) in triggers, if you want to avoid alerts on random spikes and drops.
3. Use hysteresis to avoid trigger flapping—frequent changes of trigger state (OK ↔ Problem). A continuum for the problem state can be defined by adding a recovery expression (a separate condition for problem resolution).
4. Use trigger dependencies to avoid alerts, which are not related to the root cause.
5. Use trigger severity to alert on the more serious problems only.
6. Define maintenance windows.

Hysteresis

Sometimes an interval is needed between problem and recovery states, rather than a simple threshold. For example, if we want to define a trigger that reports a problem when server room temperature goes above 20°C and we want it to stay in the problem state until the temperature drops below 15°C, a simple trigger threshold at 20°C will not be enough.

Instead, we need to define a trigger expression for the problem event first (temperature above 20°C). Then we need to define an additional recovery condition (temperature below 15°C). This is done by defining a *Recovery expression* when configuring a trigger.

In this case, problem recovery takes place in two steps:

- First, the problem expression (temperature above 20°C) has to evaluate to FALSE
- Second, the recovery expression (temperature below 15°C) has to evaluate to TRUE

The recovery expression is evaluated only after the problem event is resolved. The recovery expression being TRUE alone does not resolve a problem if the problem expression is still TRUE!

Example

Temperature in server room is too high.

Problem expression:

```
last(/server/temp)>20
```

Recovery expression:

```
last(/server/temp)<=15
```

> **Note:**
> It is unproductive to use the {TRIGGER.VALUE} macro in a recovery expression because this expression is only evaluated when the trigger is in the ”Problem” state. Consequently, {TRIGGER.VALUE} will always resolve to ”1” (which indicates a ”Problem” state) while evaluating the expression.

Trigger dependencies

Trigger dependencies can be used to avoid alerts, which are not related to the root cause.

See all best practices.

Overview

Sometimes the availability of one host depends on another. A server that is behind a router will become unreachable if the router goes down. With triggers configured for both, you might get notifications about two hosts down - while only the router was the guilty party.

This is where some dependency between hosts might be useful. With dependency set, notifications of the dependents could be withheld and only the notification on the root problem sent.

While Zabbix does not support dependencies between hosts directly, they may be defined with another, more flexible method - trigger dependencies. A trigger may have one or more triggers it depends on.

So in our simple example we open the server trigger configuration form and set that it depends on the respective trigger of the router. With such dependency, the server trigger will not change its state as long as the trigger it depends on is in the 'PROBLEM' state - and thus no dependent actions will be taken and no notifications sent.

If both the server and the router are down, and a dependency between the "server down" trigger and the "router down" trigger is set, Zabbix will not execute actions for the dependent trigger.

While the parent trigger is in the PROBLEM state, its dependents may report values that cannot be trusted. Therefore dependent triggers will not be re-evaluated until the parent trigger (the router in the example above):

- goes back from 'PROBLEM' to 'OK' state;
- changes its state from 'PROBLEM' to 'UNKNOWN';
- is closed manually, by correlation or with the help of date and time and/or nodata() functions;
- is resolved by a value of an item not involved in the dependent trigger;
- is disabled, has a disabled item or a disabled item host

In all of the cases mentioned above, the dependent trigger (server) will be re-evaluated only when a new metric for it is received. This means that the dependent trigger may not be updated immediately.

Also:

- Trigger dependency may be added from any host trigger to any other host trigger, as long as it doesn't result in a circular dependency.
- Trigger dependency may be added from one template to another. If some trigger from template A depends on some trigger from template B, template A may only be linked to a host (or another template) together with template B, but template B may be linked to a host (or another template) alone.
- Trigger dependency may be added from a template trigger to a host trigger. In this case, linking such a template to a host will create a host trigger that depends on the same trigger template that the trigger was depending on. This allows you, for example, to have a template in which some triggers depend on the router (host) triggers. All hosts linked to this template will depend on that specific router.
- Trigger dependency may not be added from a host trigger to a template trigger.
- Trigger dependency may be added from a trigger prototype to another trigger prototype (within the same low-level discovery rule) or a real trigger. A trigger prototype may not depend on a trigger prototype from a different LLD rule or on a trigger created from trigger prototype. A host trigger prototype cannot depend on a trigger from a template.

Configuration

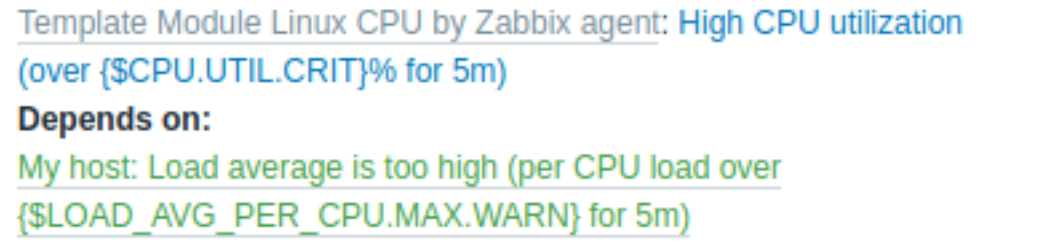
To define a dependency:

1. Open the *Dependencies* tab in the trigger configuration form.
2. Click *Add* in the *Dependencies* section and select one or more triggers that the trigger will depend on.



3. Click *Update*.

Now the trigger has the indication of its dependency in the list.



Example of several dependencies

For example, the Host is behind the Router2 and the Router2 is behind the Router1.

```
Zabbix - Router1 - Router2 - Host
```

If the Router1 is down, then obviously the Host and the Router2 are also unreachable, yet receiving three notifications about the Host, the Router1 and the Router2 all being down is excessive.

So in this case we define two dependencies:

```
the 'Host is down' trigger depends on the 'Router2 is down' trigger
the 'Router2 is down' trigger depends on the 'Router1 is down' trigger
```

Before changing the status of the 'Host is down' trigger, Zabbix will check for the corresponding trigger dependencies. If such are found and one of those triggers is in the **Problem** state, then the trigger state will not be changed, the actions will not be executed and no notifications will be sent.

Zabbix performs this check recursively. If the Router1 or the Router2 is unreachable, the Host trigger won't be updated.

Trigger severity

Trigger severity can be used to avoid alert noise on less serious issues.

See all best practices.

Overview

Severity represents the relative importance of a trigger.



Assigning severity is not only useful for visual representation, but also for avoiding notifications on trivial issues.

Trigger severities are used for:

- Visual representation (distinct color per each created problem, based on the original trigger severity)
- Limiting notifications based on severity (see action conditions)
- User media - different media (notification channel) for different severities (for example, SMS for triggers of *High* and *Disaster* severity, and Email for triggers of other severities)
- Audio in global alarms - different audio for different severities

The following trigger severities are supported by default:

Severity	Color	Description
Not classified	Gray	Can be used where the severity level of an event is unknown, has not been determined, is not part of the regular monitoring scope, etc., for example, during initial configuration, as a placeholder for future assessment, or as part of an integration process.
Information	Light blue	Can be used for informational events that do not require immediate attention, but can still provide valuable insights.
Warning	Yellow	Can be used to indicate a potential issue that might require investigation or action, but that is not critical.
Average	Orange	Can be used to indicate a significant issue that should be addressed relatively soon to prevent further problems.
High	Light red	Can be used to indicate critical issues that need immediate attention to avoid significant disruptions.
Disaster	Red	Can be used to indicate a severe incident that requires immediate action to prevent, for example, system outages or data loss.

Trigger severity names and colors can be **customized** in *Administration > General > Trigger displaying options*. Colors are shared among all frontend themes.

Translating customized severity names

> **Attention:**
> If Zabbix frontend translations are used, custom severity names will override translated names by default.

Default trigger severity names are available for translation in all locales. If a severity name is changed, a custom name is used in all locales and additional manual translation is needed.

Custom severity name translation procedure:

- set required custom severity name, for example, 'Important'
- edit <frontend_dir>/locale/<required_locale>/LC_MESSAGES/frontend.po
- add 2 lines:

```
msgid "Important"
msgstr "<translation string>"
```

and save file.

- create .mo files as described in <frontend_dir>/locale/README

Here **msgid** should match the new custom severity name and **msgstr** should be the translation for it in the specific language.

This procedure should be performed after each severity name change.

Visualization

Overview

Visualization provides different ways to display monitoring data collected by Zabbix. Graphs show trends and metrics, network maps display infrastructure topology, and dashboards combine customizable widgets to provide comprehensive monitoring views.

Visualization tools

- Graphs - view item data as simple auto-generated graphs, create custom graphs, or compare multiple items in ad-hoc graphs.
- Network maps - create visual representations of network infrastructure with customizable elements and indicators.
- Dashboards - customize layouts combining a variety of dashboard widgets that visualize data in different formats.

Graphs

Overview

With lots of data flowing into Zabbix, it becomes much easier for the users if they can look at a visual representation of what is going on rather than only numbers.

This is where graphs come in. Graphs make it possible to grasp the data flow at a glance, correlate problems, discover when something started, or illustrate when something might turn into a problem.

Zabbix provides users with:

- Built-in simple graphs of one item data.
- The possibility to create more complex customized graphs.
- Access to a comparison of several items quickly in ad-hoc graphs.
- Modern customizable vector graphs and pie charts.

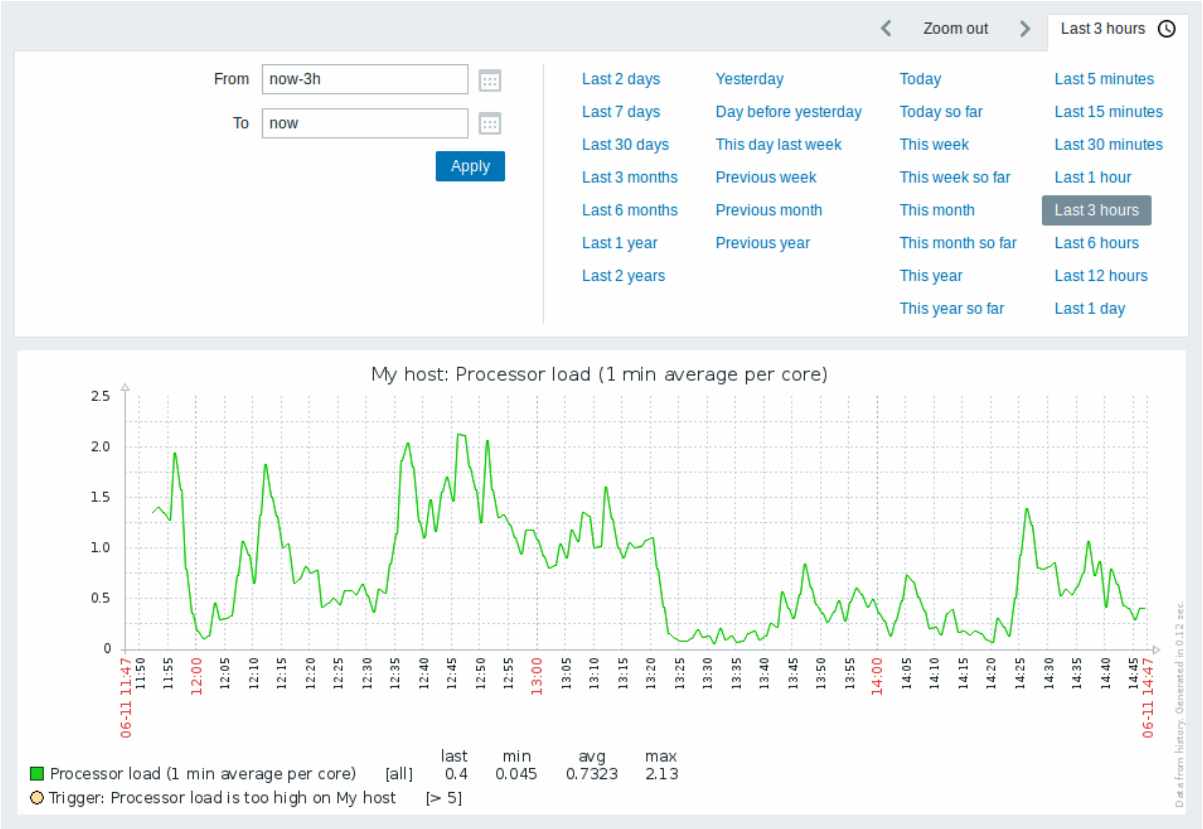
Simple graphs

Overview

Simple graphs are provided for the visualization of data gathered by items.

No configuration effort is required on the user part to view simple graphs. They are freely made available by Zabbix.

Just go to *Monitoring* > *Latest data* and click the *Graph* link for the respective item and a graph will be displayed.



Note:
Simple graphs are provided for all numeric items. For textual items, a link to *History* is available in *Monitoring* > *Latest data*.

Time period selector

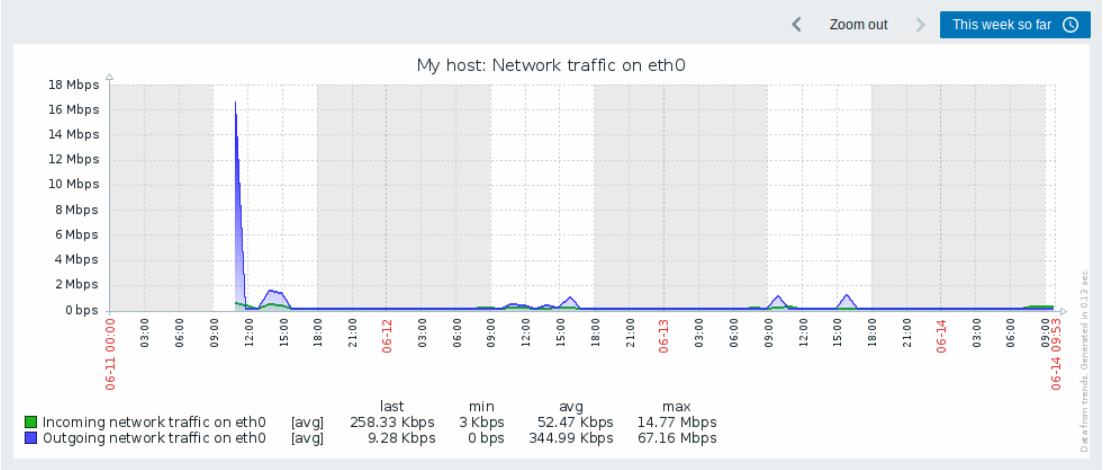
The time period selector above the graph allows you to select often required periods with one mouse click. For more information, see Time period and host selectors.

Recent data vs longer periods

For very recent data a **single** line is drawn connecting each received value. The single line is drawn as long as there is at least one horizontal pixel available for one value.

For data that show a longer period **three lines** are drawn - a dark green one shows the average, while a light pink and a light green line shows the maximum and minimum values at that point in time. The space between the highs and the lows is filled with yellow background.

Working time (working days) is displayed in graphs as a white background, while non-working time is displayed in gray (with the *Original blue* default frontend theme).

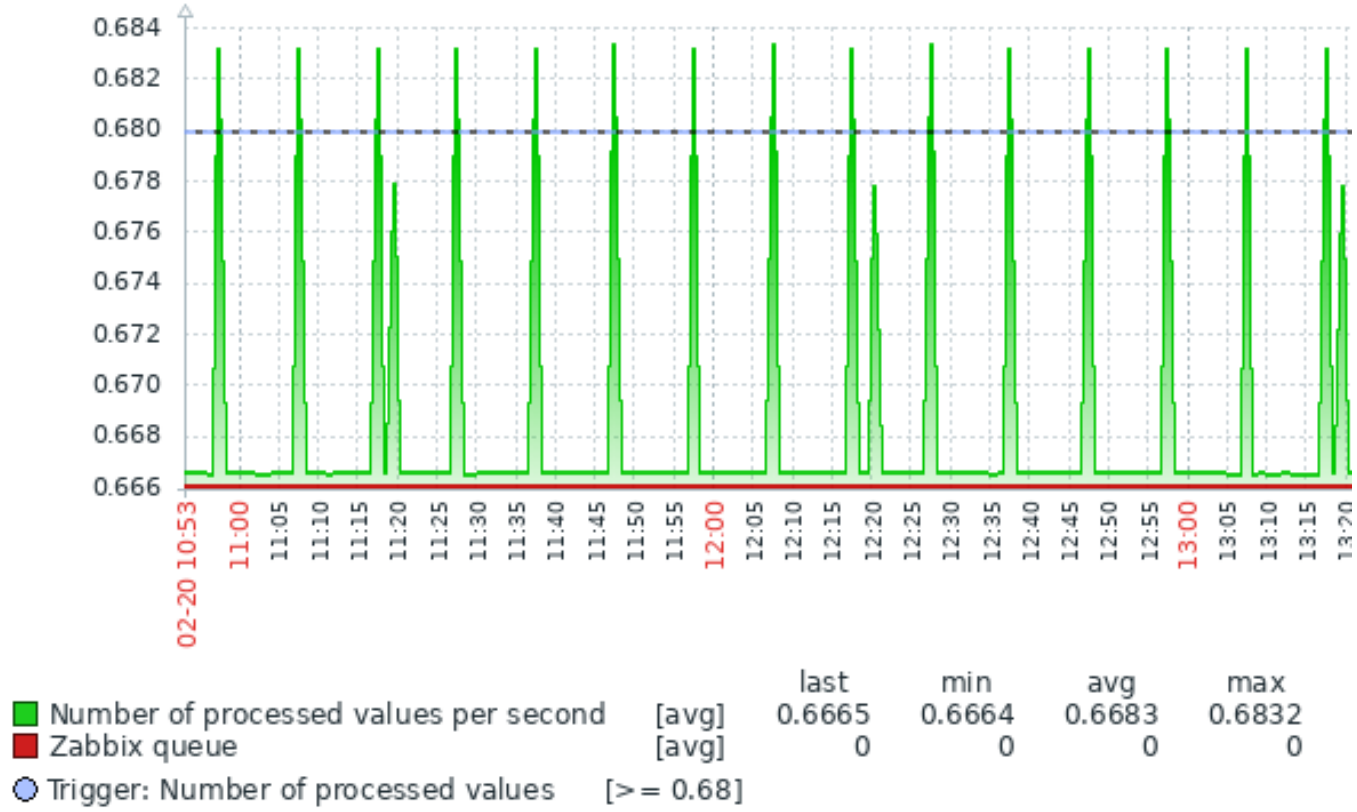


Working time is always displayed in simple graphs, whereas displaying it in custom graphs is a user preference.

Working time is not displayed if the graph shows more than 3 months.

Trigger lines

Simple triggers are displayed as lines with black dashes over trigger severity color - take note of the blue line on the graph and the trigger information displayed in the legend. Up to 3 trigger lines can be displayed on the graph; if there are more triggers then the triggers with lower severity are prioritized. Triggers are always displayed in simple graphs, whereas displaying them in custom graphs is a user preference.



Generating from history/trends

Graphs can be drawn based on either item history or trends.

For the users who have frontend debug mode activated, a gray, vertical caption is displayed at the lower-right of a graph indicating where the data come from.

Several factors influence whether history of trends is used:

- Longevity of item history. For example, item history can be kept for 14 days. In that case, any data older than the fourteen days will be coming from trends.
- Data congestion in the graph. If the amount of seconds to display in a horizontal graph pixel exceeds 3600/16, trend data is displayed (even if item history is still available for the same period).
- If trends are disabled, item history is used for graph building - if available for that period.

Absence of data

For items with a regular update interval, nothing is displayed in the graph if item data is not collected.

However, for trapper items and items with a scheduled update interval (and regular update interval set to 0), a straight line is drawn leading up to the first collected value and from the last collected value to the end of graph; the line is on the level of the first/last value respectively.

Switching to raw values

A dropdown on the upper right allows you to switch from the simple graph to the *Values/500 latest values* listings. This can be useful for viewing the numeric values making up the graph.

The values represented here are raw, i.e. no units or postprocessing of values is used. Value mapping, however, is applied.

Known issues

See known issues for graphs.

Custom graphs

Overview

Custom graphs, as the name suggests, offer customization capabilities.

While simple graphs are good for viewing data of a single item, they do not offer configuration capabilities.

Thus, if you want to change graph style or the way lines are displayed or compare several items, for example, incoming and outgoing traffic in a single graph, you need a custom graph.

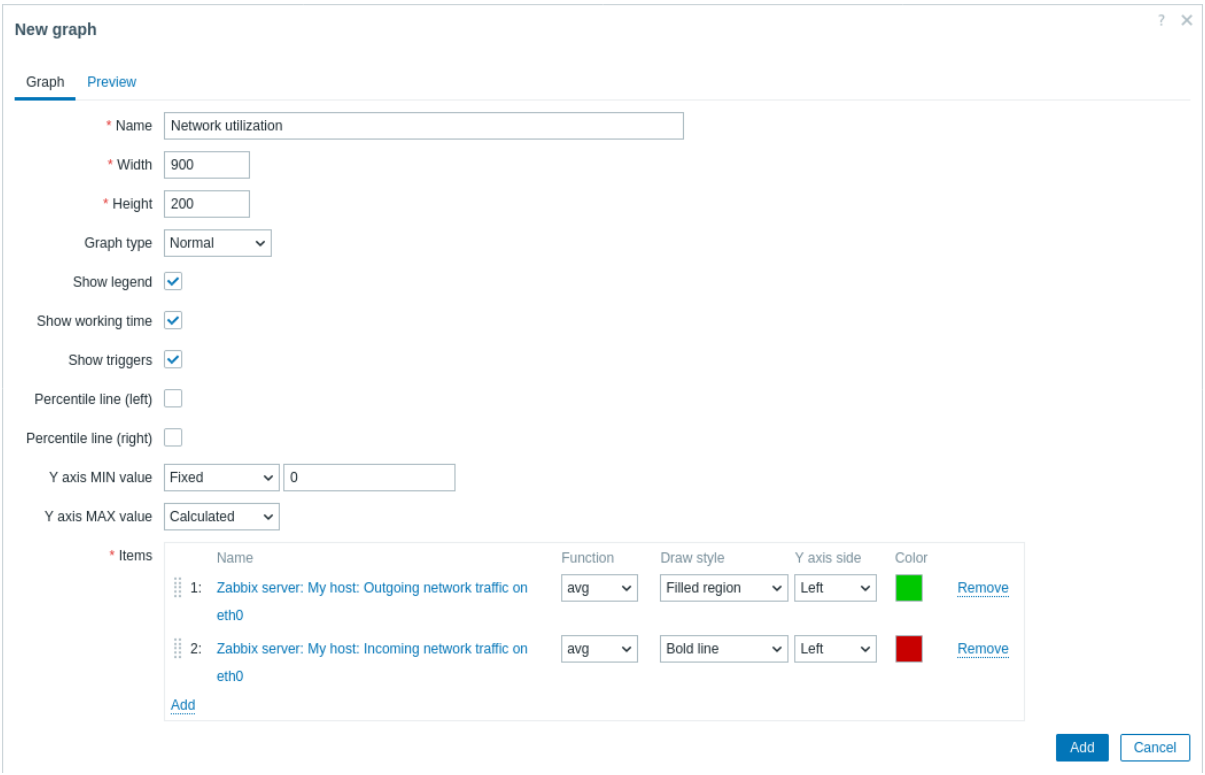
Custom graphs are configured manually.

They can be created for a host or several hosts or for a single template.

Configuring custom graphs

To create a custom graph, do the following:

1. Go to *Data collection* > *Hosts* (or *Templates*).
2. Click *Graphs* in the row next to the desired host or template.
3. In the Graphs screen, click *Create graph*.
4. Edit graph attributes.



All mandatory input fields are marked with a red asterisk.

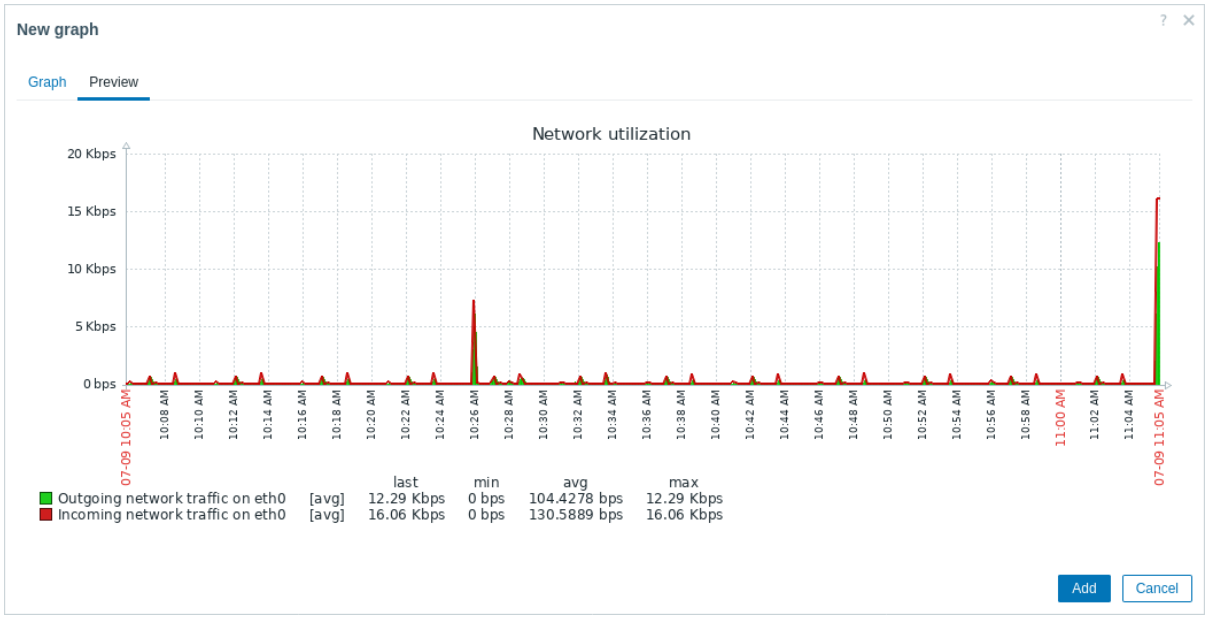
Graph attributes:

Parameter	Description
Name	Unique graph name. Expression macros are supported in this field, but only with `avg`, `last`, `min` and `max` functions, with time as parameter (for example, `{?avg(/host/key,1h)}`). {HOST.HOST<1-9>} macros are supported for the use within this macro, referencing the first, second, third, etc. host in the graph, for example, `{?avg(/{HOST.HOST2}/key,1h)}`. Note that referencing the first host with this macro is redundant, as the first host can be referenced implicitly, for example, `{?avg(//key,1h)}`.
Width	Graph width in pixels (for preview and pie/exploded graphs only).
Height	Graph height in pixels.
Graph type	Graph type: **Normal** - normal graph, values displayed as lines **Stacked** - stacked graph, filled areas displayed **Pie** - pie graph **Exploded** - ”exploded” pie graph, portions displayed as ”cut out” of the pie

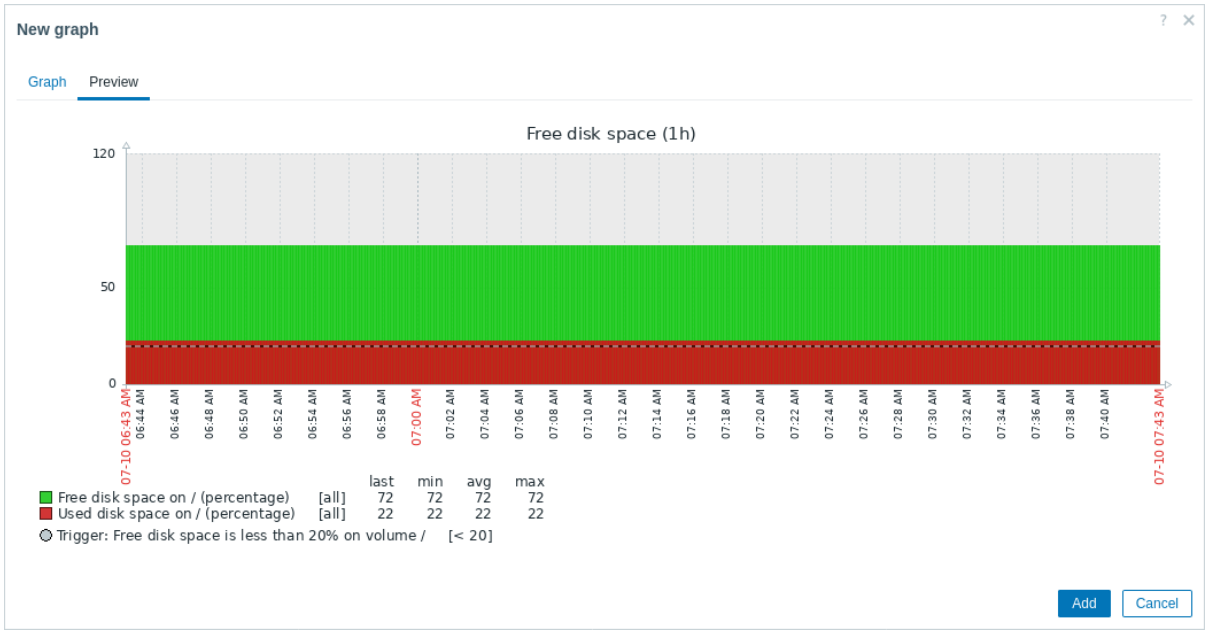
Parameter	Description
Show legend	Checking this box will set to display the graph legend.
Show working time	If selected, non-working hours will be shown with a gray background. This parameter is not available for pie and exploded pie graphs.
Show triggers	If selected, simple triggers will be displayed as lines with black dashes over trigger severity color. This parameter is not available for pie and exploded pie graphs.
Percentile line (left)	Display percentile for left Y-axis. If, for example, 95% percentile is set, then the percentile line will be at the level where 95 percent of the values fall under. Displayed as a bright green line. Only available for normal graphs.
Percentile line (right)	Display percentile for right Y-axis. If, for example, 95% percentile is set, then the percentile line will be at the level where 95 percent of the values fall under. Displayed as a bright red line. Only available for normal graphs.
Y axis MIN value	Minimum value of Y-axis: **Calculated** - Y axis minimum value will be automatically calculated. **Fixed** - fixed minimum value for Y-axis. **Item** - last value of the selected item will be the minimum value. This parameter is not available for pie and exploded pie graphs.
Y axis MAX value	Maximum value of Y-axis: **Calculated** - Y axis maximum value will be automatically calculated. **Fixed** - fixed maximum value for Y-axis. **Item** - last value of the selected item will be the maximum value This parameter is not available for pie and exploded pie graphs.
3D view	Enable 3D style. For pie and exploded pie graphs only.
Items	Items, data of which are to be displayed in this graph. Click *Add* to select items. You can also select various displaying options (function, draw style, left/right axis display, color).
Sort order (0→100)	Draw order. 0 will be processed first. Can be used to draw lines or regions behind (or in front of) another. You can drag and drop items using the icon at the beginning of a line to set the sort order or which item is displayed in front of the other.
Name	Name of the selected item is displayed as a link. Clicking the link opens the list of other available items.
Type	Type (only available for pie and exploded pie graphs): **Simple** - the value of the item is represented proportionally on the pie **Graph sum** - the value of the item represents the whole pie Note that coloring of the ”graph sum” item will only be visible to the extent that it is not taken up by ”proportional” items.
Function	Select what values will be displayed when more than one value exists per vertical graph pixel for an item: **all** - display all possible values (minimum, maximum, average) in the graph. Note that for shorter periods this setting has no effect; only for longer periods, when data congestion in a vertical graph pixel increases, ’all’ starts displaying minimum, maximum, and average values. This function is only available for *Normal* graph type. See also: Generating graphs from history/trends. **avg** - display the average values **last** - display the latest values. This function is only available if either *Pie/Exploded pie* is selected as graph type. **max** - display the maximum values **min** - display the minimum values
Draw style	Select the draw style (only available for normal graphs; for stacked graphs filled region is always used) to apply to the item data - *Line*, *Bold line*, *Filled region*, *Dot*, *Dashed line*, *Gradient line*.
Y axis side	Select the Y axis side to show the item data - *Left*, *Right*.
Color	Select the color to apply to the item data.

Graph preview

In the *Preview* tab, a preview of the graph is displayed so you can immediately see what you are creating.



Note that the preview will not show any data for template items.



In this example, pay attention to the dashed bold line displaying the trigger level and the trigger information displayed in the legend.

Note:
No more than 3 trigger lines can be displayed. If there are more triggers then the triggers with lower severity are prioritized for display.

 If graph height is set as less than 120 pixels, no trigger will be displayed in the legend.

Ad-hoc graphs

Overview

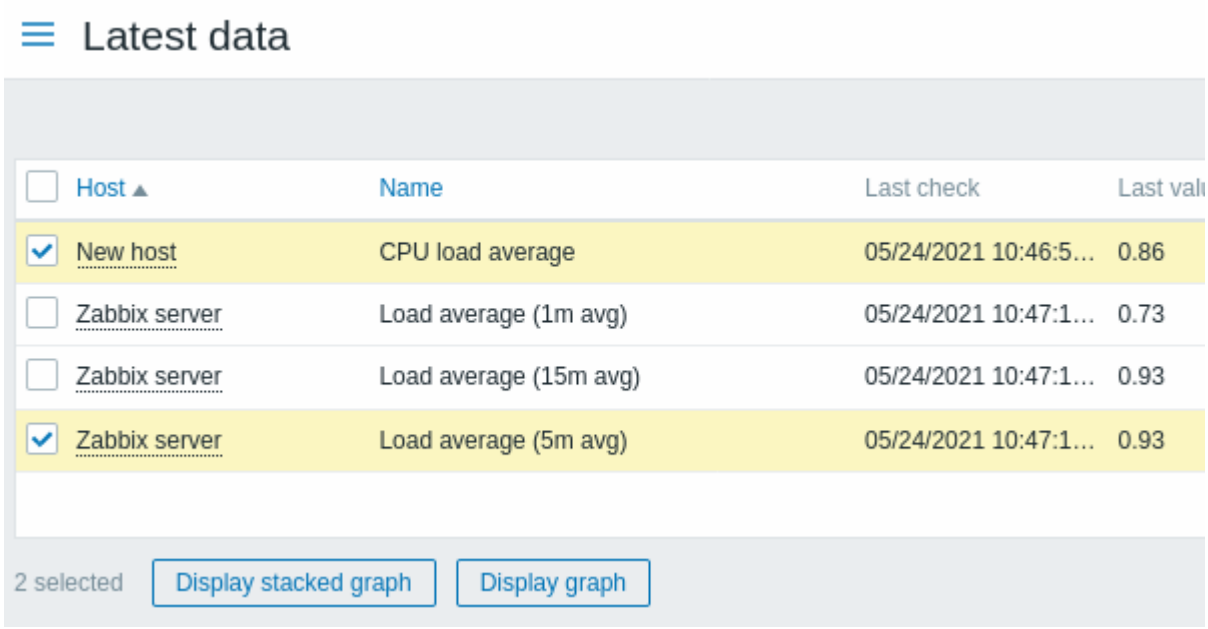
While a simple graph is useful for viewing data from a single item and custom graphs offer customization options, neither allows you to quickly create a comparison graph for multiple items with minimal effort and maintenance.

To address this issue, it is possible to create ad-hoc graphs for several items in a very quick way.

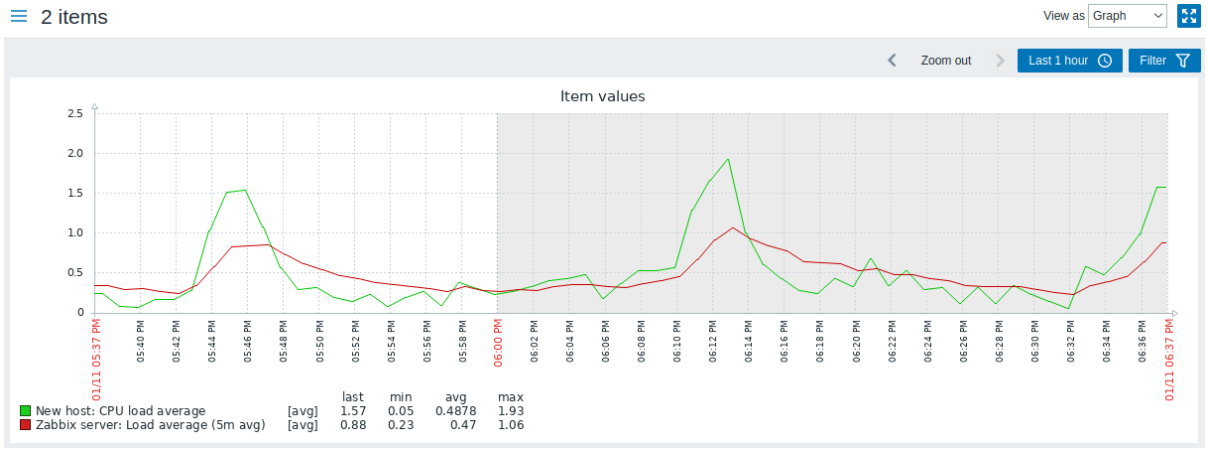
Configuration

To create an ad-hoc graph, do the following:

1. Go to *Monitoring* > *Latest data*.
2. Use filter to display items that you want.
3. Mark checkboxes of the items you want to graph.
4. Click *Display stacked graph* or *Display graph* buttons.

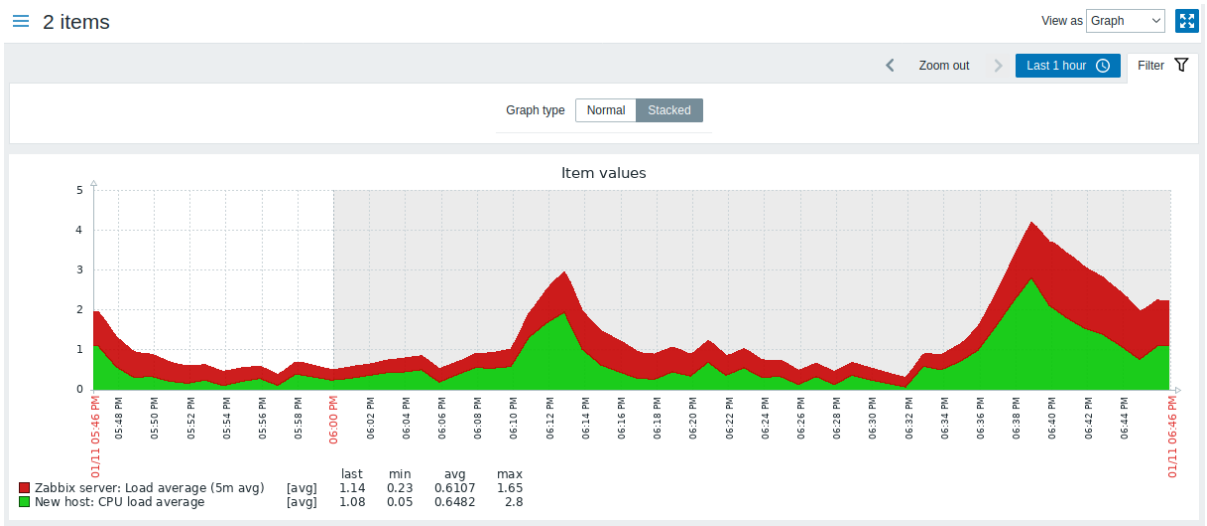


Your graph is created instantly:



Note that to avoid displaying too many lines in the graph, only the average value for each item is displayed (min/max value lines are not displayed). Triggers and trigger information is not displayed in the graph.

In the created graph window you have the time period selector available and the possibility to switch from the ”normal” line graph to a stacked one (and back).



Network maps

Overview

If you have a network to look after, you may want to have an overview of your infrastructure somewhere. For that purpose, you can create maps in Zabbix - of networks and of anything you like.

All users can create network maps. The maps can be public (available to all users) or private (available to selected users).

Proceed to configuring a network map.

Configuring a network map

Overview

Configuring a map in Zabbix requires that you first create a map by defining its general parameters and then you start filling the actual map with elements and their links.

You can populate the map with elements that are a host, a host group, a trigger, an image, or another map.

Icons are used to represent map elements. You can define the information that will be displayed with the icons and set that recent problems are displayed in a special way. You can link the icons and define information to be displayed on the links.

You can add custom URLs to be accessible by clicking the icons. Thus you may link a host icon to host properties or a map icon to another map.

The problem count in maps is displayed for cause problems only.

Maps are managed in *Monitoring* > *Maps*, where they can be configured, managed, and viewed. In the monitoring view, you can click the icons and take advantage of the links to some scripts and URLs.

Network maps are based on vector graphics (SVG).

Public and private maps

All users in Zabbix (including non-admin users) can create network maps. Maps have an owner - the user who created them. Maps can be made public or private.

- *Public* maps are visible to all users, although to see them, the user must have read access to at least one map element. Public maps can be edited in case a user/user group has read-write permissions for this map and at least read permissions to all elements of the corresponding map, including triggers in the links.
- *Private* maps are visible only to their owner and the users/user groups the map is shared with by the owner. Regular (non-Super admin) users can only share with the groups they belong to and users who belong to those groups. Admin level users can see private maps regardless of being the owner or belonging to the shared user list. Private maps can be edited by the owner of the map and in case a user/user group has read-write permissions for this map and at least read permissions to all elements of the corresponding map, including triggers in the links.

Map elements that the user does not have read permission to are displayed with a grayed-out icon and all textual information on the element is hidden. However, the trigger label is visible even if the user has no permission to the trigger.

To add an element to the map, the user must also have at least read permission to the element.

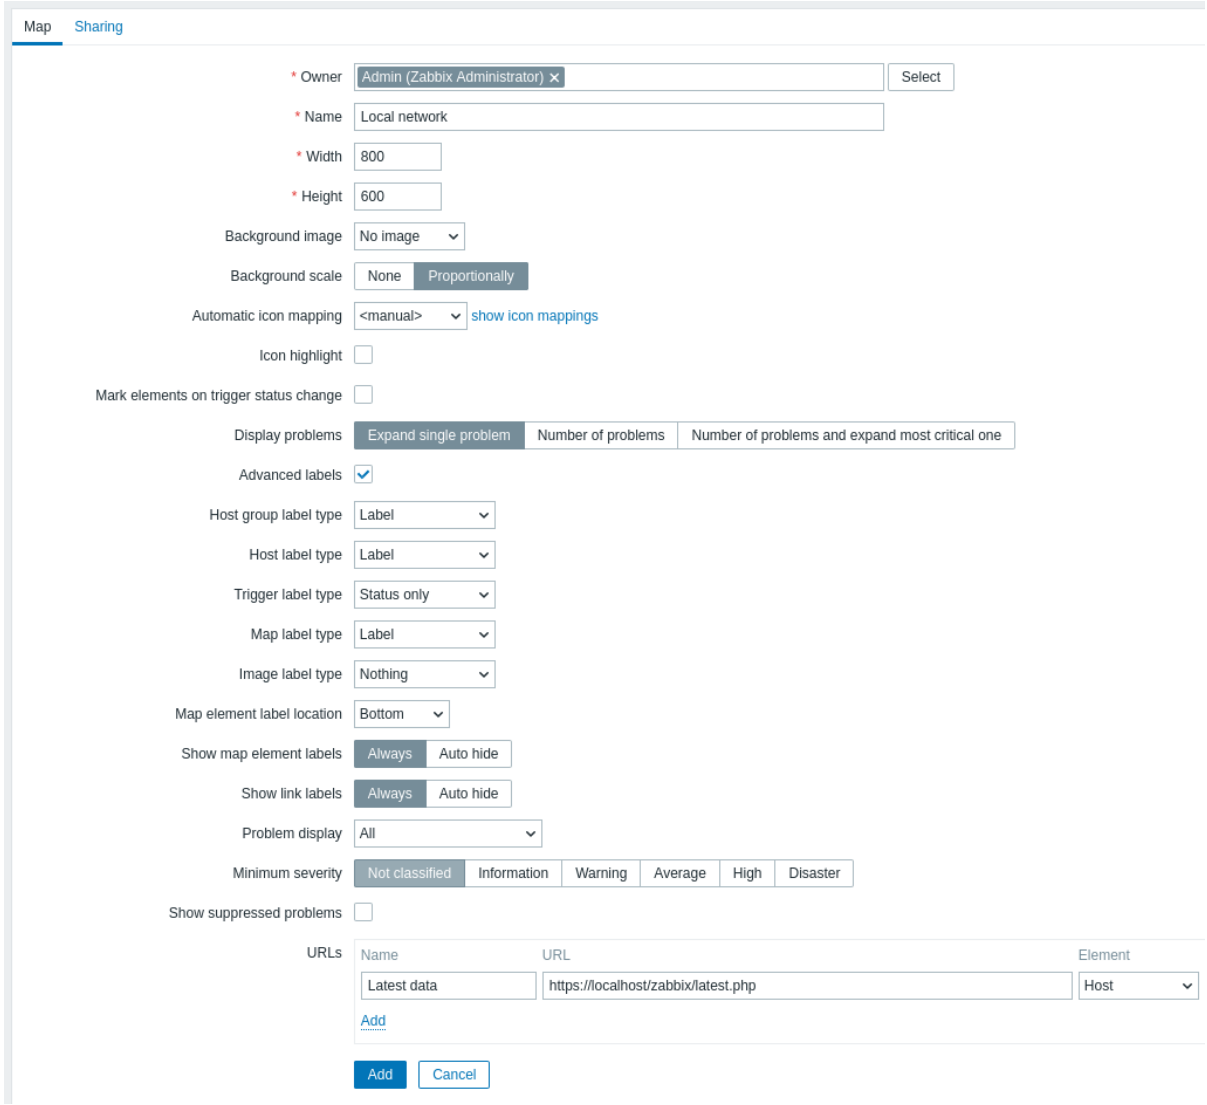
Creating a map

To create a map, do the following:

1. Go to *Monitoring* > *Maps*.
2. Go to the view with all maps.
3. Click *Create map*.
4. Configure the map details.

You can also use the *Clone* button in the configuration form of an existing map to create a new map. This map will have all of the properties of the existing map, including general layout attributes, as well as the elements of the existing map.

The **Map** tab contains general map attributes:



All mandatory input fields are marked with a red asterisk.

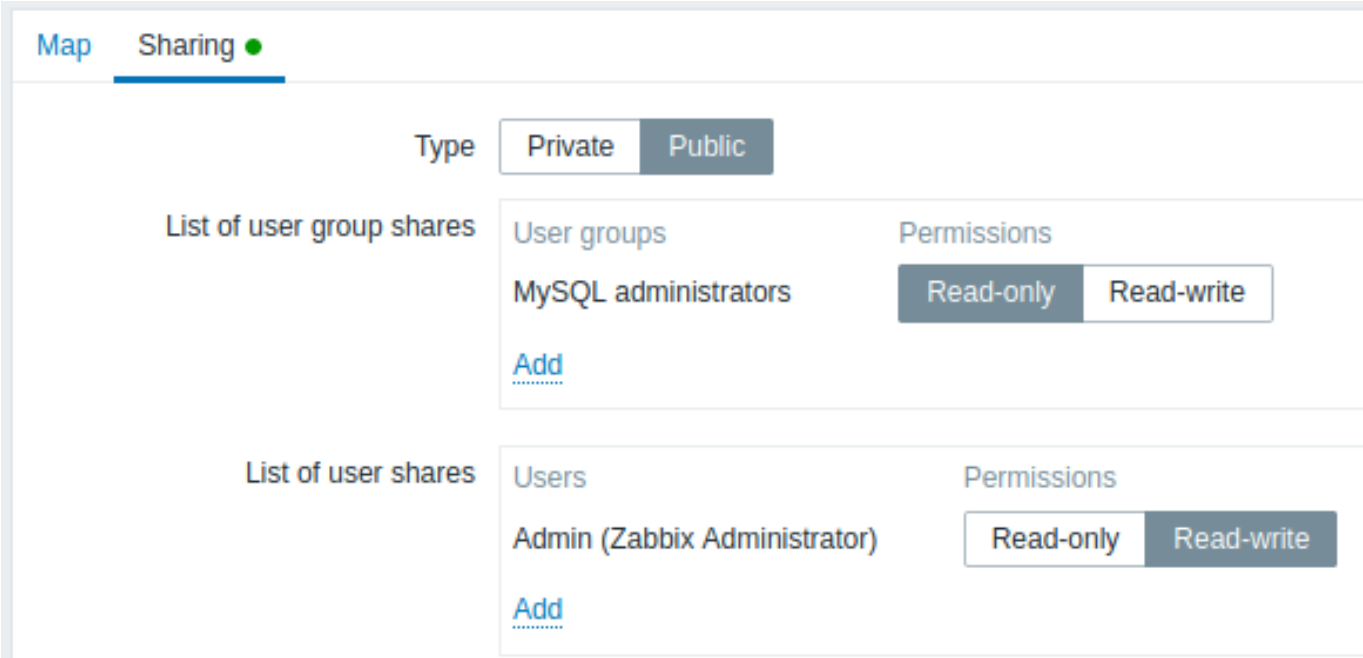
General map attributes:

Parameter	Description
Owner	Name of map owner.
Name	Unique map name.
Width	Map width in pixels.
Height	Map height in pixels.
Background image	Use background image: **No image** - no background image (white background) **Image** - selected image to be used as a background image. You may use a geographical map or any other image to enhance your map.
Background scale	Use background scaling: **None** - no scaling; **Proportionally** - cover the entire map background, but keep the image proportions (default).
Automatic icon mapping	You can set to use an automatic icon mapping, configured in *Administration > General > Icon mapping*. Icon mapping allows mapping certain icons against certain host inventory fields.
Icon highlight	If you mark this checkbox, map elements will be highlighted. Elements with an active trigger will receive a round background, in the same color as the highest severity trigger. Moreover, a thick green line will be displayed around the circle if all problems are acknowledged. Elements with ”disabled” or ”in maintenance” status will get a square background, gray and orange respectively. See also: Viewing maps
Mark elements on trigger status change	A recent change of trigger state (recent problem or resolution) will be highlighted with markers (inward-pointing red triangles) on the three sides of the element icon that are free of the label. Markers are displayed for 30 minutes.

Parameter	Description
Display problems	Select how problems are displayed with a map element: **Expand single problem** - if there is only one problem, the problem name is displayed. Otherwise, the total number of problems is displayed. **Number of problems** - the total number of problems is displayed **Number of problems and expand most critical one** - the name of the most critical problem and the total number of problems is displayed. 'Most critical' is determined based on problem severity and, if equal, problem event ID (higher ID or later problem displayed first). For a *trigger map element* it is based on problem severity and if equal, trigger position in the trigger list. In case of multiple problems of the same trigger, the most recent one will be displayed.
Advanced labels	If you mark this checkbox, you will be able to define separate label types for separate element types.
Host group label type	Label type used for the map element: **Label** - map element label **IP address** - IP address **Element name** - element name (for example, host name) **Status only** - status only (OK or PROBLEM) **Nothing** - no labels are displayed
Host label type *Trigger label type* *Map label type* *Image label type*	
Map element label location	Label location in relation to the map element: **Bottom** - beneath the map element **Left** - to the left **Right** - to the right **Top** - above the map element
Show map element labels	Select how to display map element labels: **Always** - always displayed (default); **Auto hide** - hide the label if it's not hovered on or selected.
Show link labels	Select how to display link labels: **Always** - always displayed (default); **Auto hide** - hide the label if it's not hovered on or selected.
Problem display	Display problem count as: **All** - full problem count will be displayed **Separated** - unacknowledged problem count will be displayed separated as a number of the total problem count **Unacknowledged only** - only the unacknowledged problem count will be displayed
Minimum severity	Problems below the selected minimum severity level will not be displayed on the map. For example, with *Warning* selected, changes with *Information* and *Not classified* level triggers will not be reflected in the map.
Show suppressed problems	Mark the checkbox to display problems that would otherwise be suppressed (not shown) because of host maintenance.
URLs	URLs for each element type can be defined here, where: **Name** - a label for the URL; **URL** - an URL (up to 2048 characters); **Element** - the element type (e.g., *Host*, *Host group*, etc.). These will be displayed as links when a user clicks on the element in the map viewing mode. Macros can be used in map URL names and values. For a full list, see supported macros and search for 'map URL names and values'.

Sharing

The **Sharing** tab contains the map type as well as sharing options (user groups, users) for private maps:



Parameter	Description
Type	Select map type: **Private** - map is visible only to selected user groups and users **Public** - map is visible to all
List of user group shares	Select user groups that the map is accessible to. You may allow read-only or read-write access.
List of user shares	Select users that the map is accessible to. You may allow read-only or read-write access.

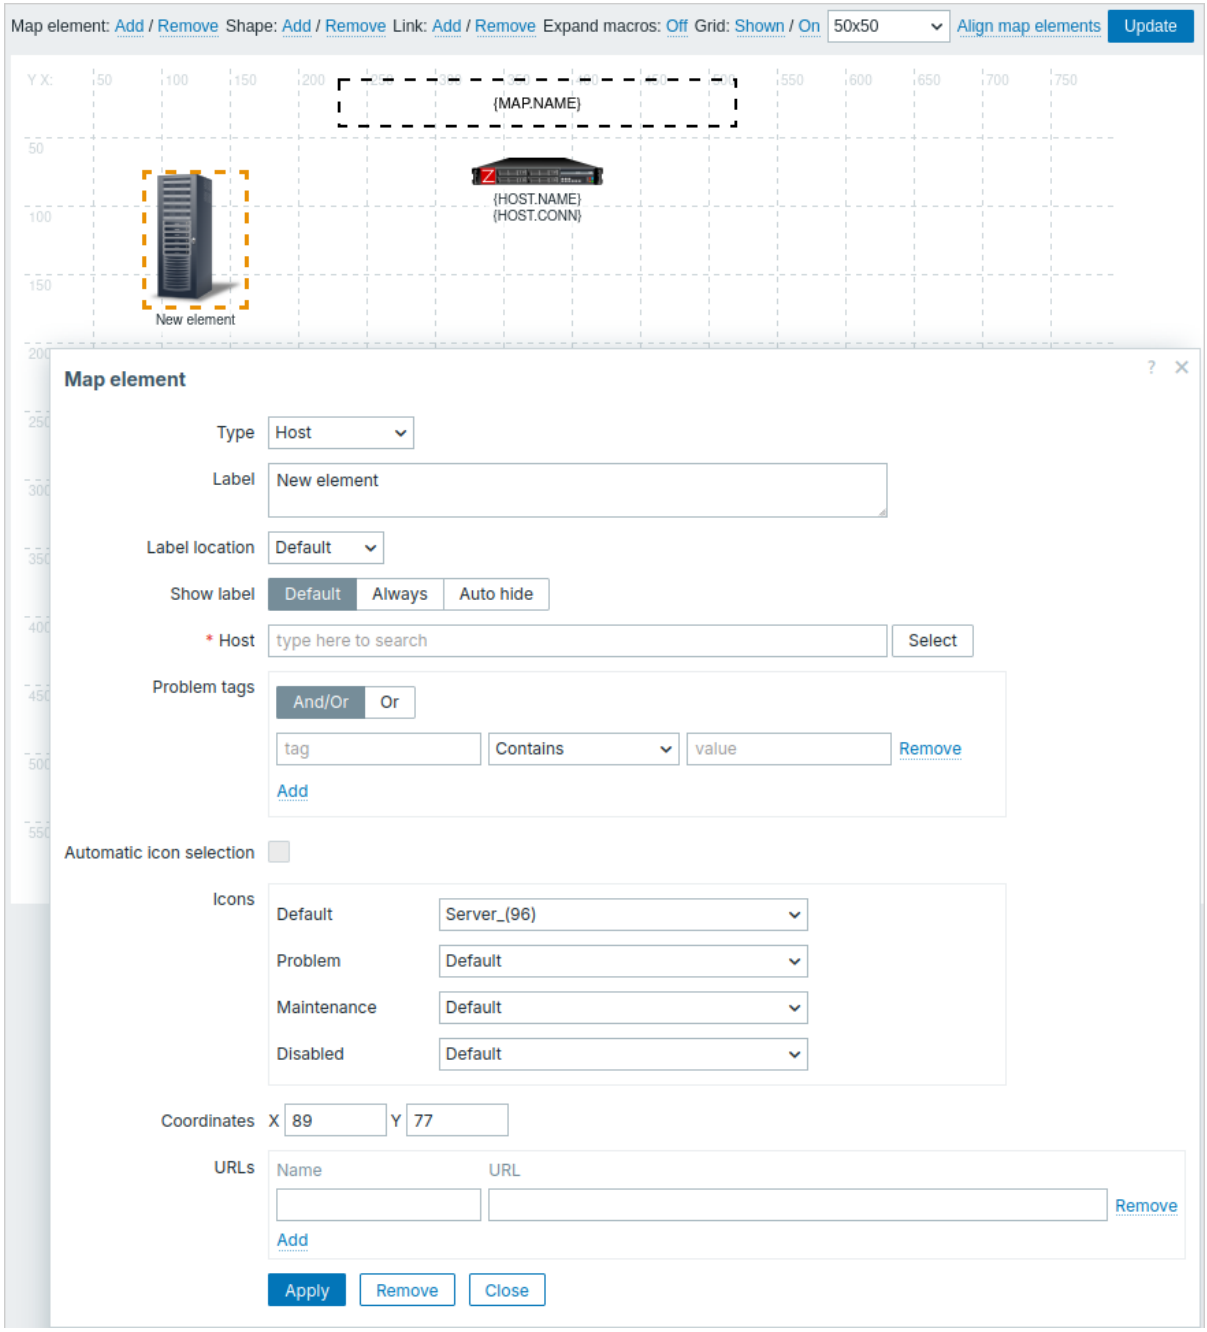
When you click *Add* to save this map, you have created an empty map with a name, dimensions, and certain preferences. Now you need to add some elements. For that, click *Edit* in the map list to open the editable area.

Adding elements

To add an element, click *Add* next to `Map element`. The new element will appear at the upper-left corner of the map. Drag and drop it wherever you like.

Note that with the Grid option "On", elements will always align to the grid (you can pick various grid sizes from the dropdown, also hide/show the grid). If you want to put elements anywhere without alignment, turn the option to "Off". (You can align random elements to the grid later, by clicking *Align map elements*.)

Now that you have some elements in place, you may want to start differentiating them by giving names, etc. By clicking the element, a form is displayed and you can set the element type, give a name, choose a different icon, etc.



Map element attributes:

Parameter	Description
Type	Type of the element: **Host** - icon representing status of all triggers of the selected host **Map** - icon representing status of all elements of a map **Trigger** - icon representing status of one or more triggers **Host group** - icon representing status of all triggers of all hosts belonging to the selected group **Image** - an icon, not linked to any resource
Label	Icon label, any string. Macros and multiline strings can be used. Expression macros are supported in this field, but only with `avg`, `last`, `min` and `max` functions, with time as parameter (for example, `{?avg(/host/key,1h)}`). For a full list of supported macros, see supported macros and search for 'map element labels'.
Label location	Label location in relation to the icon: **Default** - map's default label location; **Bottom** - beneath the icon; **Left** - to the left; **Right** - to the right; **Top** - above the icon.

Parameter	Description
Show label	Select how to display the element label: **Default** - use *Show map element labels* settings from map configuration (default); **Always** - always displayed (default); **Auto hide** - hide the label if it's not hovered on or selected.
Host	Select a host; alternatively, enter the host name. This field supports auto-completion, so starting to type the name of a host will display a dropdown of matching hosts. Scroll down to select one. Click the remove (x) icon to remove the selection. This field is available only for elements of type *Host*.
Map	Select a map; alternatively, enter the map name. This field supports auto-completion, so starting to type the name of a map will display a dropdown of matching maps. Scroll down to select one. Click the remove (x) icon to remove the selection. This field is available only for elements of type *Map*.
Triggers	Select one or more triggers in the *New triggers* field below and click *Add.*Alternatively, start typing the name of a trigger and select one from the dropdown list of matching triggers. Click the remove (x) icon to remove the selection. This field is available only for elements of type *Trigger*. The order of selected triggers can be changed, but only within the same severity of triggers. Multiple trigger selection also affects {HOST.*} macro resolution both in the editing and view modes. In edit mode, the first displayed {HOST.*} macros will be resolved depending on the first trigger in the list (based on trigger severity). View mode depends on the Display problems parameter in General map attributes: - If *Expand single problem* mode is chosen, the first displayed {HOST.*} macros will be resolved depending on the latest detected problem trigger (not mattering the severity) or the first trigger in the list (in case no problem detected). - If *Number of problems and expand most critical one* mode is chosen, the first displayed {HOST.*} macros will be resolved depending on the trigger severity.
Host group	Select a host group; alternatively, enter the host group. This field supports auto-completion, so starting to type the name of a host group will display a dropdown of matching host groups. Scroll down to select one. Click the remove (x) icon to remove the selection. This field is available only for elements of type *Host group*.
Problem tags	Specify tags to limit the number of problems displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names **Equals** - include the specified tag names and values (case-sensitive) **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive) **Does not exist** - exclude the specified tag names **Does not equal** - exclude the specified tag names and values (case-sensitive) **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive) There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition **Or** - enough if one condition is met This field is available only for elements of type *Host* and *Host group*.
Automatic icon selection	When the checkbox is marked, the icon mapping will be used to determine which icon to display.
Icons	You can choose to display different icons for the element in these cases: *Default*, *Problem*, *Maintenance*, and *Disabled*. This field is available only for elements of type *Host*.
Coordinates	**X** and **Y** coordinates of the map element.

Parameter	Description
URLs	Element-specific URLs can be set for the element here, where: **Name** - a label for the URL; **URL** - an URL (up to 2048 characters). These will be displayed as links when a user clicks on the element in the map viewing mode. If the element has its own URLs and there are map level URLs for its type defined, they will be combined in the same menu. Macros can be used in map element names and values. For a full list, see supported macros and search for 'map URL names and values'.

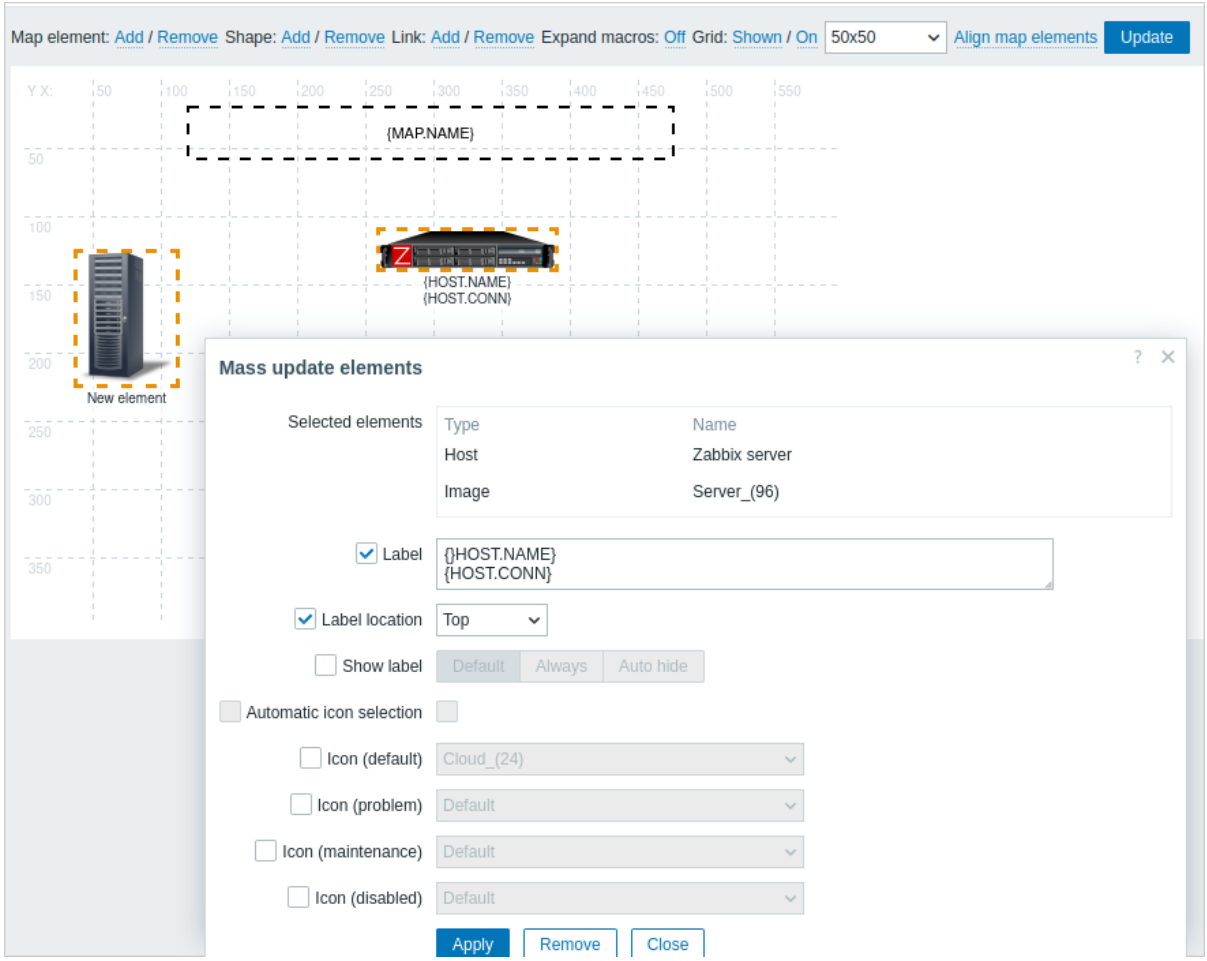
Attention:
Added elements are not automatically saved. If you navigate away from the page, all changes may be lost.
Therefore, it is a good idea to click the **Update** button in the upper-right corner. Once clicked, the changes are saved regardless of what you choose in the following popup.
Selected grid options are also saved with each map.

Selecting elements

To select elements, select one and then hold down *Ctrl* to select the others.

You can also select multiple elements by dragging a rectangle in the editable area and selecting all elements in it.

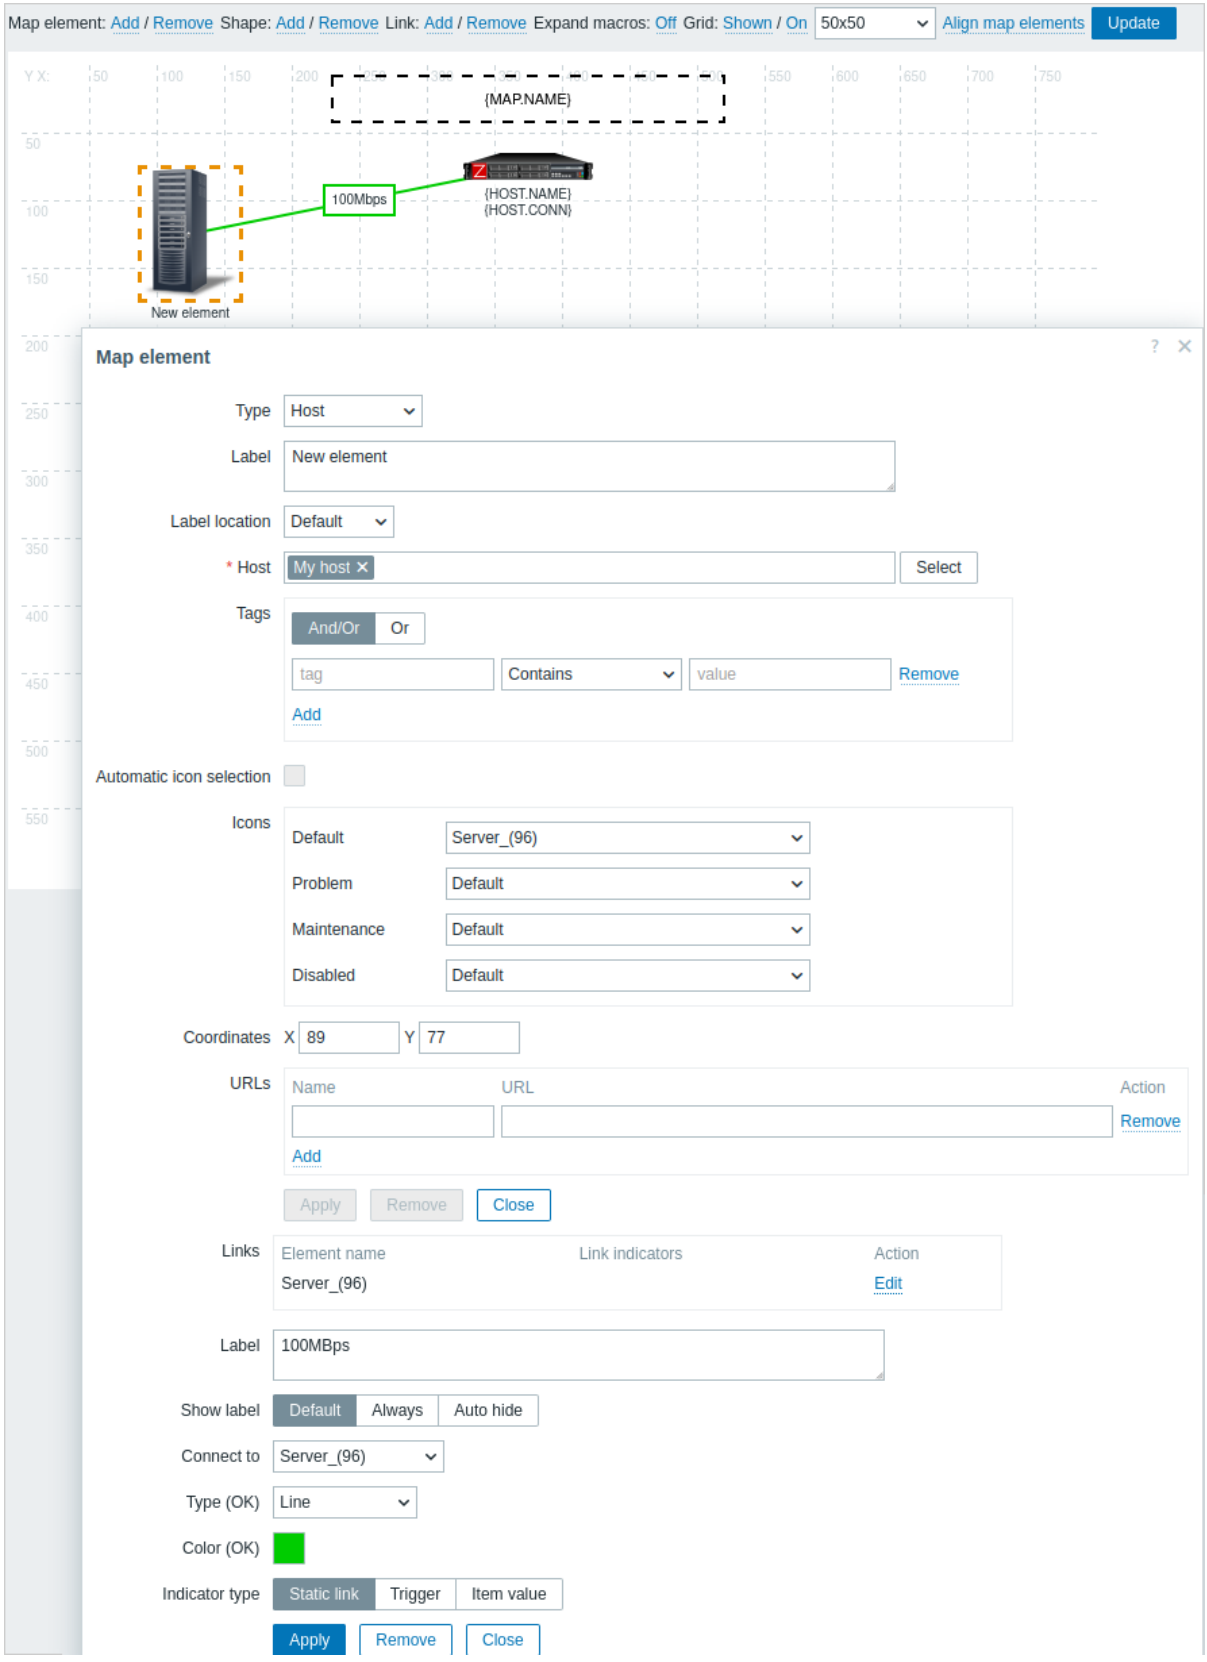
Once you select more than one element, the element property form shifts to the mass-update mode so you can change attributes of selected elements in one go. To do so, mark the attribute using the checkbox and enter a new value for it. You may use macros here (for example, {HOST.NAME} for the element label).



Linking elements

Once you have put some elements on the map, it is time to start linking them. To link two elements you must first select them. With the elements selected, click *Add* next to Link.

With a link created, the single element form now contains an additional *Links* section. Click *Edit* to edit link attributes.



Link attributes:

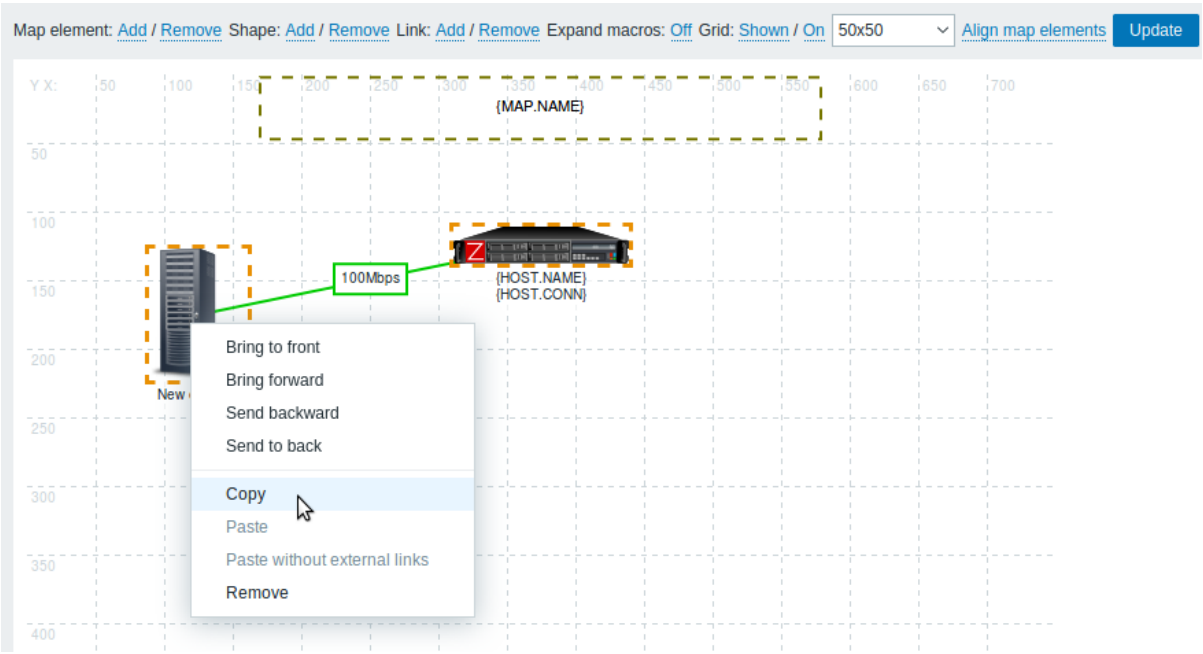
Parameter	Description
Label	Label that will be rendered on top of the link. Expression macros are supported in this field, but only with `avg`, `last`, `min` and `max` functions, with time as parameter (for example, `{?avg(/host/key,1h)}`).
Show label	Select how to display the link label: **Default** - use *Show map element labels* settings from map configuration (default); **Always** - always displayed (default); **Auto hide** - hide the label if it's not hovered on or selected.

Parameter	Description
Connect to	The element that the link connects to.
Type (OK)	Default link style. Select the style: *Line*, *Bold line*, *Dot*, or *Dashed line*.
Color (OK)	Default link color. Click the color swatch to select another color.
Indicator type	Select the link indicator type: **Static link** - no indicators; **Trigger** - allow trigger-based link indicators; **Item value** - allow item value-based link indicators.
Item	Select the item. The value of this item will affect the link style. This field is available when *Item value* is selected for *Indicator type*.
Indicators	The list of link indicators. This field is available when *Indicator type* is set to *Trigger* or *Item value*. If set to *Trigger*, triggers can be assigned as link indicators. When a trigger is in the *Problem* state, its style (selected color and line type) will be applied to the link. If set to *Item value* and an item is selected in the *Item* field, item values can be assigned as link indicators. When the item value reaches the specified threshold (for a numeric item) or matches the pattern (for a text item), its style (selected color and line type) is applied to the link.

Moving and copy-pasting elements

Several selected elements can be **moved** to another place in the map by clicking one of the selected elements, holding down the mouse button, and moving the cursor to the desired location.

One or more elements can be **copied** by selecting the elements, then clicking a selected element with the right mouse button and selecting *Copy* from the menu.

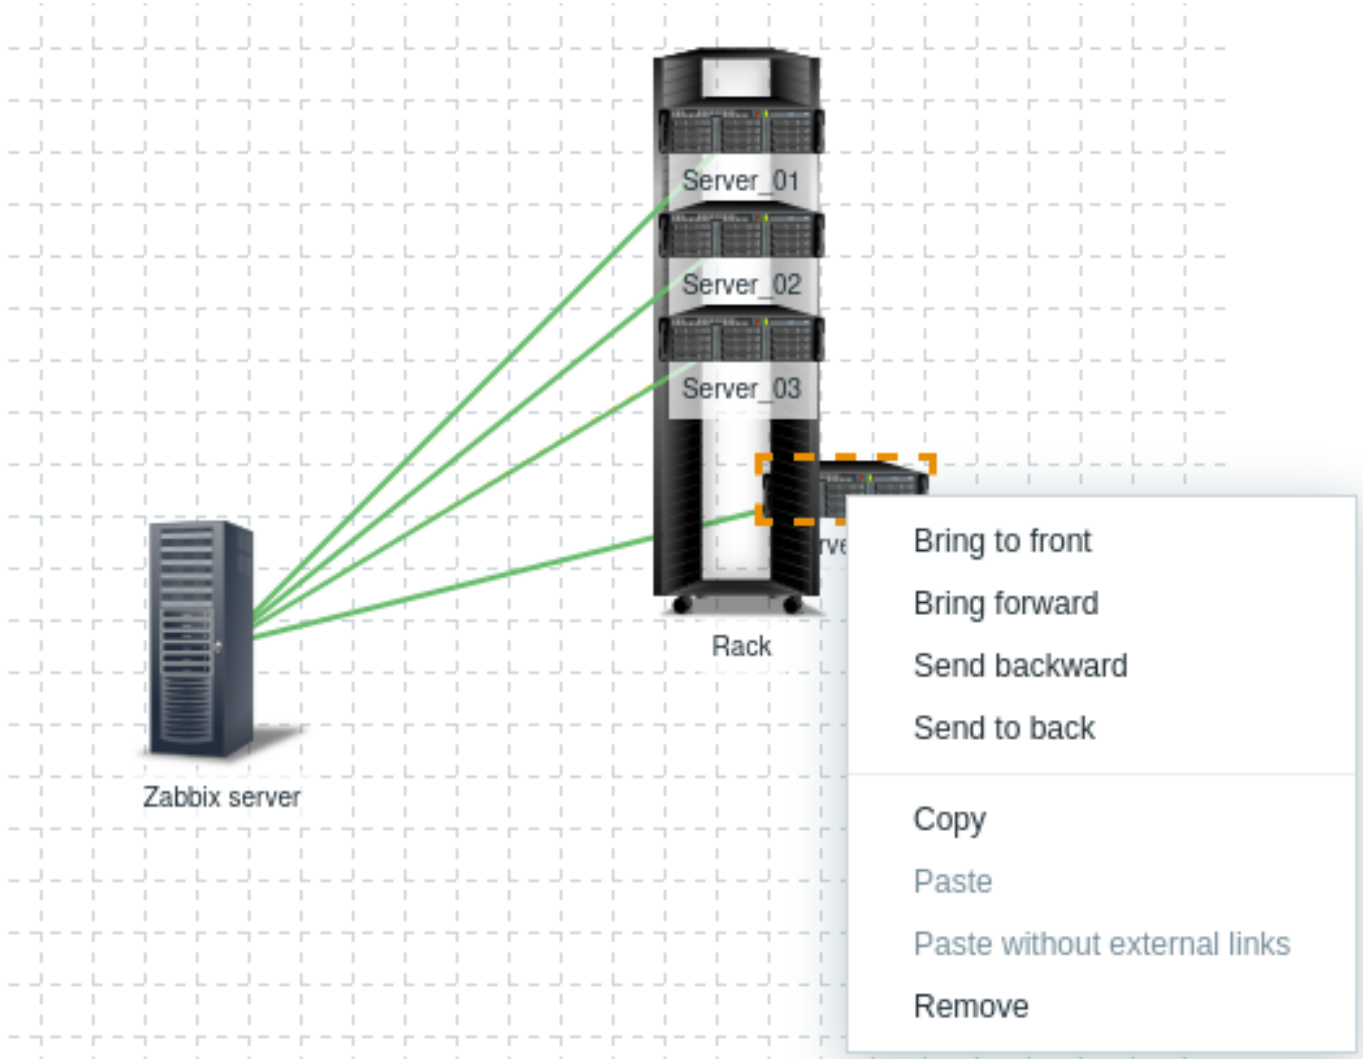


To paste the elements, click a map area with the right mouse button and select *Paste* from the menu. The *Paste without external links* option will paste the elements retaining only the links that are between the selected elements.

Copy-pasting works within the same browser window. Keyboard shortcuts are not supported.

Ordering elements

To bring one element in front of the other (or vice versa), click the element with the right mouse button and select *Bring forward*/*Bring to front* or *Send backward*/*Send to back*.

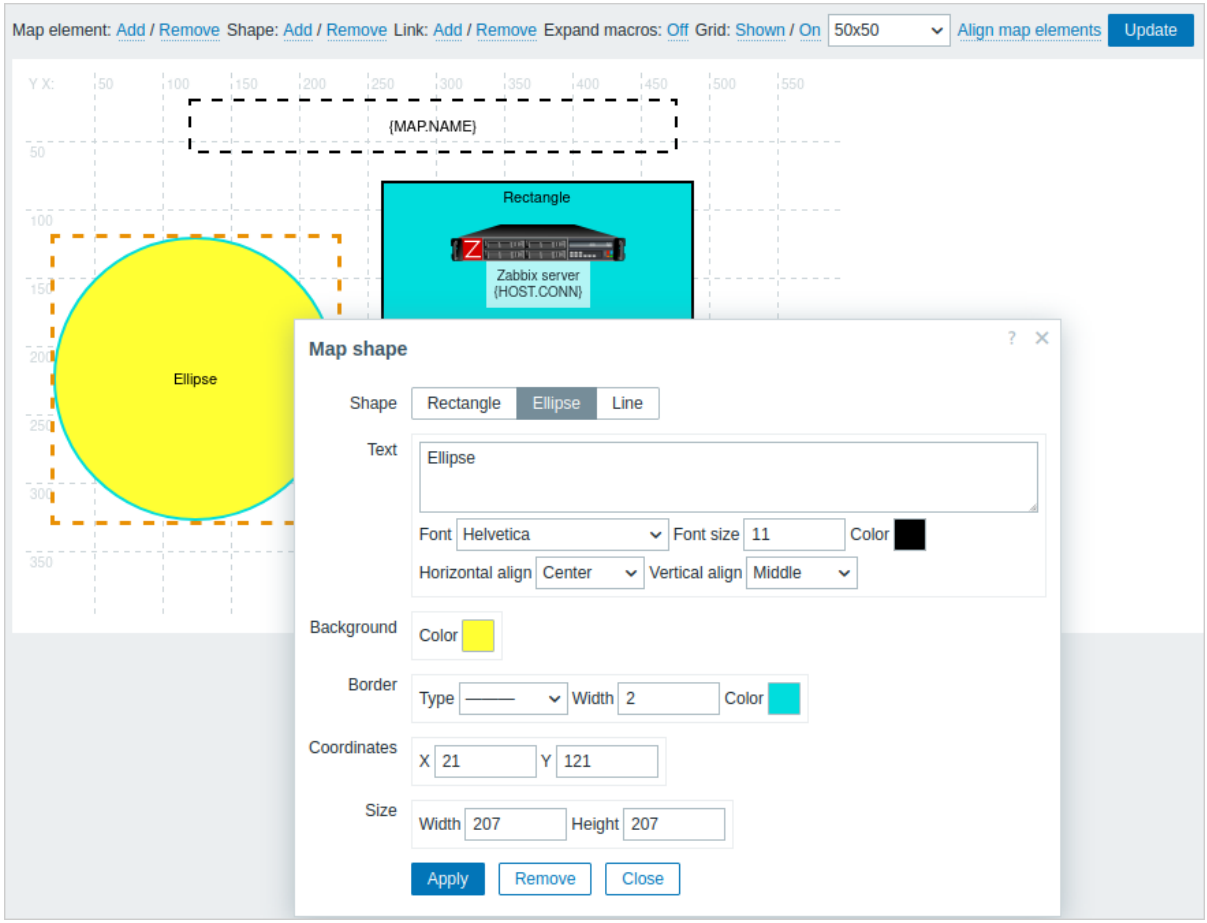


Adding shapes

In addition to map elements, it is also possible to add some shapes. Shapes are not map elements; they are just a visual representation. For example, a rectangle shape can be used as a background to group some hosts. Rectangle and ellipse shapes can be added.

To add a shape, click *Add* next to Shape. The new shape will appear at the upper-left corner of the map. Drag and drop it wherever you like.

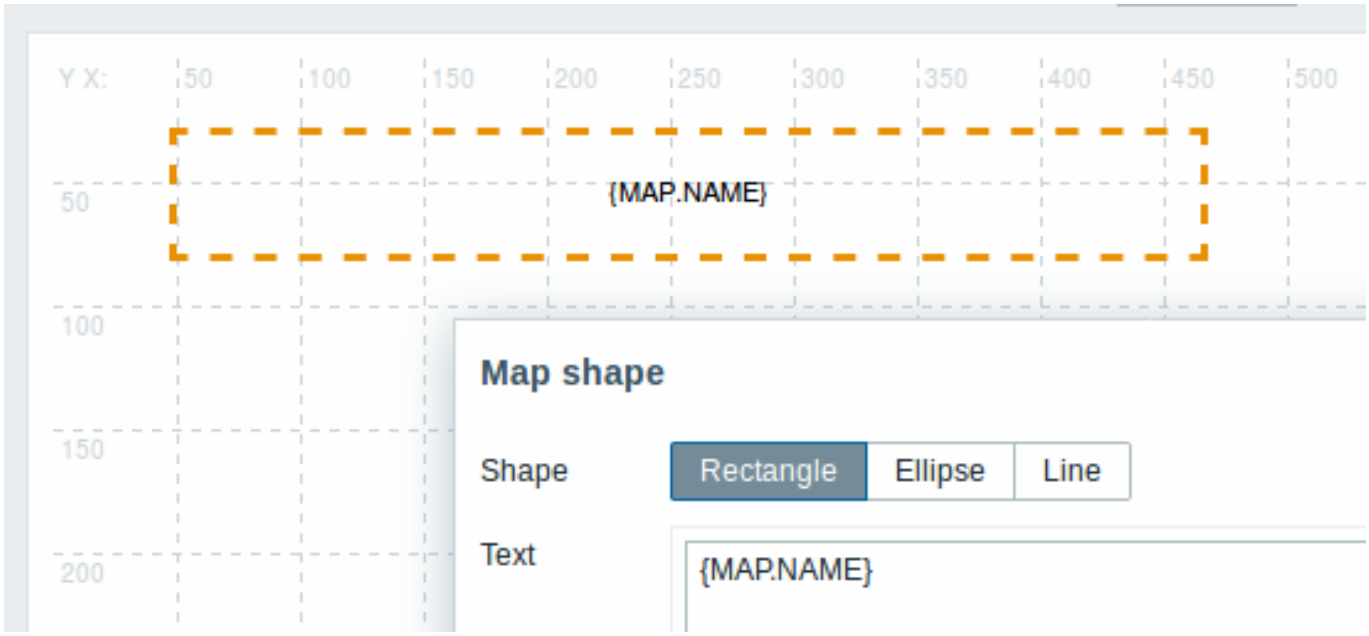
A new shape is added with default colors. By clicking the shape, a form is displayed and you can customize the way a shape looks, add text, etc.



To select shapes, select one and then hold down *Ctrl* to select the others. With several shapes selected, common properties can be mass updated, similarly as with elements.

Text can be added in the shapes. Expression macros are supported in the text, but only with `avg`, `last`, `min` and `max` functions, with time as parameter (for example, `{?avg(/host/key,1h)}`).

To display text only, the shape can be made invisible by removing the shape border (select 'None' in the *Border* field). For example, take note of how the {MAP.NAME} macro, visible in the screenshot above, is actually a rectangle shape with text, which can be seen when clicking the macro:



{MAP.NAME} resolves to the configured map name when viewing the map.

If hyperlinks are used in the text, they become clickable when viewing the map.

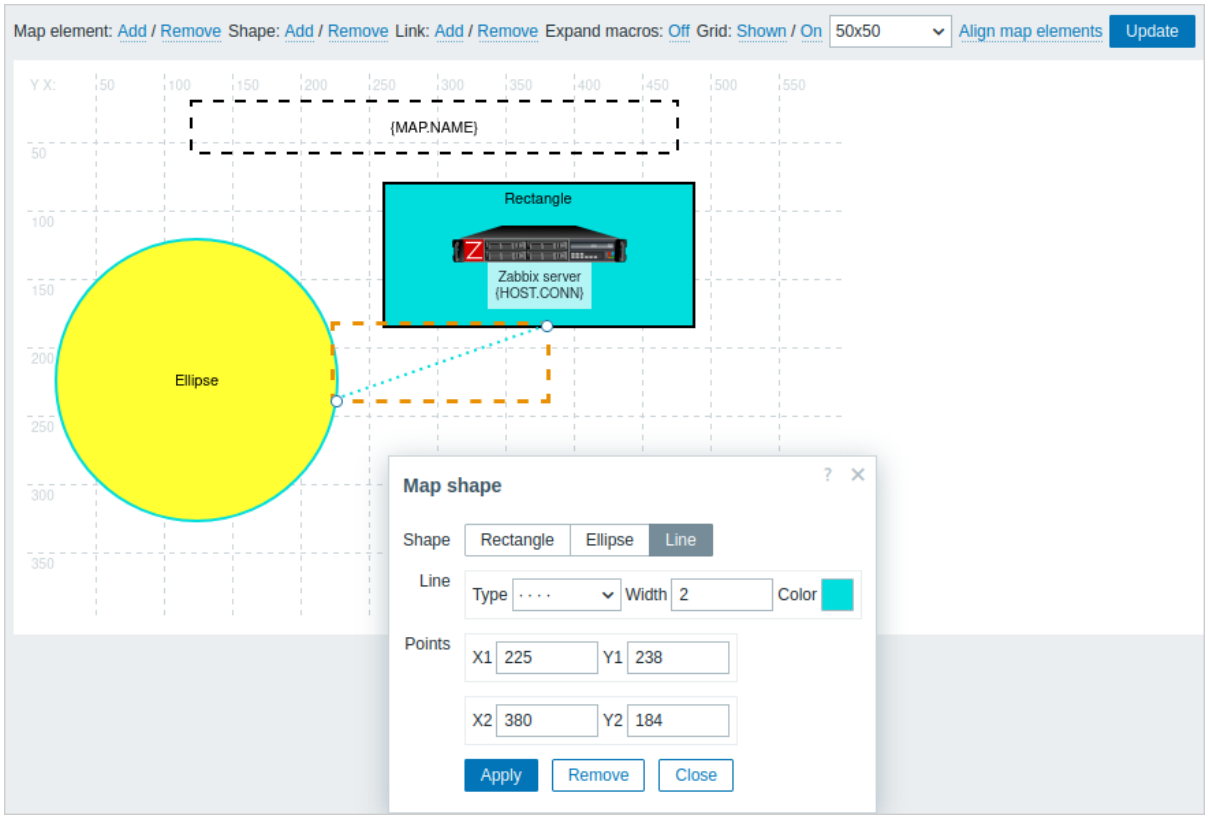
Line wrapping for text is always "on" within shapes. However, within an ellipse, the lines are wrapped as though the ellipse were a rectangle. Word wrapping is not implemented, so long words (words that do not fit the shape) are not wrapped, but are masked

(on map editing page) or clipped (other pages with maps).

Adding lines

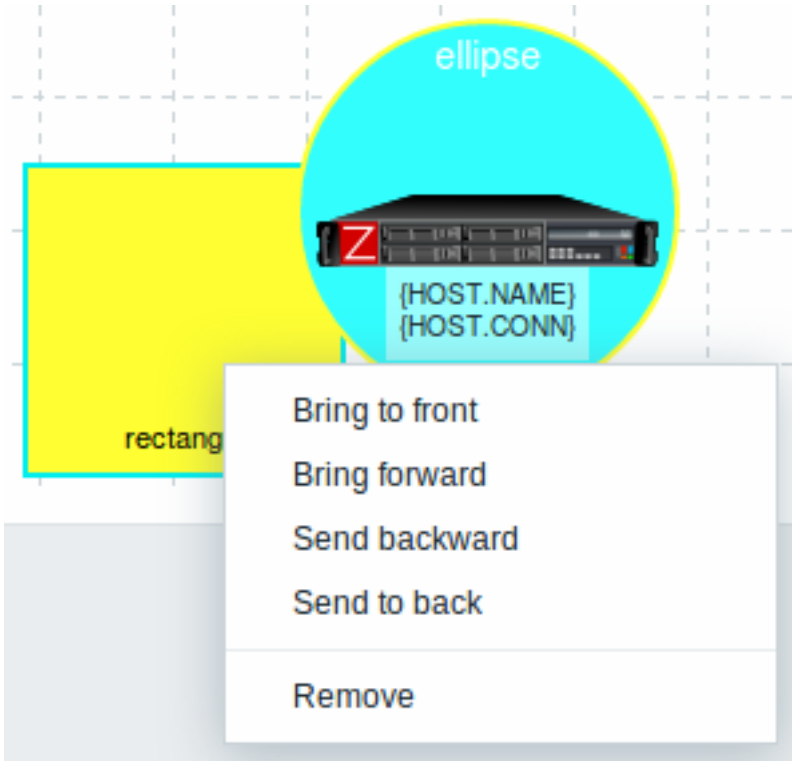
In addition to shapes, it is also possible to add some lines. Lines can be used to link elements or shapes in a map.

To add a line, click *Add* next to Shape. A new shape will appear at the upper-left corner of the map. Select it and click *Line* in the editing form to change the shape into a line. Then adjust line properties, such as line type, width, color, etc.



Ordering shapes and lines

To bring one shape in front of the other (or vice versa), click the shape with the right mouse button bringing up the map shape menu.

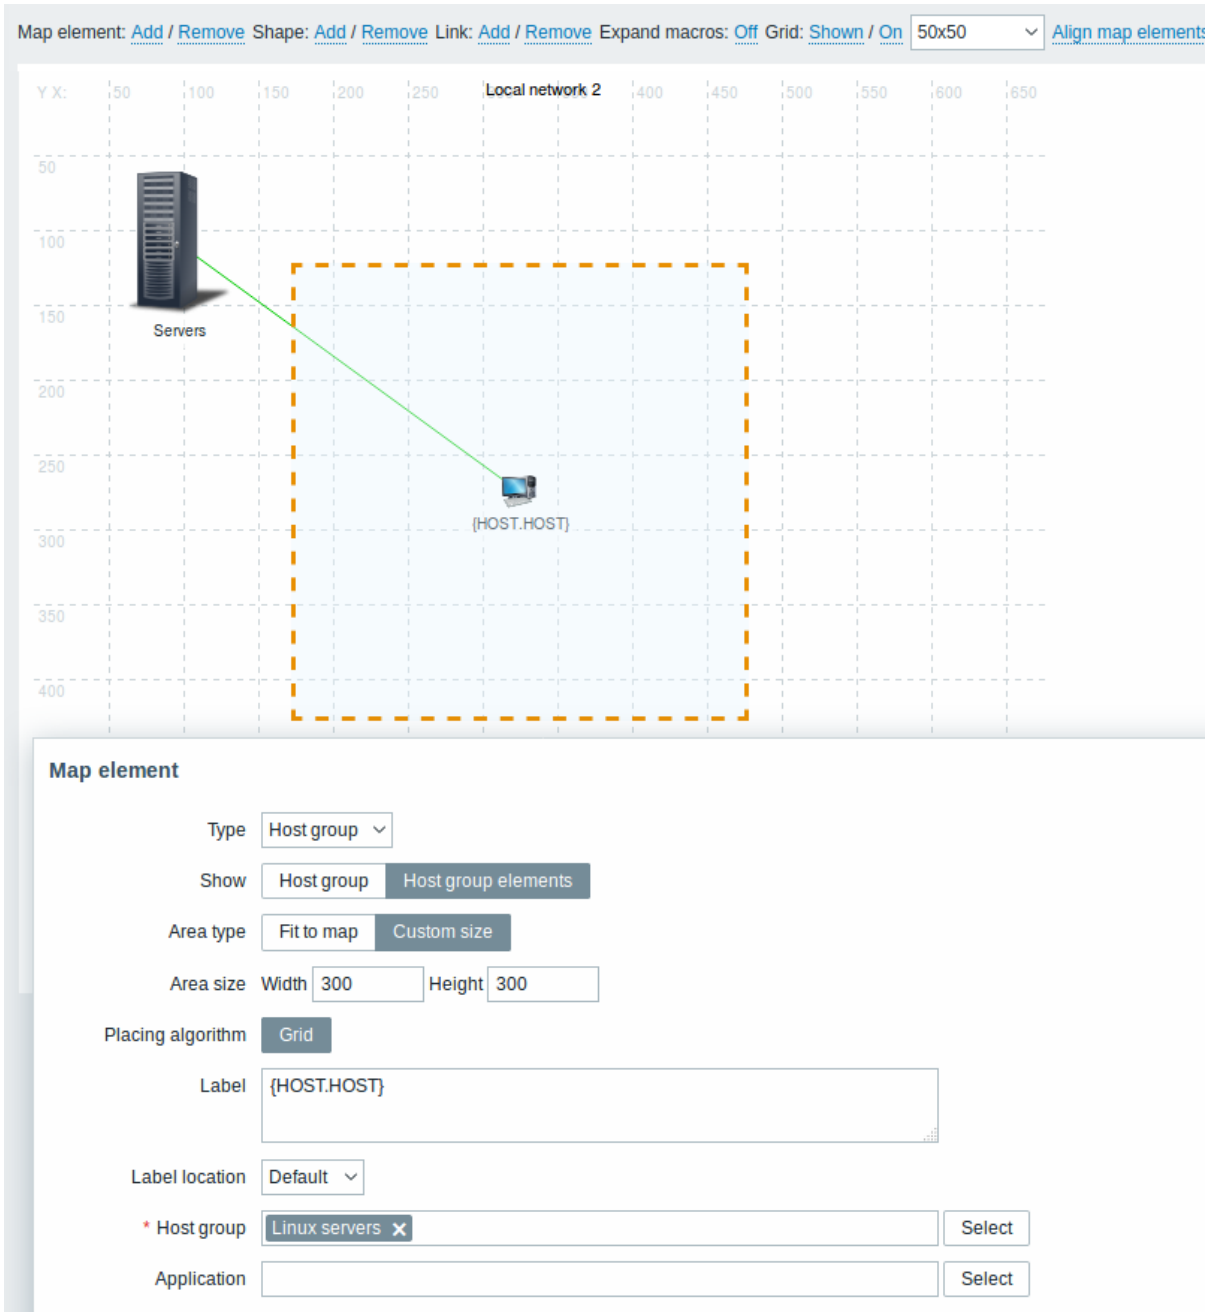


Host group elements

Overview

This section explains how to add a *Host group* type element when configuring a network map.

Configuration



All mandatory input fields are marked with a red asterisk.

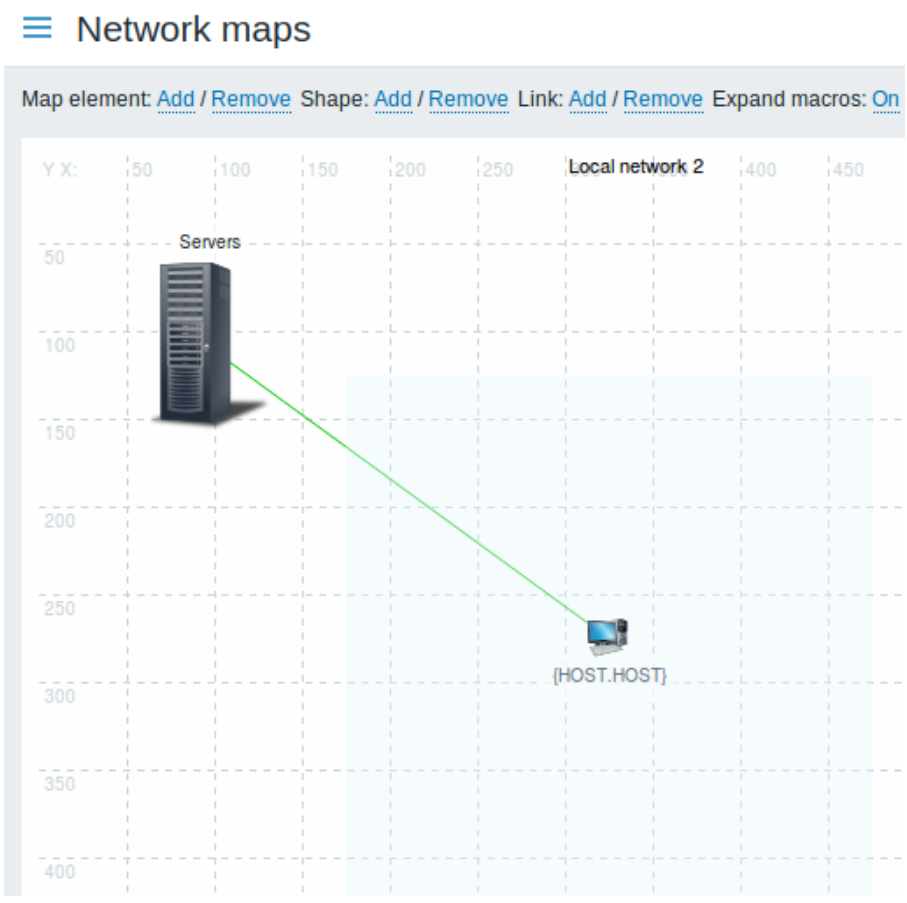
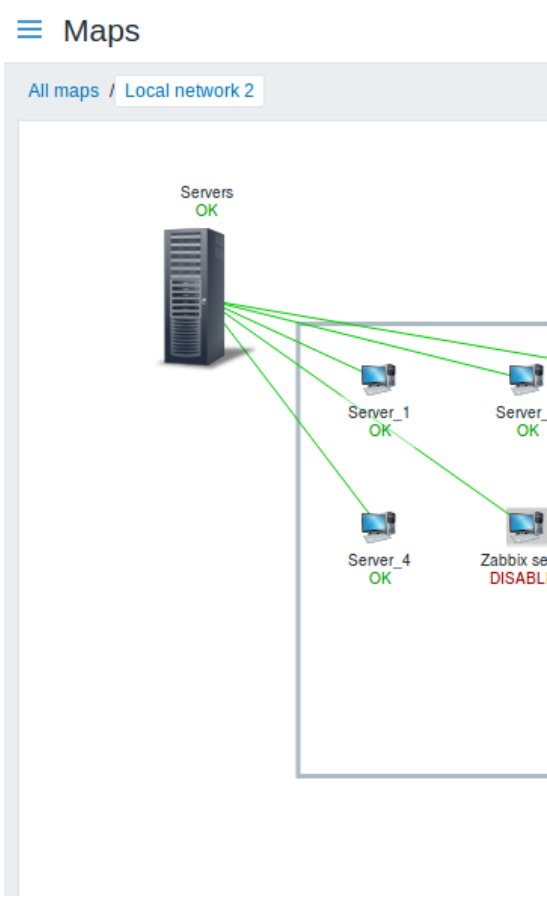
This table consists of parameters typical for *Host group* element type:

Parameter	Description
Type	Select Type of the element: **Host group** - Icon representing the status of all triggers of all hosts belonging to the selected group.
Show	Show options: **Host group** - Selecting this option will result as one single icon displaying corresponding information about the certain host group. **Host group elements** - Selecting this option will result as multiple icons displaying corresponding information about every single element (host) of the certain host group.

Parameter	Description
Area type	This setting is available if *Show* is set to *Host group elements*: **Fit to map** - All host group elements are equally placed within the map. **Custom size** - A manual setting of the map area for all the host group elements to be displayed.
Area size	This setting is available if *Area type* is set to *Custom size*: **Width** - Value in pixels to specify the map area width. **Height** - Value in pixels to specify the map area height.
Placing algorithm	**Grid** - Only available option of displaying all the host group elements.
Label	Icon label, any string. Macros and multiline strings can be used in labels. If the type of the map element is *Host group* specifying certain macros has an impact on the map view displaying corresponding information about every single host. For example, if `{HOST.IP}` macro is used, the edit map view will only display the macro `{HOST.IP}` itself while map view will include and display each host's unique IP address.

Viewing host group elements

This option is available if the *Host group elements* show option is chosen. When selecting *Host group elements* as the *show* option, you will at first see only one icon for the host group. However, when you save the map and then go to the map view, you will see that the map includes all the elements (hosts) of the certain host group:

Map editing view	Map view
	

Notice how the `{HOST.HOST}` macro is used. In map editing, the macro name is unresolved, while in map view all the unique names of the hosts are displayed.

Link indicators

Overview

You can assign indicators to a link between elements in a network map.

The indicators can be based on triggers or item values. It's possible to display different link style and color:

- when triggers go into a problem state;
- when item value:
  - reaches a threshold (for numeric items);
  - matches a regular expression (for text items).

When you configure a link, you set the default link type and color. By assigning indicators to a link, it becomes possible to make the link style and color depend on trigger state or item value.

For example, should any of the assigned triggers go into a problem state, the link style and color will change to reflect that. So maybe your default link was a green line. Now, with the trigger in the problem state, your link may become bold red (if you have defined it so).

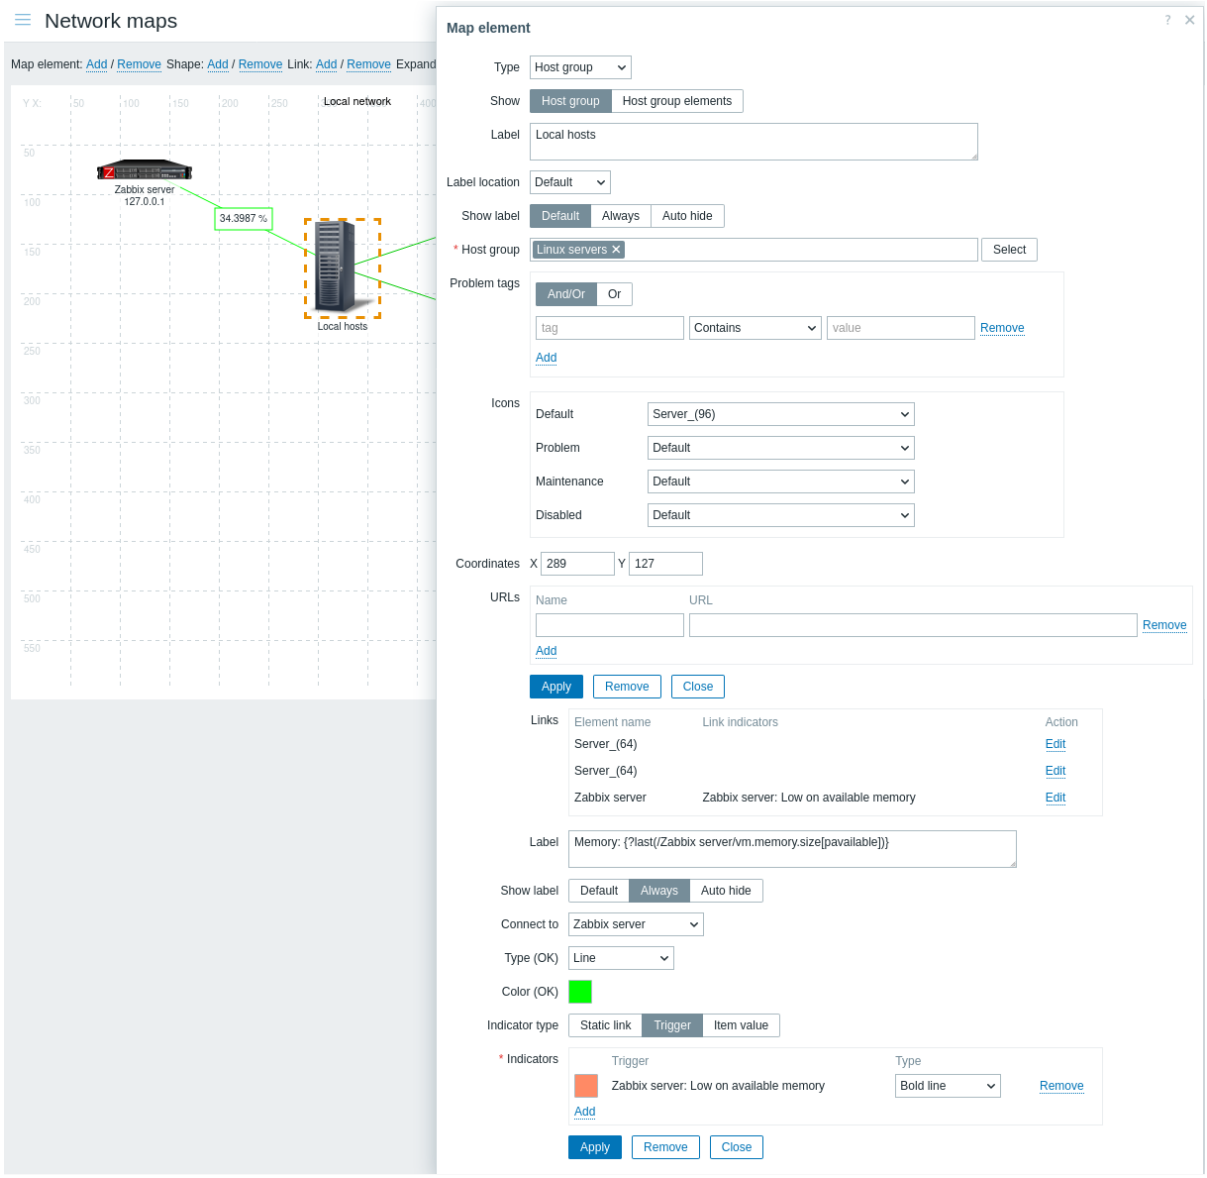
Similarly, if an item value reaches a specified threshold or matches a specified regular expression, the link style may reflect that.

Configuration

Triggers

To assign triggers as link indicators, do the following:

- select a map element
- click *Edit* in the *Links* section for the appropriate link
- select *Trigger* as the indicator type
- click *Add* in the *Indicators* block and select one or more triggers

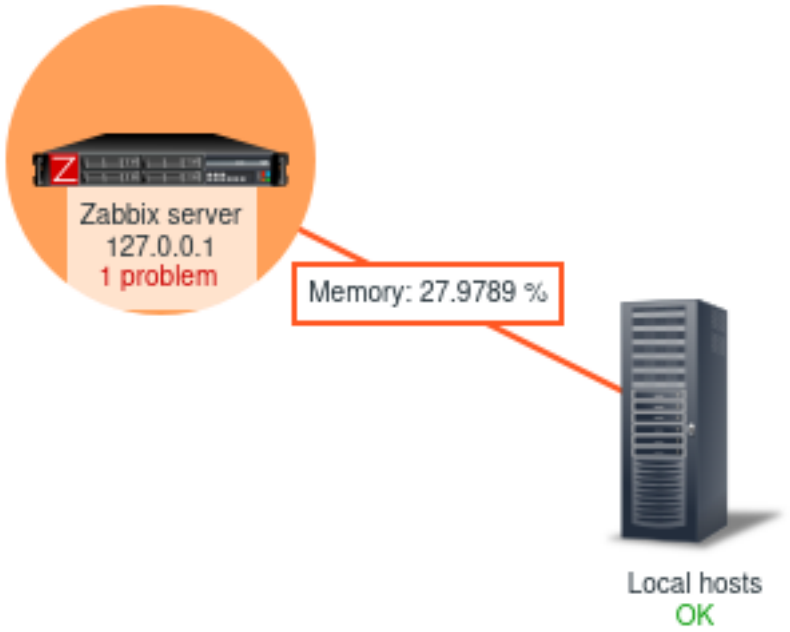


All mandatory input fields are marked with a red asterisk.

Added triggers can be seen in the *Indicators* list.

You can set the link type and color for each indicator directly from the list. When done, click *Apply*, close the form and click *Update* to save the map changes.

In *Monitoring* > *Maps* the respective link type and color will be displayed if the **trigger** goes into a problem state.



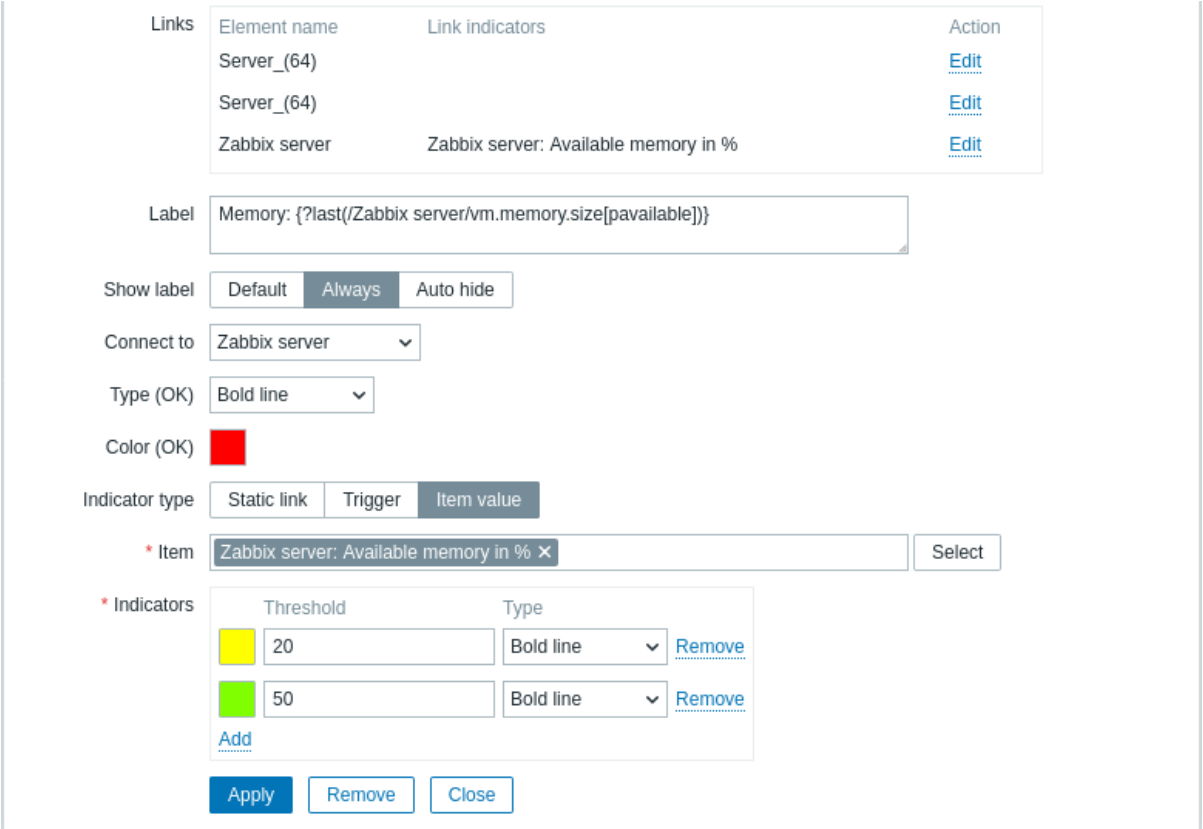
Note:
If multiple triggers go into a problem state, the problem with the highest severity will determine the link style and color. If multiple triggers with the same severity are assigned to the same map link, the one with the lowest ID takes precedence. Note also that:

1. *Minimum trigger severity* and *Show suppressed problem* settings from map configuration affect which problems are taken into account.
2. In the case of triggers with multiple problems (multiple problem generation), each problem may have a severity that differs from trigger severity (changed manually), may have different tags (due to macros), and may be suppressed.

Item values

To assign item values as link indicators, do the following:

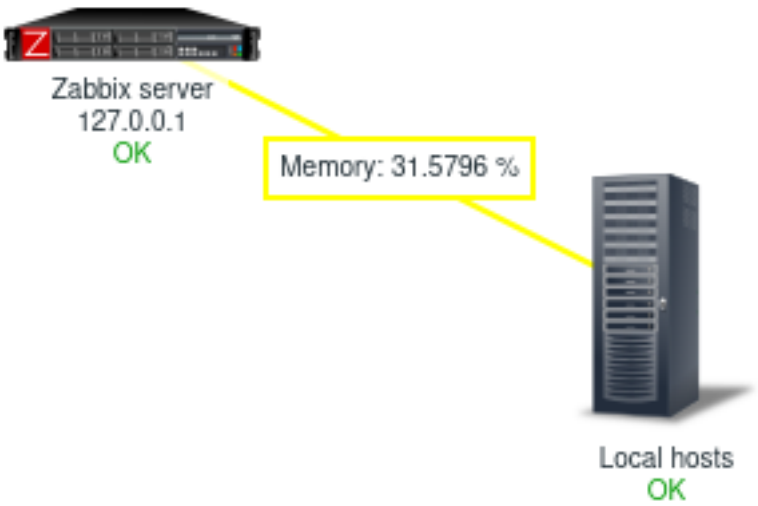
- select a map element
- click *Edit* in the *Links* section for the appropriate link
- select *Item value* as the indicator type
- select the item
- add one or more item value thresholds or patterns in the *Indicators* block



Added item thresholds/patterns can be seen in the *Indicators* list.

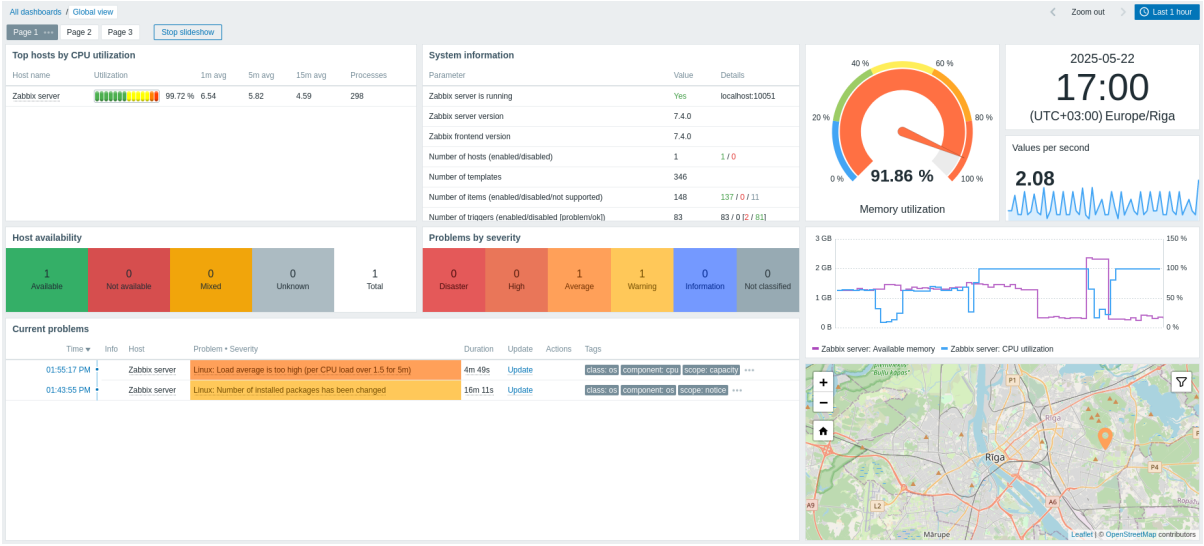
You can set the link type and color for each indicator directly from the list. When done, click *Apply*, close the form and click *Update* to save the map changes.

In *Monitoring* > *Maps* the respective link type and color will be displayed if the **item value** reaches the set threshold (for numeric data types) or matches the regular expression pattern (for text data types).



Dashboards

Dashboards - both global dashboards and host dashboards - provide a strong visualization platform with such widgets and tools as modern graphs, maps, slideshows, and more.



Templates and template groups

Overview

The use of templates is an excellent way of reducing one's workload and streamlining the Zabbix configuration. A template is a set of entities that can be conveniently applied to multiple hosts.

The entities may be:

- items
- triggers
- graphs
- dashboards
- low-level discovery rules
- web scenarios

As many hosts in real life are identical or fairly similar so it naturally follows that the set of entities (items, triggers, graphs,...) you have created for one host, may be useful for many. Of course, you could copy them to each new host, but that would be a lot of manual work. Instead, with templates you can copy them to one template and then apply the template to as many hosts as needed.

When a template is linked to a host, all entities (items, triggers, graphs,...) of the template are added to the host. Templates are assigned to each individual host directly (and not to a host group).

Templates are often used to group entities for particular services or applications (like Apache, MySQL, PostgreSQL, Postfix...) and then applied to hosts running those services.

Another benefit of using templates is when something has to be changed for all the hosts. Changing something on the template level once will propagate the change to all the linked hosts.

Templates are organized in template groups.

Proceed to creating and configuring a template.

Templates out of the box

Overview

Zabbix provides a growing set of pre-configured templates to simplify and speed up the setup of monitoring targets.

All out-of-the-box templates are available in *Data collection* > *Templates*.

When you upgrade Zabbix, templates are not upgraded automatically. This avoids overwriting any custom changes you have made. To upgrade or add templates manually, see the *Template upgrade* section below.

Please use the sidebar to access information about specific template types and operation requirements.

Global dashboards designed to work with existing Zabbix templates can be downloaded from the Dashboards repository and imported into your Zabbix environment. A detailed description of each dashboard, including prerequisites, is available in the dashboard's `README` file.

See also:

- Linking a template
- Known issues for templates

Template upgrade

To upgrade or add a **single template**:

1. Open the `templates` directory in Zabbix Git repository.
2. Choose the branch (under the *Source* label) that matches your Zabbix version (e.g., *master* for the development version).
3. Find the template you want to add or upgrade (e.g., MySQL by Zabbix agent).
4. Copy the contents of the template file (e.g., `template_db_mysql_agent.yaml`) and save them to a local file.
5. In the Zabbix user interface, go to *Data collection* > *Templates* and import the local file.

> **Note:**
> Webhook media types can be upgraded similarly: get them from the `media` directory and import in *Alerts* > *Media types*.

To upgrade or add **multiple templates**:

1. Download the template bundles archive depending on your web server:
   - Nginx - `templates_bundles-8.2-1M-nginx.zip`
   - Apache - `templates_bundles-8.2-2M-apache.zip`
2. Extract the archive. It has YAML files, each with multiple templates, and an index file listing which templates are in each YAML file.
3. In the Zabbix user interface, go to *Data collection* > *Templates* and import the YAML files one by one.

> **Note:**
> Webhook media types can be upgraded similarly: take the `media_bundle-8.2-000.yaml` file from the template bundles archive and import in *Alerts* > *Media types*.

To upgrade or add **all templates** at once, you can use the `import_templates.sh` script included in the template bundles archive. The script requires `curl` and `jq` utilities, and uses the `configuration.import` Zabbix API method. When running the script, provide the directory containing the YAML files and your Zabbix API URL and authentication token.

> **Warning:**
> The `import_templates.sh` script overwrites all existing templates, with no option to review or undo template changes.

Example:

```
/tmp/templates_bundles-8.2-2M-apache/import_templates.sh \
  /tmp/templates_bundles-8.2-2M-apache \
  https://example.com/zabbix/api_jsonrpc.php \
  d8d6b5c78ee2a8333db11bfcd3e448c49c239d8ce00702759e2b56f42e24dbe4
```

Template compatibility with Host Wizard

The Host Wizard is compatible with all out-of-the-box templates.

You can also make your custom templates compatible with the Host Wizard:

1. In Zabbix frontend, go to *Data collection* > *Templates* and export the custom template.
2. Edit the exported file:
   - Add the `wizard_ready` template element to indicate compatibility with the Host Wizard.
   - Optionally, add the `readme` or `config` template element as needed.
3. In Zabbix frontend, go to *Data collection* > *Templates* and import the updated template file.

Zabbix agent template operation

Steps to ensure correct operation of templates that collect metrics with Zabbix agent:

1. Make sure that Zabbix agent is installed on the host. For active checks, also make sure that the Zabbix server/proxy address is added to the 'ServerActive' parameter of the agent configuration file.
2. Link the template to a target host (if the template is not available in your Zabbix installation, you may need to import the template first).
3. If necessary, adjust the values of template macros.
4. Configure the instance being monitored to allow data sharing with Zabbix.

A detailed description of a template, including the full list of macros, items, and triggers, is available in the template's README file (accessible by clicking a template name).

The following templates are available:

- Apache by Zabbix agent
- Apache by Zabbix agent active
- HAProxy by Zabbix agent
- IIS by Zabbix agent
- IIS by Zabbix agent active
- MariaDB by Zabbix agent
- MariaDB by Zabbix agent active
- Microsoft Exchange Server 2016 by Zabbix agent
- Microsoft Exchange Server 2016 by Zabbix agent active
- MySQL by Zabbix agent
- MySQL by Zabbix agent active
- Nginx by Zabbix agent
- Nginx by Zabbix agent active
- Percona by Zabbix agent
- Percona by Zabbix agent active
- PHP-FPM by Zabbix agent
- PHP-FPM by Zabbix agent active
- Podman by Zabbix agent
- Podman by Zabbix agent active
- PostgreSQL by Zabbix agent
- PostgreSQL by Zabbix agent active
- RabbitMQ cluster by Zabbix agent

Zabbix agent 2 template operation

Steps to ensure correct operation of templates that collect metrics with Zabbix agent 2:

1. Make sure that the agent 2 is installed on the host, and that the installed version contains the required plugin. In some cases, you may need to upgrade the agent 2 first.
2. Link the template to a target host (if the template is not available in your Zabbix installation, you may need to import the template first).
3. If necessary, adjust the values of template macros. Note that user macros can be used to override configuration parameters.
4. Configure the instance being monitored to allow data sharing with Zabbix.

> **Attention:**
> Zabbix agent 2 templates work in conjunction with the plugins. While the basic configuration can be done by simply adjusting user macros, the deeper customization can be achieved by configuring the plugin itself. For example, if a plugin supports named sessions, it is possible to monitor several entities of the same kind (e.g., MySQL1 and MySQL2) by specifying named session with own URI, username and password for each entity in the configuration file.

A detailed description of a template, including the full list of macros, items, and triggers, is available in the template's README file (accessible by clicking a template name).

The following templates are available:

- Ceph by Zabbix agent 2
- Docker
- Memcached
- MongoDB cluster by Zabbix agent 2

- MongoDB node by Zabbix agent 2
- MySQL by Zabbix agent 2
- MySQL by Zabbix agent 2 active
- Oracle by Zabbix agent 2
- PostgreSQL by Zabbix agent 2
- PostgreSQL by Zabbix agent 2 active
- Redis
- SMART by Zabbix agent 2
- SMART by Zabbix agent 2 active
- Systemd by Zabbix agent 2
- Website certificate by Zabbix agent 2

HTTP template operation

Steps to ensure correct operation of templates that collect metrics with the HTTP agent:

1. Create a host in Zabbix that corresponds to the monitoring target. Specify the IP address/DNS name of the monitoring target as the main host interface. This is needed for the `{HOST.CONN}` macro to resolve properly in the template items.
2. Link the template to the host created in Step 1 (if the template is not available in your Zabbix installation, you may need to import the template first).
3. If necessary, adjust the values of template-level macros.
4. Configure the instance being monitored to allow data sharing with Zabbix.

A detailed description of a template, including the full list of macros, items and triggers, is available in the template's `README` file (accessible by clicking a template name).

The following templates are available:

- Acronis Cyber Protect Cloud by HTTP
- Apache by HTTP
- Asterisk by HTTP
- AWS by HTTP
- AWS Cost Explorer by HTTP
- AWS EC2 by HTTP
- AWS ECS Cluster by HTTP
- AWS ECS Serverless Cluster by HTTP
- AWS ELB Application Load Balancer by HTTP
- AWS ELB Network Load Balancer by HTTP
- AWS Lambda by HTTP
- AWS RDS instance by HTTP
- AWS S3 bucket by HTTP
- Azure by HTTP
- Cisco Meraki organization by HTTP
- Cisco SD-WAN by HTTP
- Cisco Secure Firewall Threat Defense by HTTP
- Claude API by HTTP
- ClickHouse by HTTP
- Cloudflare by HTTP
- CockroachDB by HTTP
- Control-M enterprise manager by HTTP
- Control-M server by HTTP
- Cradlepoint NCM v2 by HTTP
- DELL PowerEdge R720 by HTTP
- DELL PowerEdge R740 by HTTP
- DELL PowerEdge R820 by HTTP
- DELL PowerEdge R840 by HTTP
- Domain RDAP by HTTP
- Elasticsearch Cluster by HTTP
- Envoy Proxy by HTTP
- Etcd by HTTP
- FortiGate by HTTP
- GitHub repository by HTTP
- GitLab by HTTP
- GLPI by HTTP

- Google Cloud Platform (GCP) by HTTP
- Hadoop by HTTP
- HAProxy by HTTP
- HashiCorp Consul Cluster by HTTP
- HashiCorp Consul Node by HTTP
- HashiCorp Nomad by HTTP
- HashiCorp Vault by HTTP
- Hikvision camera by HTTP
- HPE iLO by HTTP
- HPE MSA 2040 Storage by HTTP
- HPE MSA 2060 Storage by HTTP
- HPE Primera by HTTP
- HPE Synergy by HTTP
- InfluxDB by HTTP
- Jenkins by HTTP
- Kubernetes API server by HTTP, Kubernetes cluster state by HTTP, Kubernetes Controller manager by HTTP, Kubernetes kubelet by HTTP, Kubernetes nodes by HTTP, Kubernetes Scheduler by HTTP are deprecated and have been replaced by the *Kubernetes Cluster by HTTP* and *Kubernetes Node by HTTP* templates.
- Kubernetes Cluster by HTTP
- Kubernetes Node by HTTP
- MantisBT by HTTP
- Microsoft 365 reports by HTTP
- Microsoft SharePoint by HTTP
- NetApp AFF A700 by HTTP
- Nextcloud by HTTP
- NGINX by HTTP
- NGINX Plus by HTTP
- Nutanix Prism Element by HTTP
- OpenAI Platform by HTTP
- OpenStack by HTTP
- OpenWeatherMap by HTTP
- Oracle Cloud by HTTP
- Palo Alto PA-440 by HTTP
- PHP-FPM by HTTP
- Podman by HTTP
- Proxmox VE by HTTP
- Pure Storage FlashArray v1 and v2 by HTTP
- RabbitMQ cluster by HTTP
- Ribbon SBC Edge by HTTP
- Ribbon SBC SWe CE by HTTP
- Ribbon SBC SWe Core by HTTP
- TiDB by HTTP
- TiDB PD by HTTP
- TiDB TiKV by HTTP
- Travis CI by HTTP
- Veeam Backup Enterprise Manager by HTTP
- Veeam Backup and Replication by HTTP
- VeloCloud SD-WAN by HTTP
- VeloCloud SD-WAN Edge by HTTP
- YugabyteDB by HTTP
- ZooKeeper by HTTP

IPMI template operation

IPMI templates do not require any specific setup. To start monitoring, link the template to a target host (if the template is not available in your Zabbix installation, you may need to import the template first).

A detailed description of a template, including the full list of macros, items, and triggers, is available in the template's README file (accessible by clicking a template name).

Available template:

- Chassis by IPMI

JMX template operation

Steps to ensure correct operation of templates that collect metrics by JMX:

1. Make sure that Zabbix Java gateway is installed and set up properly.
2. Link the template to the target host. The host should have the JMX interface set up. If the template is not available in your Zabbix installation, you may need to import the template first.
3. If necessary, adjust the values of template macros.
4. Configure the instance being monitored to allow data sharing with Zabbix.

A detailed description of a template, including the full list of macros, items, and triggers, is available in the template's README file (accessible by clicking a template name).

The following templates are available:

- Apache ActiveMQ by JMX
- Apache Cassandra by JMX
- Apache Kafka by JMX
- Apache Tomcat by JMX
- GridGain by JMX
- Ignite by JMX
- Jira Data Center by JMX
- WildFly Domain by JMX
- WildFly Server by JMX

ODBC template operation

Steps to ensure correct operation of templates that collect metrics via ODBC monitoring:

1. Make sure that required ODBC driver is installed on Zabbix server or proxy.
2. Link the template to a target host (if the template is not available in your Zabbix installation, you may need to import the template first).
3. If necessary, adjust the values of template macros.
4. Configure the instance being monitored to allow data sharing with Zabbix.

A detailed description of a template, including the full list of macros, items, and triggers, is available in the template's README file (accessible by clicking a template name).

The following templates are available:

- MariaDB by ODBC
- MSSQL by ODBC
- MySQL by ODBC
- Oracle by ODBC
- Percona by ODBC
- PostgreSQL by ODBC

Standardized templates for network devices

Overview

In order to provide monitoring for network devices such as switches and routers, we have created two so-called models: for the network device itself (its chassis basically) and for network interface.

Templates for many families of network devices are provided. All templates cover (where possible to get these items from the device):

- Chassis fault monitoring (power supplies, fans and temperature, overall status)
- Chassis performance monitoring (CPU and memory items)
- Chassis inventory collection (serial numbers, model name, firmware version)
- Network interface monitoring with IF-MIB and EtherLike-MIB (interface status, interface traffic load, duplex status for Ethernet)

If the template is not available in your Zabbix installation, you may need to import the template first.

If you are importing the new out-of-the-box templates, you may want to also update the `@Network interfaces for discovery` global regular expression to:

```
Does not match regular expression: ^Software Loopback Interface
Does not match regular expression: ^(In)?[lL]oop[bB]ack[0-9._]*$
Does not match regular expression: ^NULL[0-9.]*$
Does not match regular expression: ^[lL]o[0-9.]*$
Does not match regular expression: ^[sS]ystem$
Does not match regular expression: ^Nu[0-9.]*$
```

to filter out loopbacks and null interfaces on most systems.

Devices

List of device families for which templates are available:

Template name	Vendor	Device family	Known models	OS	MIBs used	**Tags**
Alcatel Timetra TiMOS SNMP	Alcatel	Alcatel Timetra	ALCATEL SR 7750	TiMOS	TIMETRA-SYSTEM-MIB,TIMETRA-CHASSIS-MIB	Certified
Aruba CX 8300s by SNMP	HPE Aruba Network-ing	Aruba CX Switch Series	Aruba JL636A 8325, Aruba JL717A 8360	Aruba AOS-CX	ARUBAWIRED-FAN-MIB, ARUBAWIRED-POWERSUPPLY-MIB, ARUBAWIRED-SYSTEMINFO-MIB, ARUBAWIRED-TEMPSENSOR-MIB, OSPF-MIB	Certified
Brocade FC SNMP	Brocade	Brocade FC switches	Brocade 300 SAN Switch-	-	SW-MIB,ENTITY-MIB	Performance Fault
Brocade_Foundry Stackable SNMP	Brocade	Brocade ICX	Brocade ICX6610, Brocade ICX7250-48, Brocade ICX7450-48F		FOUNDRY-SN-AGENT-MIB, FOUNDRY-SN-STACKING-MIB	Certified
Brocade_Foundry Nonstackable SNMP	Brocade, Foundry	Brocade MLX, Foundry	Brocade MLXe, Foundry FLS648, Foundry FWSX424		FOUNDRY-SN-AGENT-MIB	Performance Fault
Check Point Next Generation Firewall by SNMP	Check Point	Next Genera-tion Firewall	-	Gaia	HOST-RESOURCES-MIB, CHECKPOINT-MIB, UCD-SNMP-MIB, SNMPv2-MIB, IF-MIB	Certified
Ciena 3906 by SNMP	Ciena	Ciena CPE	Ciena 3906	SAOS	WWP-LEOS-BLADE-MIB, WWP-LEOS-CHASSIS-MIB, WWP-LEOS-SYSTEM-CONFIG-MIB	Certified
Cisco Catalyst 3750<device model> SNMP	Cisco	Cisco Catalyst 3750	Cisco Catalyst 3750V2-24FS, Cisco Catalyst 3750V2-24PS, Cisco Catalyst 3750V2-24TS, Cisco Catalyst SNMP, Cisco Catalyst SNMP		CISCO-MEMORY-POOL-MIB, IF-MIB, EtherLike-MIB, SNMPv2-MIB, CISCO-PROCESS-MIB, CISCO-ENVMON-MIB, ENTITY-MIB	Certified

Template name	Vendor	Device family	Known models	OS	MIBs used	Tags
Cisco IOS SNMP	Cisco	Cisco IOS ver > 12.2 3.5	Cisco C2950	IOS	CISCO-PROCESS-MIB,CISCO-MEMORY-POOL-MIB,CISCO-ENVMON-MIB	Certified
Cisco IOS versions 12.0_3_T-12.2_3.5 SNMP	Cisco	Cisco IOS > 12.0 3 T and 12.2 3.5	-	IOS	CISCO-PROCESS-MIB,CISCO-MEMORY-POOL-MIB,CISCO-ENVMON-MIB	Certified
Cisco IOS prior to 12.0_3_T SNMP	Cisco	Cisco IOS 12.0 3 T	-	IOS	OLD-CISCO-CPU-MIB,CISCO-MEMORY-POOL-MIB	Certified
D-Link DES_DGS Switch SNMP	D-Link	DES/DGX switches	D-Link DES-xxxx/DGS-xxxx,DLINK DGS-3420-26SC	-	DLINK-AGENT-MIB,EQUIPMENT-MIB,ENTITY-MIB	Certified
D-Link DES 7200 SNMP	D-Link	DES-7xxx	D-Link DES 7206	-	ENTITY-MIB,MY-SYSTEM-MIB,MY-PROCESS-MIB,MY-MEMORY-MIB	Performance Fault Interfaces
Dell Force S-Series SNMP	Dell	Dell Force S-Series	S4810		F10-S-SERIES-CHASSIS-MIB	Certified
Extreme Exos SNMP	Extreme	Extreme EXOS	X670V-48x	EXOS	EXTREME-SYSTEM-MIB,EXTREME-SOFTWARE-MONITOR-MIB	Certified
FortiGate by SNMP	Fortinet	FortiGate (NGFW)	-	FortiOS	HOST-RESOURCES-MIB FORTINET-FORTIGATE-MIB FORTINET-CORE-MIB SNMPv2-MIB IF-MIB ENTITY-MIB	Performance Inventory
HP Comware HH3C SNMP	HP	HP (H3C) Comware	HP A5500-24G-4SFP HI Switch		HH3C-ENTITY-EXT-MIB,ENTITY-MIB	Certified
HP Enterprise Switch SNMP	HP	HP Enter-prise Switch	HP ProCurve J4900B Switch 2626, HP J9728A 2920-48G Switch		STATISTICS-MIB,NETSWITCH-MIB,HP-ICF-CHASSIS,ENTITY-MIB,SEMI-MIB	Certified
Huawei AR600 by SNMP	Huawei	Huawei NetEngine AR600 Series Routers	Huawei AR600/AR6100/AR6200/AR6300 series routers with VRP software (V300R019/V300R021/V300R023 or later)		HUAWEI-CBQOS-MIB, HUAWEI-NQA-MIB, ENTITY-MIB, IF-MIB, HOST-RESOURCES-MIB, SNMPv2-MIB	Certified

Template name	Vendor	Device family	Known models	OS	MIBs used	**Tags**
Huawei OceanStor 5300 V5 by SNMP	Huawei	Huawei OceanStor Dorado	Huawei OceanStor Dorado with V5 software		HUAWEI-STORAGE-HARDWARE-MIB, HUAWEI-STORAGE-SPACE-MIB, ISM-PERFORMANCE-MIB, HOST-RESOURCES-MIB, SNMPv2-MIB	Certified
Huawei OceanStor Dorado by SNMP	Huawei	Huawei OceanStor, Huawei OceanStor Dorado	Huawei OceanStor Dorado 3000/5000/6000/8000/18000/5300V5/5500/5600/6800/18500/18800 with V6/V700 software		HUAWEI-STORAGE-HARDWARE-MIB, HUAWEI-STORAGE-SPACE-MIB, ISM-PERFORMANCE-MIB, HOST-RESOURCES-MIB, SNMPv2-MIB	Certified
Huawei VRP by SNMP	Huawei	Huawei VRP	S2352P-EI		ENTITY-MIB,HUAWEI-ENTITY-EXTENT-MIB	Certified
Intel_Qlogic Infiniband SNMP	Intel/QLogic	Intel/QLogic Infini-band devices	Infiniband 12300		ICS-CHASSIS-MIB	Fault Inventory
Juniper SNMP	Juniper	MX,SRX,EX models	Juniper MX240, Juniper EX4200-24F	JunOS	JUNIPER-MIB	Certified
Juniper MX NETCONF	Juniper	MX models	Juniper MX204 Edge Router	JunOS 24.2R1-S1.10	*Not applicable (uses NETCONF over SSH)*	Certified
Juniper MX SNMP	Juniper	MX models	Juniper MX204 Edge Router	JunOS 24.2R1-S1.10	OSPF-MIB, JUNIPER-DOM-MIB, JUNIPER-MIB, BGP4-V2-MIB-JUNIPER, OSPFV3-MIB-JUNIPER	Certified
Mellanox SNMP	Mellanox	Mellanox Infini-band devices	SX1036	MLNX-OS	HOST-RESOURCES-MIB,ENTITY-MIB,ENTITY-SENSOR-MIB,MELLANOX-MIB	Certified

Template name	Vendor	Device family	Known models	OS	MIBs used	**Tags**
MikroTik CCR<device model> SNMP	MikroTik	MikroTik Cloud Core Routers (CCR series)	Separate dedicated templates are available for MikroTik CCR1009-7G-1C-1S+, MikroTik CCR1009-7G-1C-1S+PC, MikroTik CCR1009-7G-1C-PC, MikroTik CCR1016-12G, MikroTik CCR1016-12S-1S+, MikroTik CCR1036-12G-4S-EM, MikroTik CCR1036-12G-4S, MikroTik CCR1036-8G-2S+, MikroTik CCR1036-8G-2S+EM, MikroTik CCR1072-1G-8S+, MikroTik CCR2004-16G-2S+, MikroTik CCR2004-1G-12S+2XS	RouterOS	MIKROTIK-MIB,HOST-RESOURCES-MIB	Certified

Template name	Vendor	Device family	Known models	OS	MIBs used	Tags
MikroTik CRS<device model> SNMP	MikroTik	MikroTik Cloud Router Switches (CRS series)	Separate dedicated templates are available for MikroTik CRS106-1C-5S, MikroTik CRS109-8G-1S-2HnD-IN, MikroTik CRS112-8G-4S-IN, MikroTik CRS112-8P-4S-IN, MikroTik CRS125-24G-1S-2HnD-IN, MikroTik CRS212-1G-10S-1S+IN, MikroTik CRS305-1G-4S+IN, MikroTik CRS309-1G-8S+IN, MikroTik CRS312-4C+8XG-RM, MikroTik CRS317-1G-16S+RM, MikroTik CRS326-24G-2S+IN, MikroTik CRS326-24G-2S+RM, MikroTik CRS326-24S+2Q+RM, MikroTik CRS328-24P-4S+RM, MikroTik CRS328-4C-20S-4S+RM, MikroTik CRS354-48G-4S+2Q+RM, MikroTik CRS354-48P-4S+2Q+RM	RouterOS/SwitchOS	MIKROTIK-MIB,HOST-RESOURCES-MIB	Certified
MikroTik CSS<device model> SNMP	MikroTik	MikroTik Cloud Smart Switches (CSS series)	Separate dedicated templates are available for MikroTik CSS326-24G-2S+RM, MikroTik CSS610-8G-2S+IN	RouterOS	MIKROTIK-MIB,HOST-RESOURCES-MIB	Certified
MikroTik FiberBox SNMP	MikroTik	MikroTik FiberBox	MikroTik FiberBox	RouterOS	MIKROTIK-MIB,HOST-RESOURCES-MIB	Certified
MikroTik hEX <device model> SNMP	MikroTik	MikroTik hEX	Separate dedicated templates are available for MikroTik hEX, MikroTik hEX lite, MikroTik hEX PoE, MikroTik hEX PoE lite, MikroTik hEX S	RouterOS	MIKROTIK-MIB,HOST-RESOURCES-MIB	Certified

Template name	Vendor	Device family	Known models	OS	MIBs used	**Tags**
MikroTik netPower <device model> SNMP	MikroTik	MikroTik net-Power	Separate dedicated templates are available for MikroTik netPower 15FR, MikroTik netPower 16P SNMP, MikroTik netPower Lite 7R	RouterOS/SwitchOS, SwitchOS Lite	MIKROTIK-MIB,HOST-RESOURCES-MIB	Certified
MikroTik PowerBox <device model> SNMP	MikroTik	MikroTik Power-Box	Separate dedicated templates are available for MikroTik PowerBox, MikroTik PowerBox Pro	RouterOS	MIKROTIK-MIB,HOST-RESOURCES-MIB	Certified
MikroTik RB<device model> SNMP	MikroTik	MikroTik RB series routers	Separate dedicated templates are available for MikroTik RB1100AHx4, MikroTik RB1100AHx4 Dude Edition, MikroTik RB2011iL-IN, MikroTik RB2011iL-RM, MikroTik RB2011iLS-IN, MikroTik RB2011UiAS-IN, MikroTik RB2011UiAS-RM, MikroTik RB260GS, MikroTik RB3011UiAS-RM, MikroTik RB4011iGS+RM, MikroTik RB5009UG+S+IN	RouterOS	MIKROTIK-MIB,HOST-RESOURCES-MIB	Certified
MikroTik SNMP	MikroTik	MikroTik RouterOS devices	MikroTik CCR1016-12G, MikroTik RB2011UAS-2HnD, MikroTik 912UAG-5HPnD, MikroTik 941-2nD, MikroTik 951G-2HnD, MikroTik 1100AHx2	RouterOS	MIKROTIK-MIB,HOST-RESOURCES-MIB	Certified
Netgear Fastpath SNMP	Netgear	Netgear Fastpath	M5300-28G		FASTPATH-SWITCHING-MIB,FASTPATH-BOXSERVICES-PRIVATE-MIB	Fault Inventory

Template name	Vendor	Device family	Known models	OS	MIBs used	**Tags**
QTech QSW SNMP	QTech	Qtech devices	Qtech QSW-2800-28T	-	QTECH-MIB,ENTITY-MIB	Performance Inventory
Stormshield SNS by SNMP	Stormshield	Stormshield Network Security (SNS) firewalls	SN3100	Stormshield SNS	HOST-RESOURCES-MIB, UCD-SNMP-MIB, STORMSHIELD-ASQ-STATS-MIB, STORMSHIELD-AUTOUPDATE-MIB, STORMSHIELD-HA-MIB, STORMSHIELD-PROPERTY-MIB, STORMSHIELD-HEALTH-MONITOR-MIB, STORMSHIELD-IF-MIB, STORMSHIELD-SYSTEM-MONITOR-MIB, STORMSHIELD-IPSEC-STATS-MIB	Certified
TP-LINK SNMP	TP-LINK	TP-LINK	T2600G-28TS v2.0		TPLINK-SYSMONITOR-MIB,TPLINK-SYSINFO-MIB	Performance Inventory
Ubiquiti AirOS SNMP	Ubiquiti	Ubiquiti AirOS wireless devices	NanoBridge,NanoStation,Unifi	AirOS	FROGFOOT-RESOURCES-MIB,IEEE802dot11-MIB	Performance
Vyatta Virtual Router by SNMP	Ciena	Vyatta	Vyatta Virtual Router 1908e	Vyatta 1908e	SNMPv2-MIB, HOST-RESOURCES-MIB, UCD-SNMP-MIB, IF-MIB, DISMAN-EVENT-MIB	Performance Inventory

Template design

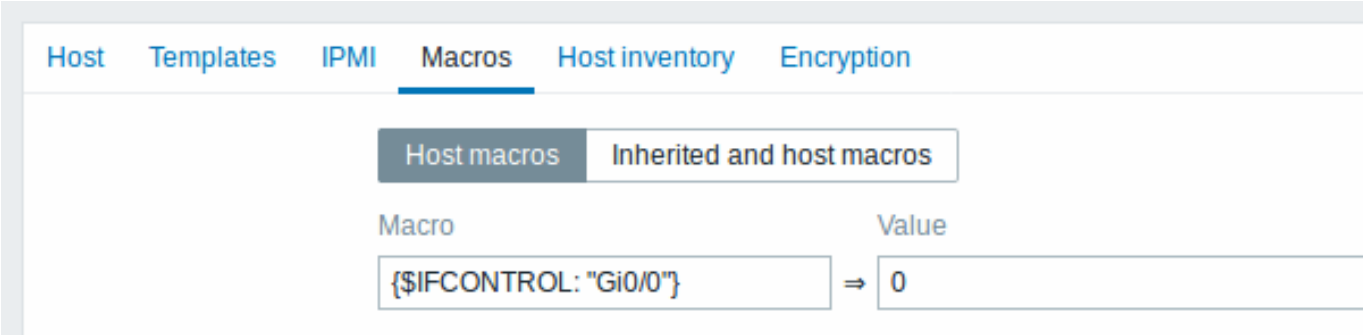
Templates were designed with the following in mind:

- User macros are used as much as possible so triggers can be tuned by the user;
- Low-level discovery is used as much as possible to minimize the number of unsupported items;
- All templates depend on Template ICMP Ping so all devices are also checked by ICMP;
- Items don't use any MIBs - SNMP OIDs are used in items and low-level discoveries. So it's not necessary to load any MIBs into Zabbix for templates to work;
- Loopback network interfaces are filtered when discovering as well as interfaces with ifAdminStatus = down(2)
- 64bit counters are used from IF-MIB::ifXTable where possible. If it is not supported, default 32bit counters are used instead.

All discovered network interfaces have a trigger that monitors its operational status (link), for example:

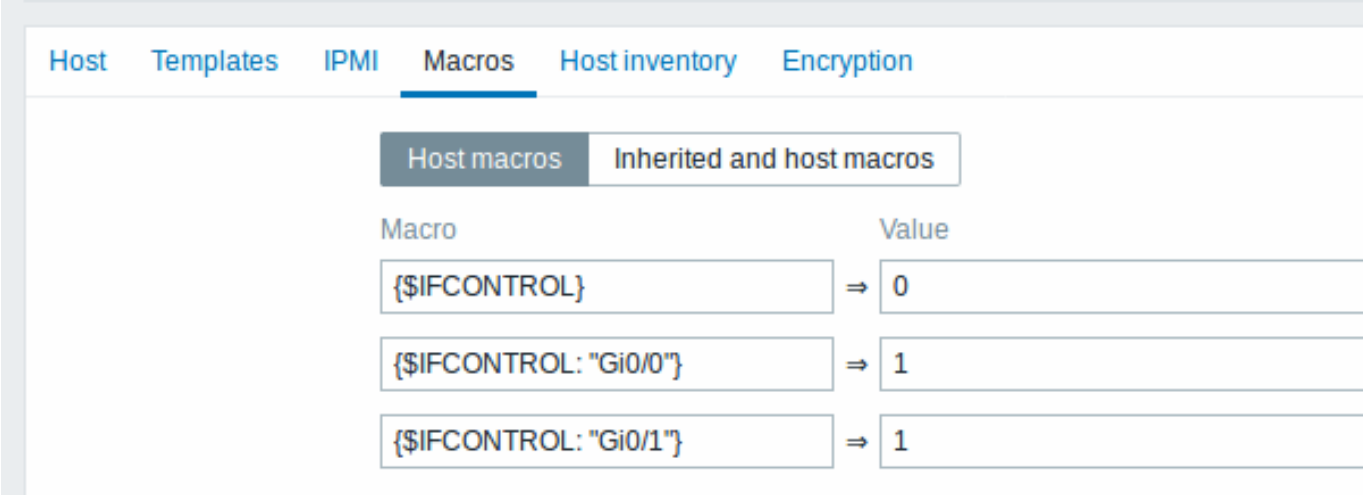
```
{$IFCONTROL:"{#IFNAME}"}=1 and last(/Alcatel Timetra TiMOS SNMP/net.if.status[ifOperStatus.{#SNMPINDEX}])
```

- If you do no want to monitor this condition for a specific interface create a user macro with context with the value 0. For example:



where Gi0/0 is {#IFNAME}. That way the trigger is not used any more for this specific interface.

- You can also change the default behavior for all triggers not to fire and activate this trigger only to limited number of interfaces like uplinks:



Tags

- Performance - device family MIBs provide a way to monitor CPU and memory items;
- Fault - device family MIBs provide a way to monitor at least one temperature sensor;
- Inventory - device family MIBs provide a way to collect at least the device serial number and model name;
- Certified - all three main categories above are covered.

VMware template operation

Overview

Zabbix includes a range of ready-to-use templates for virtual machine monitoring, designed for VMware vCenter and ESXi hypervisors. The available templates are contained in two separate template sets:

- VMware - uses UUID data for corresponding macros
- VMware FQDN - uses FQDN data for corresponding macros

The templates designed for monitoring VMware vCenter or ESXi hypervisors contain pre-configured low-level discovery rules along with various built-in checks for monitoring virtual installations.

Note:
For the correct functioning of the *VMware FQDN* template, each monitored virtual machine should have a unique OS name adhering to FQDN rules. Additionally, VMware Tools/Open Virtual Machine tools must be installed on every machine. If these prerequisites are met, using the *VMware FQDN* template is recommended. The *VMware FQDN* template has been available since Zabbix 5.2 with the introduction of the ability to create hosts with custom interfaces.

 A classic *VMware* template is also available and can be used if FQDN requirements are unmet. However, the *VMware* template has a known issue. Hosts for discovered virtual machines are created with names that are saved in vCenter (for example, ”VM1”, ”VM2”, etc.). If Zabbix agent is installed on these hosts, and active Zabbix agent autoregistration is enabled, the autoregistration process will read host names as they were registered during launch (for example, ”vm1.example.com”, ”vm2.example.com”, etc.). This can lead to the creation of new hosts for existing virtual machines (since no name matches have been found), resulting in duplicate hosts with different names.

Please note:

- The template *VMware Hypervisor* can be manually linked to a host as well as used in discovery;
- The template *VMware Guest* should not be manually linked to a host and can only be used in discovery.

Host macro configuration

To use VMware simple checks, the host must have the following user macros defined:

- {$VMWARE.URL} - VMware service (vCenter or ESXi hypervisor) SDK URL (https://servername/sdk)
- {$VMWARE.USERNAME} - VMware service user name
- {$VMWARE.PASSWORD} - VMware service {$VMWARE.USERNAME} user password

OTLP template operation

Zabbix provides a range of ready-to-use templates for collecting and monitoring OpenTelemetry traces, metrics, and logs using *Telemetry query* items.

Steps to ensure correct operation of the templates:

1. Check the template's `README` file for any template-specific configuration requirements.
2. Create a host in Zabbix that corresponds to the monitoring target.
3. Ensure the `TelemetryProvider` parameter is configured in the Zabbix server or Zabbix proxy configuration file.
4. Link the template to the created host (if the template is not available in your Zabbix installation, you may need to import the template first).
5. If necessary, adjust the values of template-level macros.

A detailed description of templates, including the full list of macros, items, and triggers, is available in the template's `README` file, which you can access by clicking the template name.

The following templates are available:

- Envoy Proxy by OTLP
- Generic OpenTelemetry by OTLP
- OpenRouter traces by OTLP
- Proxmox VE by OTLP

Notifications upon events

Overview

Assuming that we have configured some items and triggers and now are getting some events happening as a result of triggers changing state, it is time to consider some actions.

To begin with, we would not want to stare at the triggers or events list all the time. It would be much better to receive notification if something significant (such as a problem) has happened. Also, when problems occur, we would like to see that all the people concerned are informed.

That is why sending notifications is one of the primary actions offered by Zabbix. Who and when should be notified upon a certain event can be defined.

To be able to send and receive notifications from Zabbix you have to:

- define some media
- configure an action that sends a message to one of the defined media

Actions consist of *conditions* and *operations*. Basically, when conditions are met, operations are carried out. The two principal operations are sending a message (notification) and executing a remote command.

For discovery and autoregistration created events, some additional operations are available. Those include adding or removing a host, linking a template etc.

Media types

Overview

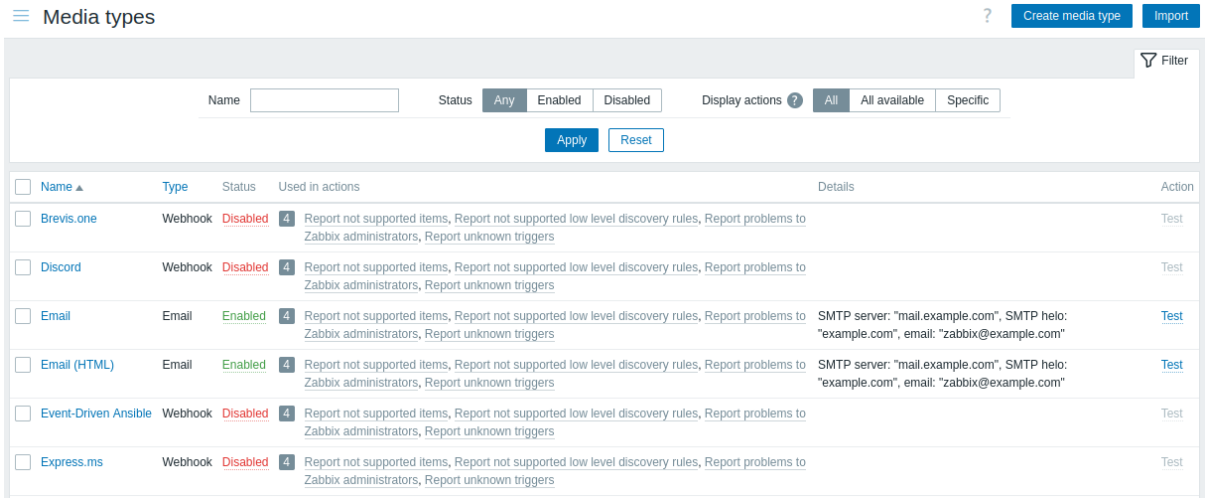
Media types are the delivery channels used for sending notifications and alerts from Zabbix.

Media types support the following delivery methods:

- Email

- Push notification
- SMS
- Custom script
- Webhook

Media types are maintained in *Alerts* > *Media types*. Some media types come pre-defined in the default dataset. You just need to fine-tune their parameters to get them working.



To see how media types fit within the alerting process, let's look at the three requirements to deliver notifications from Zabbix to end users:

1. An action operation must be defined that sends notifications.
2. A working media type must be defined (such as *Email* that send alerts using SMTP).
3. User-level delivery details (such as e-mail addresses, phone numbers, etc) must be defined in user media.

Media type testing

To test if a configured media type works, click *Test* in the media type list.

The testing request will be sent to Zabbix server. Zabbix server will attempt to send an alert using the specified media type and will return the result to frontend. The frontend will wait for the server to return the results. Media type testing has a 65-second timeout by default (configurable in *Administration* > *General* > *Timeouts*).

See also testing details for:

- Email
- Push notification
- Webhook
- Script

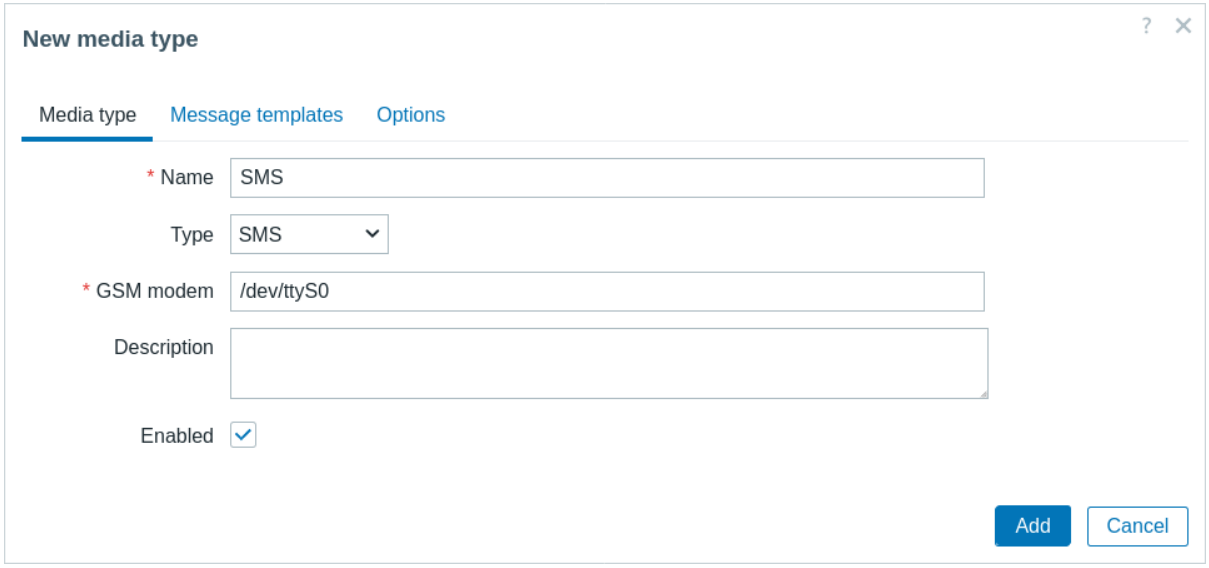
Configuration

To create a media type in Zabbix frontend:

1. Go to *Alerts* > *Media types*.
2. Click *Create media type*.
3. Enter parameters of the media type in the form.

Some parameters are **common** for all delivery methods.

Note:
If one or more delivery methods are specified in the `$ZBX_FEATURE_FLAGS['media_type_denylist']` feature flag in the frontend configuration file (`zabbix.conf.php`), those delivery methods are not available in the *Type* list.

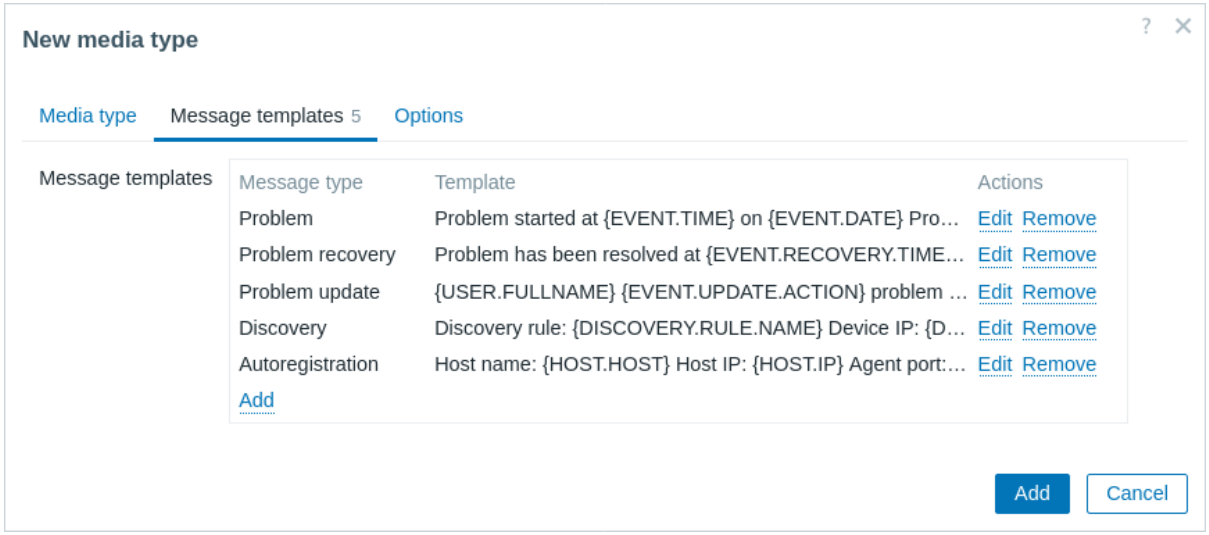


Parameter	Description
Name	Name of the media type.
Type	Select the delivery method for the media type.
Description	Enter a description for the media type.
Enabled	Mark the checkbox to enable the media type.

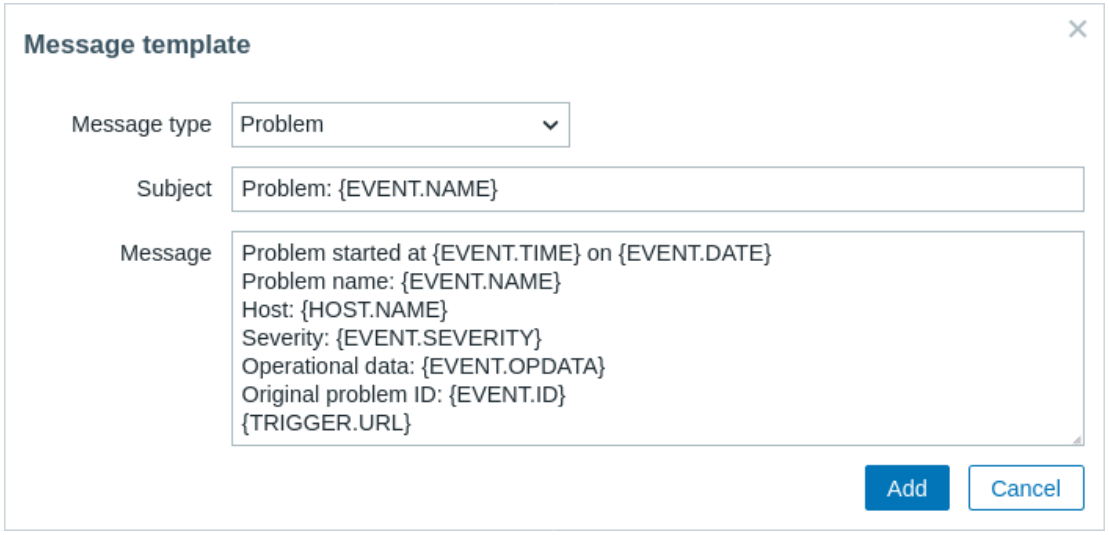
For method-specific parameters, see email, push notification, SMS, custom alertscript, or webhook pages.

Message templates

The **Message templates** tab contains default messages per event type (problem, problem recovery, discovery, etc).



Click **Add** to define a default message (or **Edit** to update an existing message):



Parameter	Description
Message type	Type of an event for which the default message should be used. Only one default message can be defined for each event type.
Subject	Subject of the default message. The subject may contain macros. It is limited to 255 characters. Subject is not available for SMS media type.
Message	The default message. It is limited to certain amount of characters depending on the database type (see Sending messages for more information). The message may contain supported macros. In problem and problem update messages, expression macros are supported (for example, `{?avg(/host/key,1h)}`).

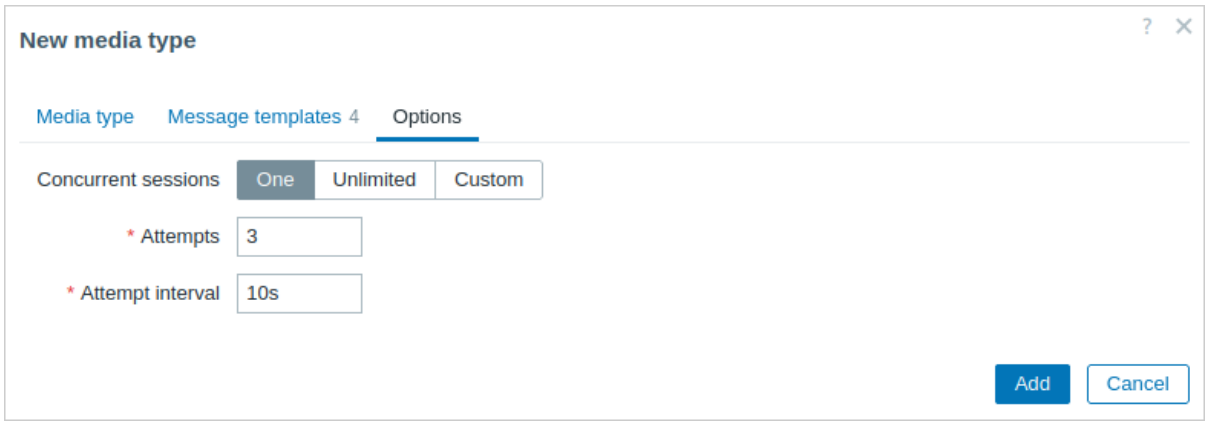
Note that default messages are overridden by custom messages, if defined in action operations.

Warning:
Defining message templates is mandatory for all delivery methods, including webhooks or custom alert scripts that do not use default messages for notifications. For example, the action "Send message to Pushover webhook" will fail to send problem notifications, if the problem message for the Pushover webhook is not defined.

Options

The **Options** tab contains alert processing settings. The same set of options is configurable for each media type.

All media types are processed in parallel. While the maximum number of concurrent sessions is configurable per media type, the total number of alerter processes on the server can only be limited by the StartAlerters parameter. Alerts generated by one trigger are processed sequentially. So multiple notifications may be processed simultaneously only if they are generated by multiple triggers.



Parameter	Description
Concurrent sessions	Select the number of parallel alerter sessions for the media type: **One** - one session **Unlimited** - unlimited number of sessions **Custom** - select a custom number of sessions Unlimited/high values mean more parallel sessions and increased capacity for sending notifications. Unlimited/high values should be used in large environments where lots of notifications may need to be sent simultaneously. If more notifications need to be sent than there are concurrent sessions, the remaining notifications will be queued; they will not be lost.
Attempts	Number of attempts for trying to send a notification. Up to 100 attempts can be specified; the default value is '3'. If '1' is specified, Zabbix will send the notification only once and will not retry if the sending fails.
Attempt interval	Frequency of trying to resend a notification in case the sending failed, in seconds (0-3600). If '0' is specified, Zabbix will retry immediately. Time suffixes are supported (e.g., `30s`, `1m`, `1h`).

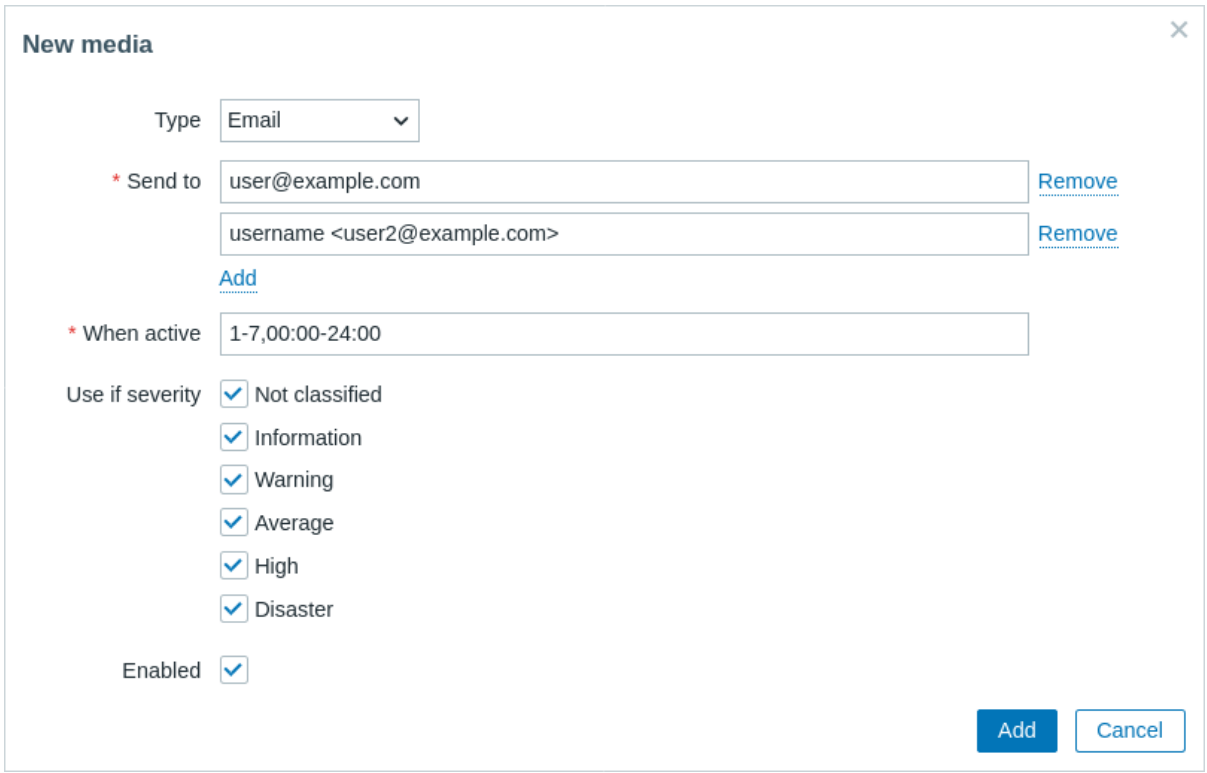
User media

While media types define **how** a notification will be sent, user media define **where** the notification must be sent.

User media (e.g., email address, webhook user ID, etc.) must be defined in the user profile regardless of the delivery method. An action sending messages to *Admin* user using webhook *X* will fail to deliver if the webhook *X* delivery details are not defined in the Admin user profile.

To define user media:

1. Go to *Users* > *Users* and open the user properties form. Alternatively, go to *User settings* > *Notifications* of your own user profile.
2. Click *Add* in the *Media* tab.
3. Enter the user media parameters in the form.



Parameter	Description
Type	The drop-down list contains the names of enabled media types. Note that when editing a medium of a disabled media type, the type will be displayed in red.

Parameter	Description
Send to	Enter the contact information where messages should be sent. For the email media type, multiple addresses can be added by clicking the Add button below the address field. In this case, notifications will be sent to all listed addresses. For address examples, see the *Email* parameter description for the email media type. For the push notification media type, you can select from the user's active mobile devices, or add one or more specific devices.
When active	You can limit the time when messages are sent, for example, set the working days only (1-5,09:00-18:00). Note that this limit is based on the user time zone. If the user time zone is changed and is different from the system time zone this limit may need to be adjusted accordingly so as not to miss important messages. See the Time period specification page for description of the format. User macros are supported.
Use if severity	Mark the checkboxes of trigger severities that you want to receive notifications for. *Note* that the default severity ('Not classified') **must be** checked if you want to receive notifications for non-trigger events. After saving, the selected trigger severities will be displayed in the corresponding severity colors, while unselected ones will be grayed out.
Status	Status of the user media. **Enabled** - is in use. **Disabled** - is not being used.

Email

Overview

To configure email as the delivery channel for messages, you need to configure email as the media type and assign specific addresses to users.

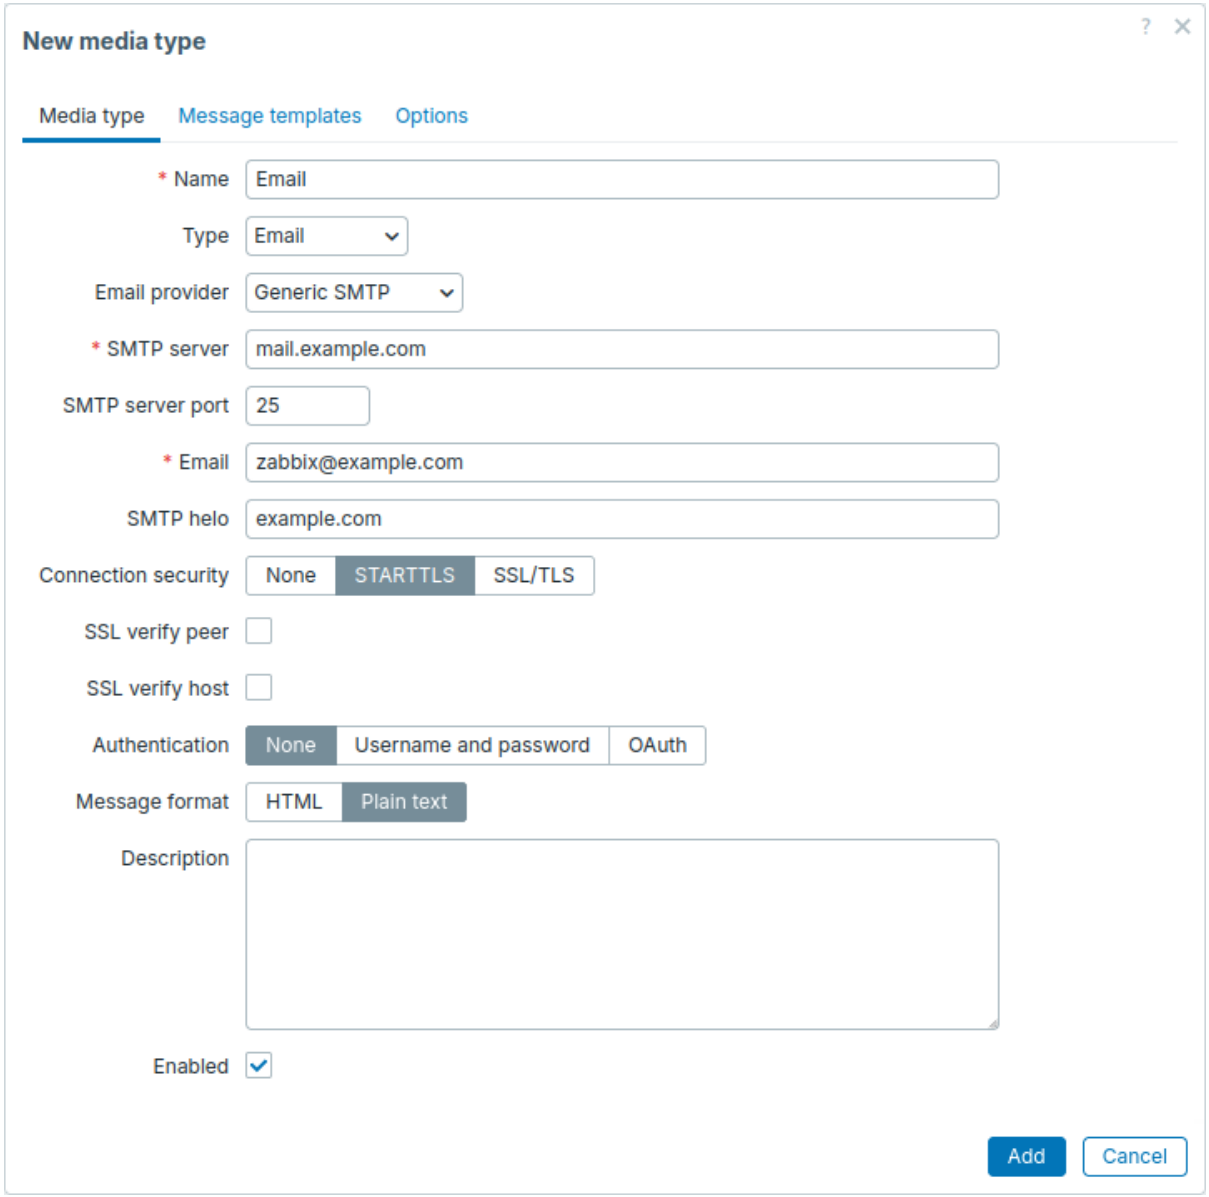
> **Note:**
> Multiple notifications for single event will be grouped together on the same email thread.

Configuration

To configure email as the media type:

1. Go to *Alerts* > *Media types*.
2. Click *Create media type* (or click *Email* in the list of pre-defined media types).
3. Enter the email media type parameters in the form.

The **Media type** tab contains general media type attributes:



All mandatory input fields are marked with a red asterisk.

Warning:
Password is exported in clear text when exporting the email media type configuration.

The following parameters are specific for the email media type:

Parameter	Description
Email provider	Select the email provider: *Generic SMTP*, *Gmail*, *Gmail relay*, *Office365*, or *Office365 relay*. If you select the Gmail/Office365-related options, you will only need to supply the sender email address and password; such options as *SMTP server*, *SMTP server port*, *SMTP helo*, and *Connection security* will be automatically filled by Zabbix. See also: Automated Gmail/Office365 media types.
SMTP server	Set an SMTP server to handle outgoing messages. This field is available if *Generic SMTP* is selected as the email provider.
SMTP server port	Set the SMTP server port to handle outgoing messages. This field is available if *Generic SMTP* is selected as the email provider.

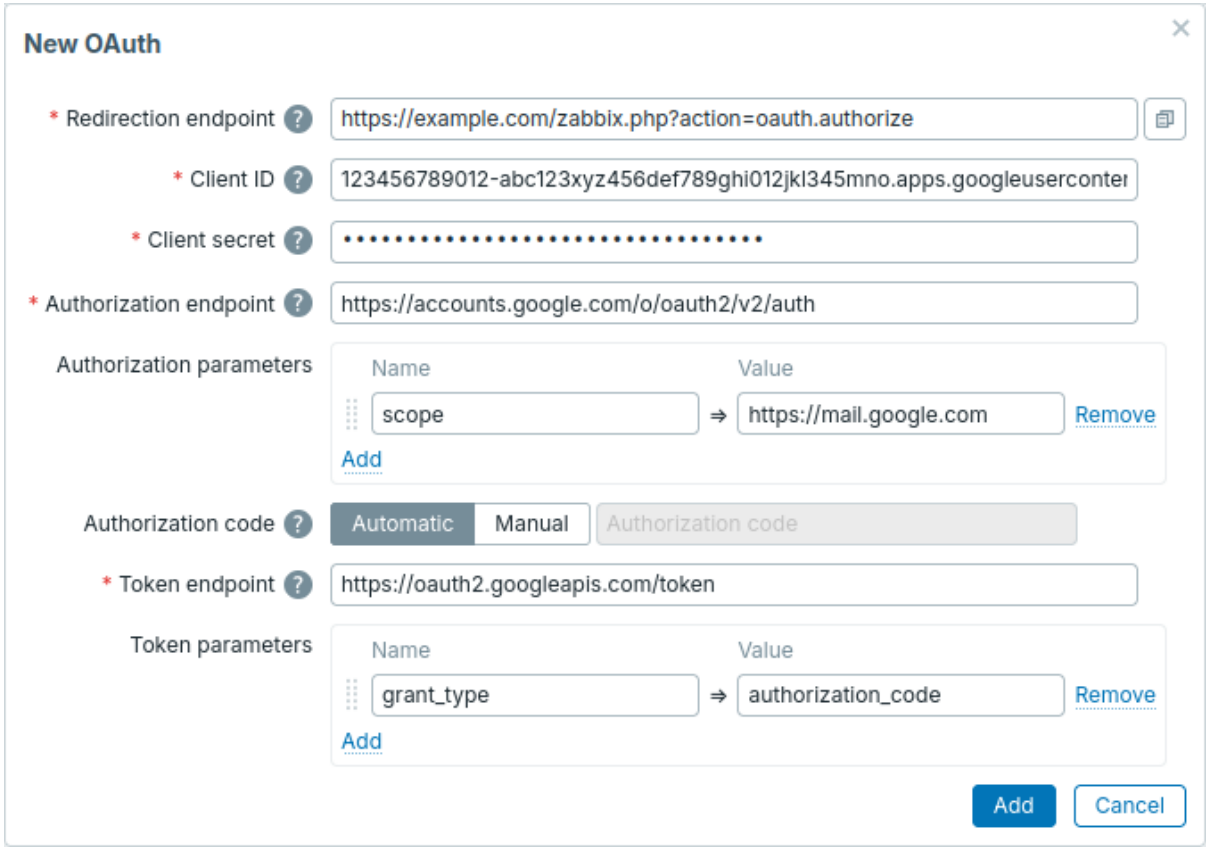
Parameter	Description
Email	The address entered here will be used as the **From** address for the messages sent. Adding a sender display name with the actual email address (for example, `Zabbix_info <zabbix@company.com>`) is supported. There are some restrictions on display names in Zabbix emails compared with those allowed by RFC 5322, as illustrated by the following examples. Valid examples: `zabbix@company.com` (only an email address; angle brackets are not required) `Zabbix_info <zabbix@company.com>` (a display name followed by the email address in angle brackets) `Ω-monitoring <zabbix@company.com>` (UTF-8 characters in the display name) Invalid examples: `Zabbix HQ zabbix@company.com` (a display name present but no angle brackets around the email address) `"Zabbix@<H(comment)Q>" <zabbix@company.com>` (although valid by RFC 5322, quoted pairs and comments are not supported in Zabbix emails).
SMTP helo	Set a correct SMTP helo value, normally a domain name. If empty, the domain name of the email will be sent (i.e., what comes after `@` in the *Email* field). If it is impossible to fetch the domain name, a debug-level warning will be logged and the server hostname will be sent as the domain for HELO command. This field is available if *Generic SMTP* is selected as the email provider.
Connection security	Select the level of connection security: **None** - do not use the CURLOPT_USE_SSL option **STARTTLS** - use the CURLOPT_USE_SSL option with CURLUSESSL_ALL value **SSL/TLS** - use of CURLOPT_USE_SSL is optional.
SSL verify peer	Mark the checkbox to verify the SSL certificate of the SMTP server. The value of "SSLCALocation" server configuration directive should be put into CURLOPT_CAPATH for certificate validation. This sets cURL option CURLOPT_SSL_VERIFYPEER.
SSL verify host	Mark the checkbox to verify that the *Common Name* field or the *Subject Alternate Name* field of the SMTP server certificate matches. This sets cURL option CURLOPT_SSL_VERIFYHOST.
Authentication	Select the level of authentication: **None** - no cURL options are set **Username and password** - implies "AUTH=*" leaving the choice of authentication mechanism to cURL **OAuth** - OAuth authentication OAuth authentication is not supported for the *Office365 relay* email provider.
Username	User name to use in authentication. This sets the value of CURLOPT_USERNAME. User macros supported.
Password	Password to use in authentication. This sets the value of CURLOPT_PASSWORD. User macros supported. The value entered here is exported as is when media type configuration is exported. To avoid exposing sensitive credentials in exported files, use a user macro (for example, `{$EMAIL_SMTP_PASSWORD}`) instead of a literal password—the macro reference will be exported while the secret value should be (re)configured on the destination system.
OAuth tokens	Click *Configure* to set up parameters for retrievieng OAuth tokens in a new window. This field is only available if "OAuth" is selected in the *Authentication* field.
Message format	Select message format: **HTML** - send as HTML; **Plain text** - send as plain text.

Attention:
To enable SMTP authentication, Zabbix server must be compiled with the `--with-libcurl` compilation option, which requires libcurl (see server requirements for version details).

See also common media type parameters for details on how to configure default messages and alert processing options.

OAuth tokens

The following parameters are required for retrieving OAuth tokens:



Warning:
Client secret is exported in clear text when exporting the email media type configuration.

Parameter	Description
Redirection endpoint	Enter Zabbix frontend URL where the OAuth service will redirect back the OAuth authorization (using format `https://<zabbix-frontend-url>/zabbix.php?action=oauth.authorize`). It is set automatically for a new email media type when the frontend URL is defined.
Client ID	Enter unique identifier of the client app registered within the OAuth authorization server.
Client secret	Enter private secret of the client app registered within the OAuth authorization server.
Authorization endpoint	Enter the OAuth authorization server URL for requesting user authorization.
Authorization parameters	Enter parameters for the authorization endpoint.
Authorization code	Enter the authorization code: **Automatic** - the code will be retrieved automatically through a redirection page; **Manual** - enter the code manually if automatic retrieval fails.
Token endpoint	Enter the OAuth authorization server URL to exchange the authorization code for access and refresh tokens.
Token parameters	Enter parameters for the access token.

Note:
The retrieval of OAuth tokens has automated features for **Gmail**, **Gmail relay** and **Office365** email providers. It is only required to supply *Redirection endpoint*, *Client ID* and *Client secret* parameter values. Zabbix will automatically fill the other required values (see OAuth URL defaults by provider).

OAuth access and refresh token retrieval

After submitting OAuth parameters:

1. A browser popup is opened navigating user to the *Authorization endpoint*.
2. User authorizes Zabbix in the OAuth service.

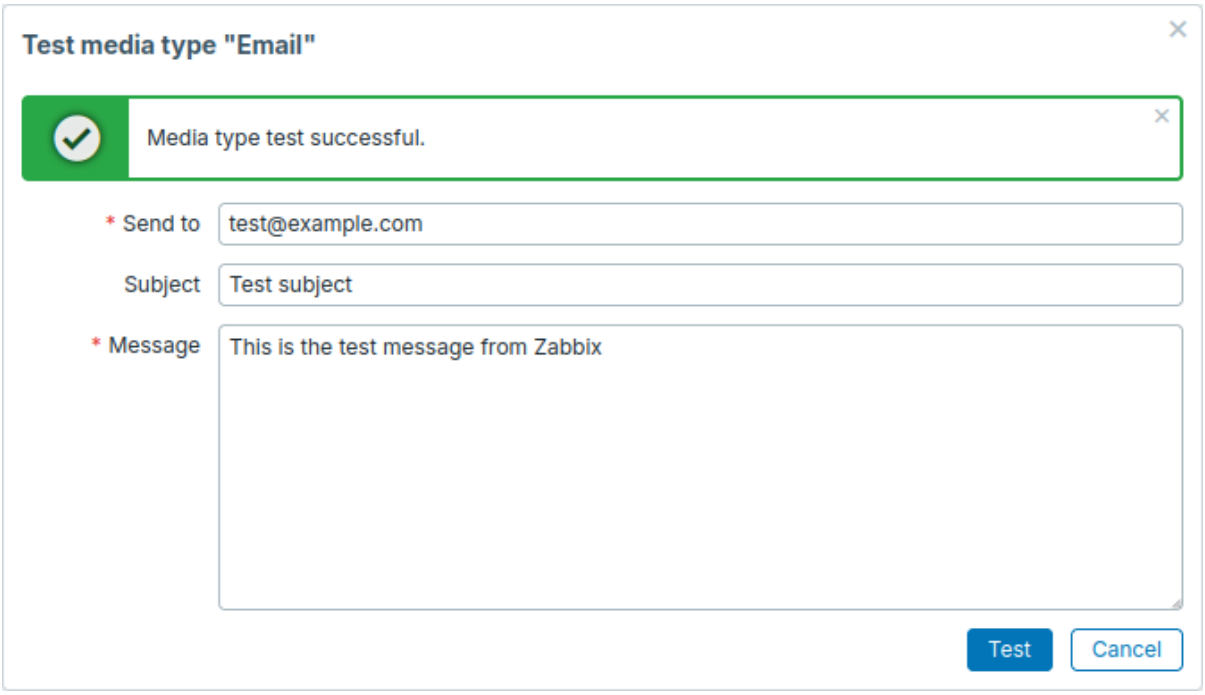
3. The OAuth service redirects user to the Zabbix frontend action `oauth.authorize` with the authorization code and scope value.
4. In response, Zabbix will make a request to *Token endpoint* to exchange the authorization code for access and refresh tokens.

Testing

To test whether a configured email media type works correctly:

1. Locate the relevant email in the list of media types.
2. Click *Test* in the last column of the list (a testing window will open).
3. Enter a *Send to* recipient address, message body and, optionally, subject.
4. Click *Test* to send a test message.

Test success or failure message will be displayed in the same window:



User media

Once the email media type is configured, go to the *Users* > *Users* section and edit user profile to assign email media to the user. Steps for setting up user media, being common for all media types, are described on the Media types page.

Automated Gmail/Office365 media types

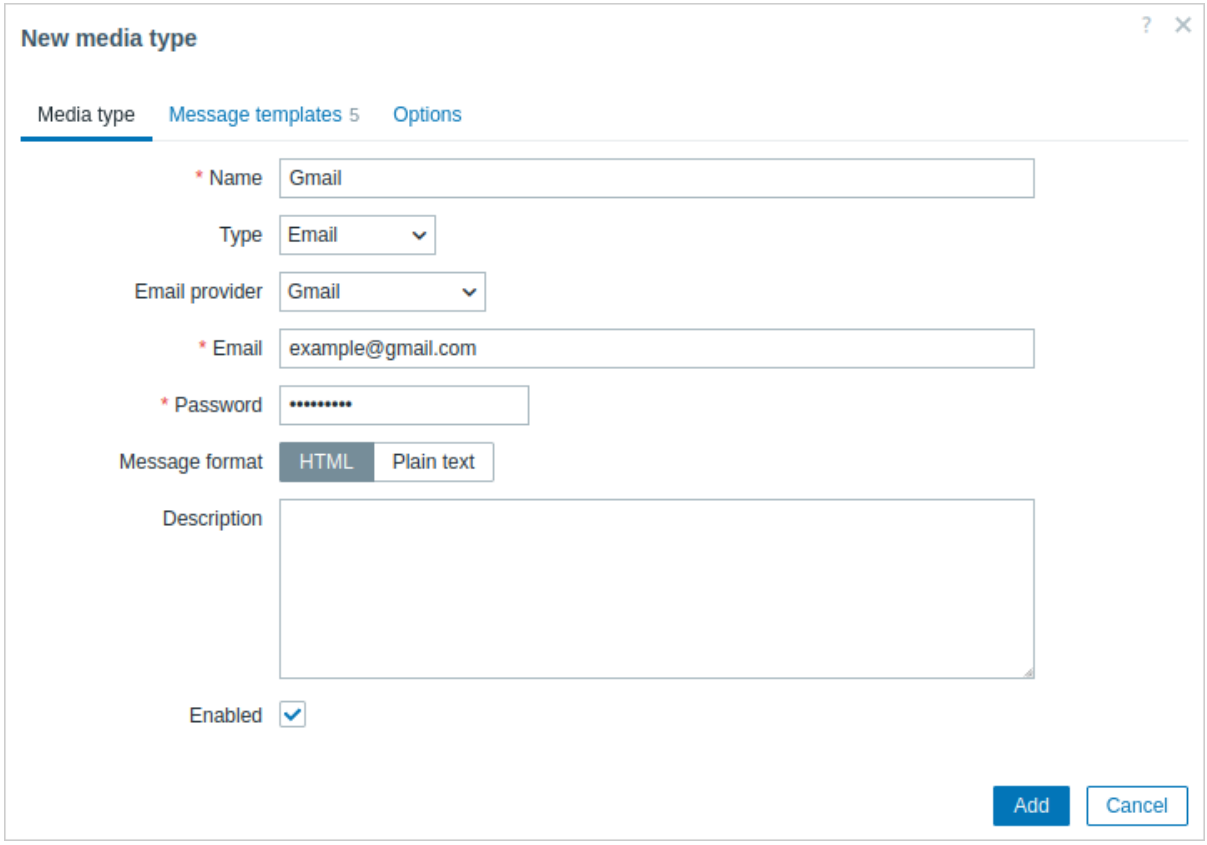
Overview

Gmail or Office365 users may benefit from automated features in media type configuration.

Configuration

The *Email provider* field in the email media type configuration allows you to select pre-configured options for Gmail and Office 365.

When selecting the Gmail/Office365 related options, it is only required to supply the sender email address/password to create a working media type.



As soon as the email address/password is supplied, Zabbix will be able to automatically fill all required settings for Gmail/Office365 media types with the actual/recommended values, i.e., *SMTP server*, *SMTP server port*, *SMTP helo*, and *Connection security*. Because of this automation, these fields are not even shown, however, it is possible to see the SMTP server and email details in the media type list (see the *Details* column).

Note also that:

- The password is not required for the relay options.
- For Office365 relay, the domain name of the provided email address will be used to dynamically fill the SMTP server (i.e., replace "example.com" in `example-com.mail.protection.outlook.com` with the real value).

OAuth tokens

The retrieval of OAuth tokens has automated features for *Gmail*, *Gmail relay* and *Office365* email providers.



It is only necessary to supply a limited set of parameters for OAuth access and refresh token retrieval - *Redirection endpoint*, *Client ID* and *Client secret*. Zabbix will automatically fill the other required values (see OAuth URL defaults by provider).

Note that it is also possible to use the Generic SMTP form for setting up OAuth authorization for these providers.

Note:
SmtpClientAuthentication must be enabled either per user/shared mailbox or per whole tenant in Office365 (not enabled by default settings).

OAuth URL defaults by provider

The following table lists default OAuth URL values and parameters per provider that are used by Zabbix.

Parameter	Gmail	Office365	Generic SMTP
Authorization endpoint	`https://accounts.google.com/o/oauth2/auth`	`https://login.microsoftonline.com/common/oauth2/v2.0/authorize`	No default
scope	`https://mail.google.com`	`https://outlook.office.com/SMTP.Send, offline_access`	No default
access_type	`offline`	Not used	No default
prompt	`consent`	Not used	No default
redirect_uri	Value of *Redirection endpoint* field is added by code automatically. Not listed in *Authorization endpoint* parameters.		
state	Unique hash to identify media being updated with OAuth token is added by code automatically. Not listed in *Authorization endpoint* parameters.		
Token endpoint	`https://oauth2.googleapis.com/token`	`https://login.microsoftonline.com/common/oauth2/v2.0/token`	No default
grant_type	`authorization_code`	`authorization_code`	No default
redirect_uri	Value of *Redirection endpoint* field is added by code automatically. Not listed in *Token endpoint* parameters.		
state	Unique hash to identify media being updated with OAuth token is added by code automatically. Not listed in *Token endpoint* parameters.		

Office365 workaround for personal accounts

Office365 personal accounts do not support query strings in the redirect URL.

To work around this issues the Apache web server administrator can add the following rewrite rule to the `.htaccess` file:

```
RewriteEngine On
RewriteRule ^/oauth\.authorize$ /zabbix.php?action=oauth.authorize [QSA,L,PT]
```

This enables the use of a simplified redirect URL like:

```
http://server-name/zabbix/oauth.authorize
```

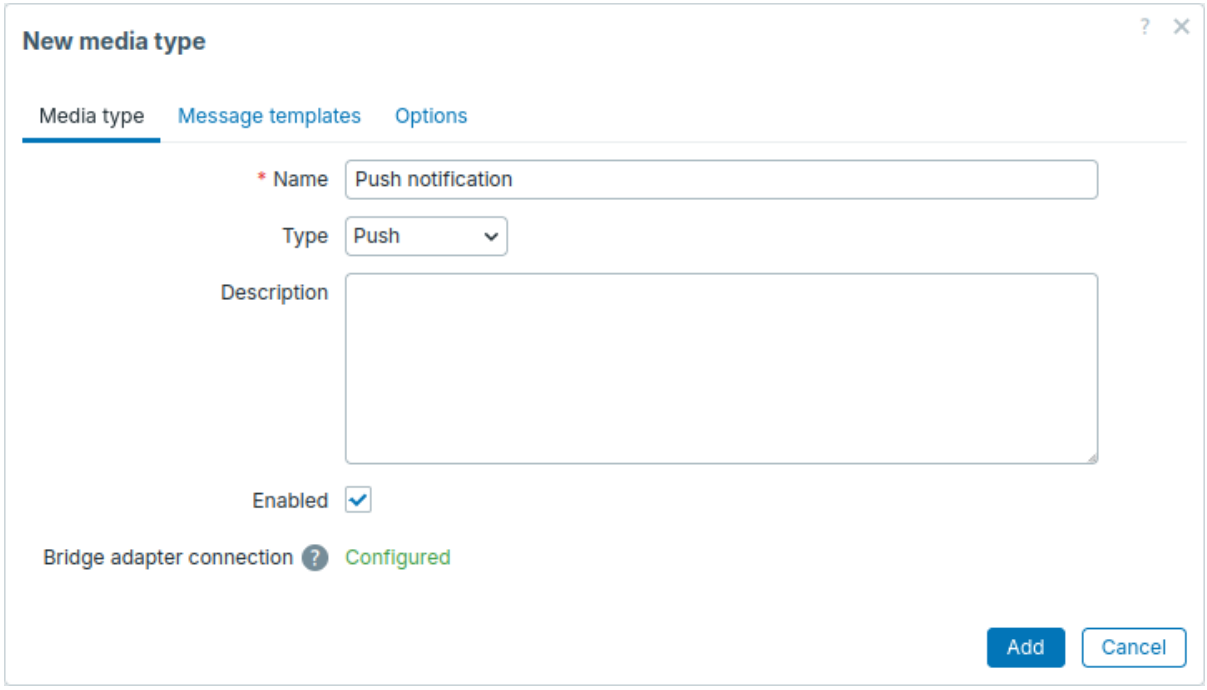
instead of a URL with query parameters, ensuring compatibility with Office365 personal accounts.

Push notification

This media type lets Zabbix server send push notifications to Zabbix mobile app on your mobile device.

Configuration

When creating a media type, select *Push* as its *Type*:



The configuration form contains common media type parameters, message templates, and options.

The form also shows whether Zabbix bridge adapter is configured.
 If the adapter is not configured, you can still create the media type, but it will not work.

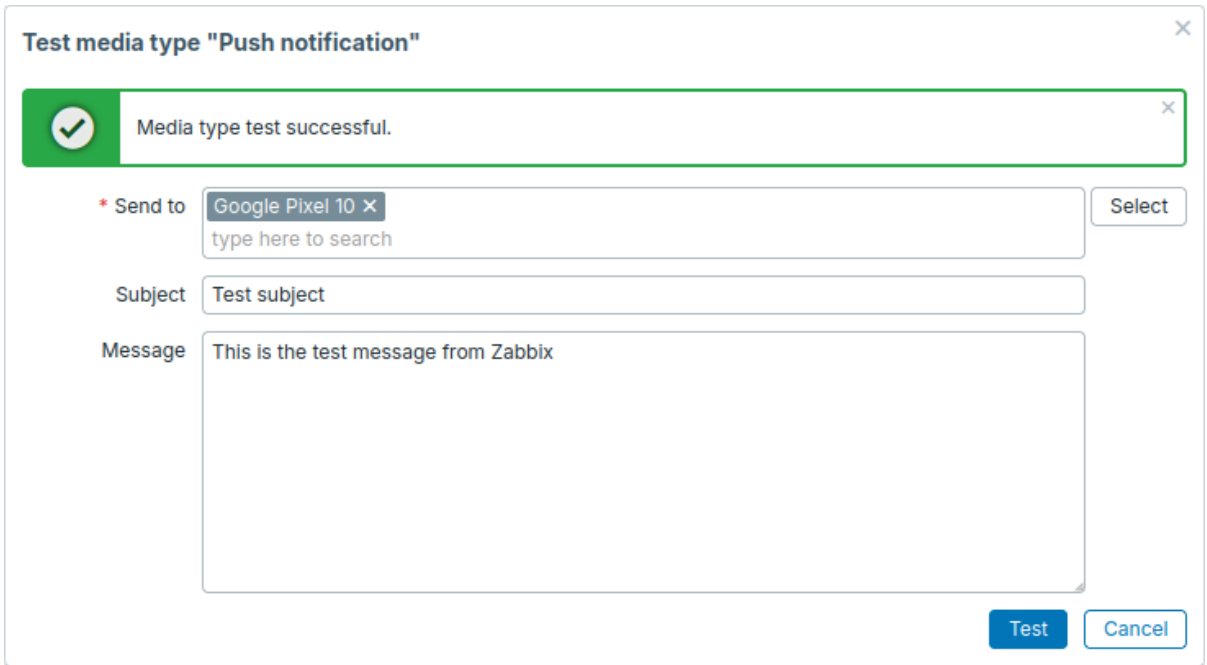
After you have created and enabled the media type, add a device for a user to automatically add the media type to that user (or add user media manually). This lets them receive push notifications on their mobile device.

Testing

To test whether this media type works:

1. Find the media type in the list of media types.
2. Click *Test* in the last column of the list to open the testing form.
3. In the *Send to* field, select a device.
4. Click *Test* to send a test push notification to the selected device.

A success or failure message will appear in the same window:



SMS

Overview

Zabbix supports the sending of SMS messages using a serial GSM modem connected to Zabbix server's serial port.

Make sure that:

- The speed of the serial device (normally /dev/ttyS0 under Linux) matches that of the GSM modem. Zabbix does not set the speed of the serial link. It uses default settings.
- The 'zabbix' user has read/write access to the serial device. Run the command ls -l /dev/ttyS0 to see current permissions of the serial device.
- The GSM modem has PIN entered and it preserves it after power reset. Alternatively you may disable PIN on the SIM card. PIN can be entered by issuing command AT+CPIN="NNNN" (NNNN is your PIN number, the quotes must be present) in a terminal software, such as Unix minicom or Windows HyperTerminal.

Zabbix has been tested with these GSM modems:

- Siemens MC35
- Teltonika ModemCOM/G10

To configure SMS as the delivery channel for messages, you also need to configure SMS as the media type and enter the respective phone numbers for the users.

Configuration

To configure SMS as the media type:

1. Go to *Alerts* > *Media types*.
2. Click *Create media type* (or click *SMS* in the list of pre-defined media types).
3. Enter the SMS media type parameters in the form.

The following parameters are specific for the SMS media type:

Parameter	Description
GSM modem	Set the serial device name of the GSM modem. The path entered here will be validated against the SMSDevices server parameter (if specified).

See common media type parameters for details on how to configure default messages and alert processing options. Note that parallel processing of sending SMS notifications is not possible.

User media

Once the SMS media type is configured, go to the *Users* > *Users* section and edit user profile to assign SMS media to the user. Steps for setting up user media, being common for all media types, are described on the Media types page.

Custom alert scripts

Overview

If you are not satisfied with the existing media types for sending alerts, there is an alternative way to do that. You can create a script that will handle the notification your way.

Custom alert scripts are executed on Zabbix server. These scripts must be located in the directory specified in the server configuration file AlertScriptsPath parameter.

Here is an example of a custom alert script:

```
#####!/bin/bash

to=$1
subject=$2
body=$3
host=$4
value=$5

cat <<EOF | mail -s "$subject" "$to"
$body

Host: $host
Value: $value
EOF
```

Attention:
Zabbix checks for the exit code of the executed commands and scripts. Any exit code, which is different from **0**, is considered as a command execution error. In such cases, Zabbix will try to repeat failed execution.

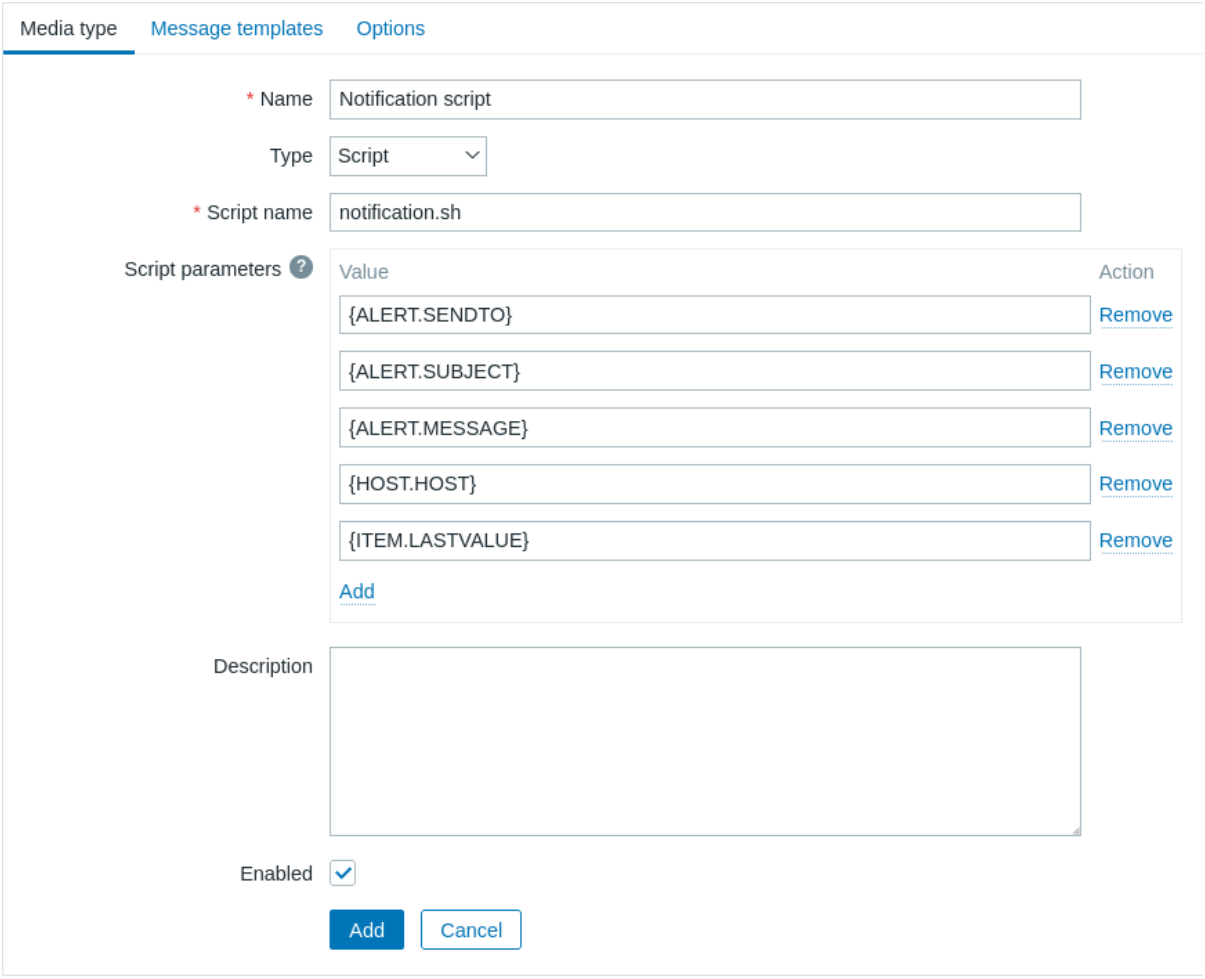
Environment variables are not preserved or created for the script, so they should be handled explicitly.

Configuration

To configure custom alert scripts as a media type:

1. Go to *Alerts* > *Media types*.
2. Click *Create media type*.
3. Enter the scripts media type parameters in the form.

The **Media type** tab contains general media type attributes:



All mandatory input fields are marked with a red asterisk.

The following parameters are specific for the script media type:

Parameter	Description
Script name	Enter the name of the script file (e.g., notification.sh) that is located in the directory specified in the AlertScriptsPath server configuration parameter.
Script parameters	Add optional script parameters that will be passed to the script as command-line arguments in the order in which they are defined. Script parameters support {ALERT.SENDTO}, {ALERT.SUBJECT}, {ALERT.MESSAGE} macros, and all macros that are supported in notifications, as well as user macros.

See common media type parameters for details on how to configure default messages and alert processing options.

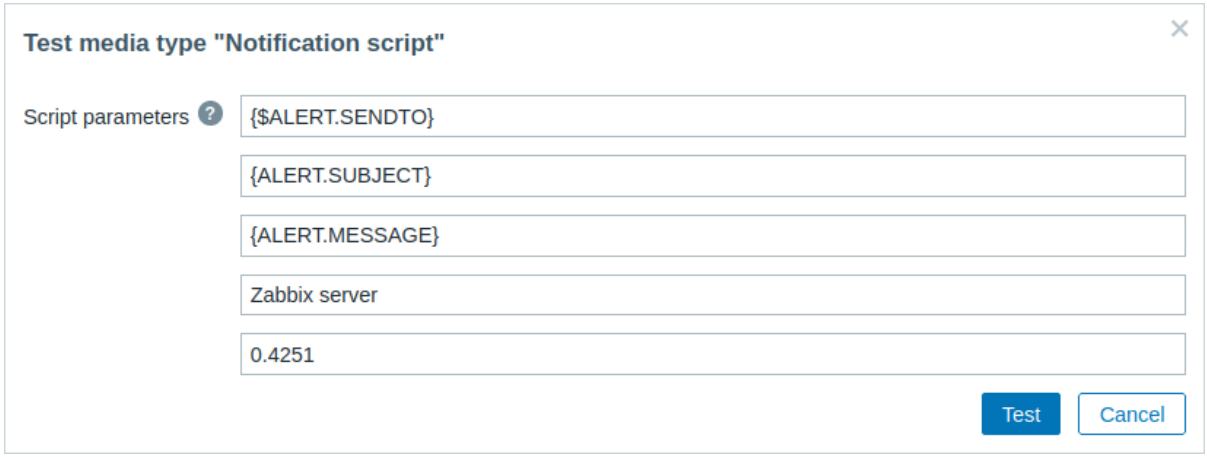
Warning:
Even if an alert script does not use default messages, the message templates for operation types used by this media type must still be defined. Otherwise, a notification will not be sent.

Attention:
If more than one script media type is configured, these scripts may be processed in parallel by the alerter processes. The total number of alerter processes is limited by the server configuration file `StartAlerters` parameter.

Testing

To test a configured script media type:

1. Locate the relevant script in the list of media types.
2. Click *Test* in the last column of the list; a testing form will open in a pop-up window. The testing form will contain the same number of parameters that are configured for the script media type.
3. Edit the script parameter values if needed. Editing only affects the test procedure; the actual values will not be changed.
4. Click *Test*.



Note:
When testing a configured script media type, {ALERT.SENDTO}, {ALERT.SUBJECT}, {ALERT.MESSAGE} and user macros will resolve to their values, but macros that are related to events (e.g., {HOST.HOST}, {ITEM.LASTVALUE}, etc.) will not resolve, as during testing there is no related event to get the details from. Note that macros within {ALERT.SUBJECT} and {ALERT.MESSAGE} macros will also not resolve. For example, if the value of {ALERT.SUBJECT} is composed of "Problem: {EVENT.NAME}" then the {EVENT.NAME} macro will not be resolved.

User media

Once the media type is configured, go to the *Users* > *Users* section and edit a user profile by assigning this media type to the user. Steps for setting up user media, being common for all media types, are described on the Media types page.

Note that when defining the user media, the *Send to* field cannot be empty. If this field is not used in the alert script, enter any combination of supported characters to bypass validation requirements.

Webhook

Overview

The webhook media type is useful for making HTTP calls using custom JavaScript code for straightforward integration with external software such as helpdesk systems, chats, or messengers. You may choose to import an integration provided by Zabbix or create a custom integration from scratch.

Integrations

The following integrations are available, allowing predefined webhook media types to be used to push Zabbix notifications to:

- brevis.one
- Discord
- Event-Driven Ansible
- Express.ms messenger

- GitHub
- GLPI
- IBM Maximo Service Request
- iLert
- iTop
- Jira
- Jira Service Management
- ManageEngine ServiceDesk
- Mantis Bug Tracker
- Mattermost
- MS Teams Workflows
- LINE
- Opsgenie
- OTRS CE
- Pagerduty
- Pushover
- Redmine
- Rocket.Chat
- ServiceNow
- SIGNL4
- Slack
- SolarWinds
- SysAid
- Telegram
- TOPdesk
- VictorOps
- Zammad
- Zendesk

Note:
In addition to the services listed here, Zabbix can be integrated with **Spiceworks** (no webhook is required). To convert Zabbix notifications into Spiceworks tickets, create an email media type and enter Spiceworks helpdesk email address (e.g., help@zabbix.on.spiceworks.com) in the profile settings of a designated Zabbix user.

Configuration

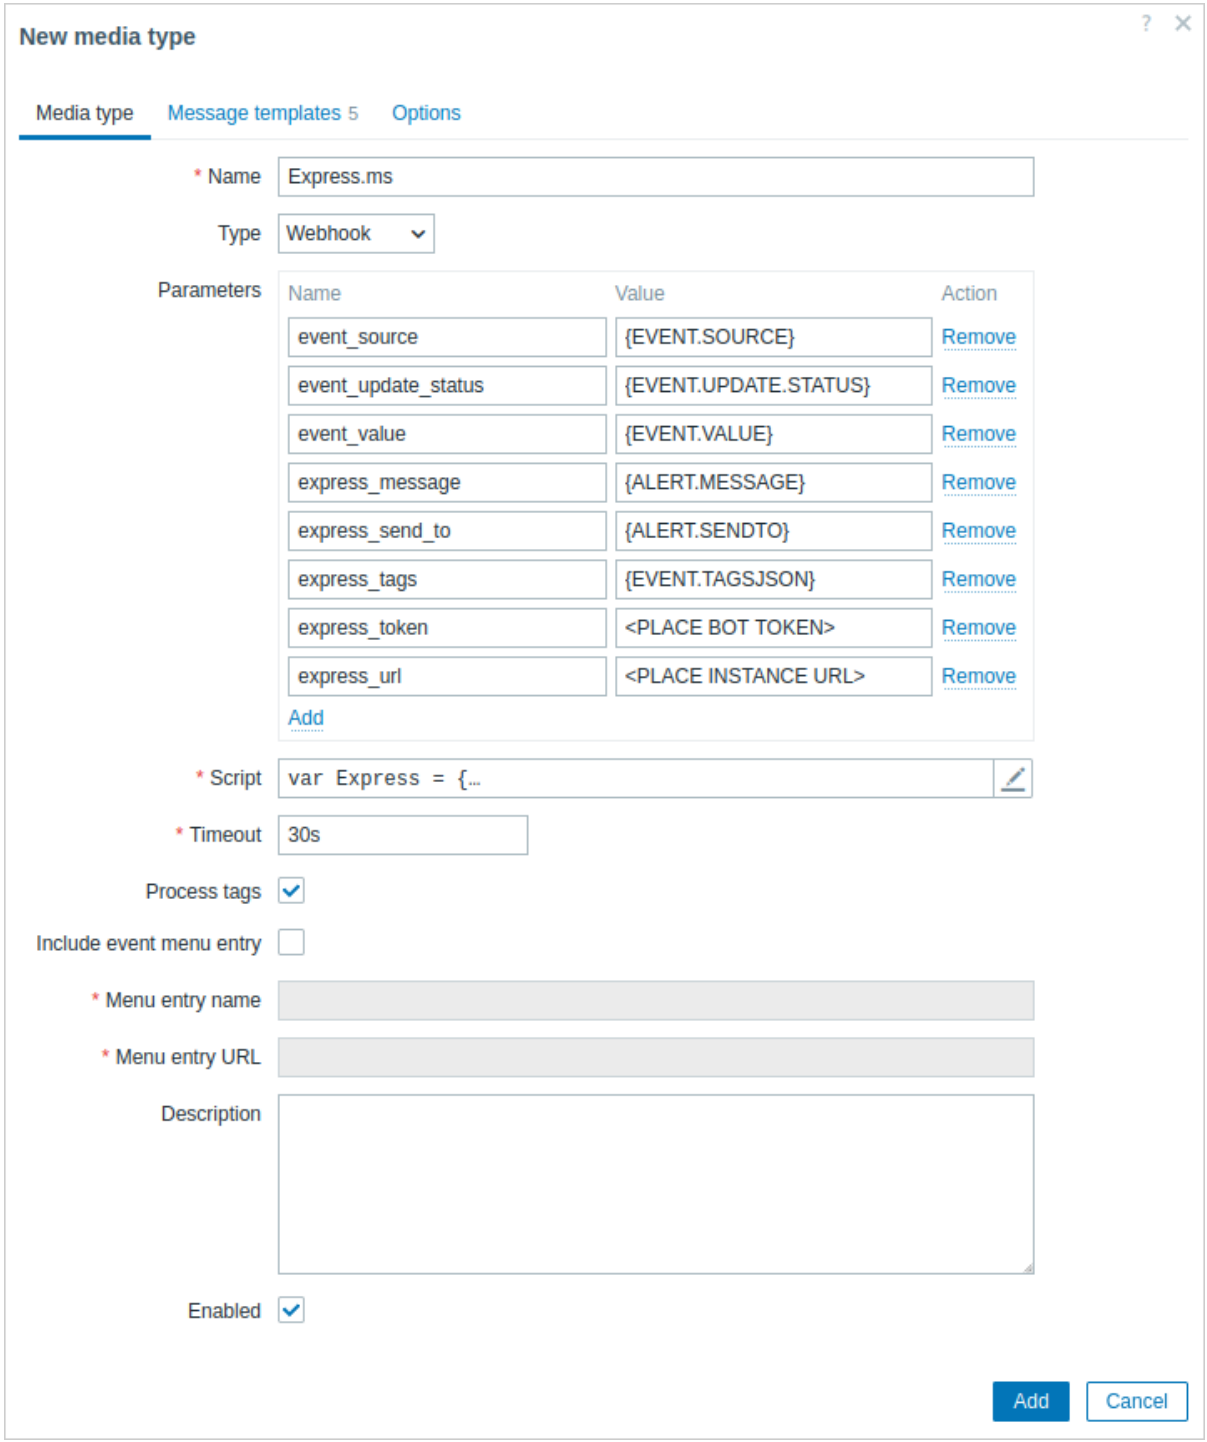
To start using a webhook integration:

1. Locate the required `.yaml` file in the `templates/media` directory of the downloaded Zabbix version or download it from Zabbix git repository.
2. Import the file into your Zabbix installation. The webhook will appear in the list of media types.
3. Configure the webhook according to instructions in the `Readme.md` file (you may click a webhook's name above to quickly access `Readme.md`).

To create a custom webhook from scratch:

1. Go to *Alerts* > *Media types*.
2. Click *Create media type*.
3. Enter the webhook media type parameters in the form.

The **Media type** tab contains various attributes specific for this media type:



All mandatory input fields are marked with a red asterisk.

The following parameters are specific for the webhook media type:

Parameter	Description
Parameters	Specify the webhook variables as the attribute and value pairs. For preconfigured webhooks, a list of parameters varies, depending on the service. Check the webhook's *Readme.md* file for parameter description. For new webhooks, several common variables are included by default (URL:<empty>, HTTPProxy:<empty>, To:{ALERT.SENDTO}, Subject:{ALERT.SUBJECT}, Message:{ALERT.MESSAGE}), feel free to keep or remove them. Webhook parameters support user macros, all macros that are supported in problem notifications and, additionally, {ALERT.SENDTO}, {ALERT.SUBJECT}, and {ALERT.MESSAGE} macros. If you specify an HTTP proxy, the field supports the same functionality as in the item configuration HTTP proxy field. The proxy string may be prefixed with `[scheme]://` to specify which kind of proxy is used (e.g., https, socks4, socks5; see documentation).
Script	Enter JavaScript code in the modal editor that opens when clicking in the parameter field or on the pencil icon next to it. This code will perform the webhook operation. The script is a function code that accepts parameter - value pairs. The values should be converted into JSON objects using JSON.parse() method, for example: `var params = JSON.parse(value);`. The code has access to all parameters, it may perform HTTP GET, POST, PUT and DELETE requests, support additional methods such as CONNECT, PATCH, HEAD, OPTIONS and TRACE, and has control over HTTP headers and request body. The script must contain a return operator; otherwise, it will not be valid. It may return OK status along with an optional list of tags and tag values (see *Process tags* option) or an error string. Recovery events (whether generated automatically or as a result of a manual close) are created by the server and include resolved event tags (including tags inherited from templates, hosts and triggers). Webhook scripts are executed after the alert is created; therefore tags returned by a webhook script are added only after the initial alert creation and will not be present in the {EVENT.TAGS} and {EVENT.RECOVERY.TAGS} macros of the initial problem message or the immediate recovery message. *Note*: Using local variables (e.g., `var local = 1`) instead of a global one (e.g., `global = 1`) is recommended to ensure each script operates on its own data and to avoid collisions between simultaneous calls (see known issues). See also: Webhook development guidelines, Webhook script examples, Additional JavaScript objects.
Timeout	JavaScript execution timeout (1-60s, default 30s). Time suffixes are supported (e.g., 30s, 1m).
Process tags	Mark the checkbox to process returned JSON property values as tags. These tags are added to any existing problem tags. Note that when using webhook tags, the webhook must return a JSON object containing at least an empty tags object: `var result = {tags: {}};` Examples of tags that can be returned: *jira-id:prod-1234*, *responsible:John Smith*, *processed:<no value>*
Include event menu entry	Mark the checkbox to include an entry in the event menu linking to a created external ticket. An entry will be included for each webhook that is enabled and has this checkbox marked. Note that if the *Menu entry name* and *Menu entry URL* parameters contain any {EVENT.TAGS.<tag name>} macros, an entry will be included only if these macros can be resolved (that is, the event has these tags defined). If marked, the webhook should not be used for sending notifications to different users (consider creating a dedicated user instead) and should not be used in multiple alert actions for a single problem event.
Menu entry name	Specify the menu entry name. {EVENT.TAGS.<tag name>} macro is supported. This field is only mandatory if *Include event menu entry* is marked.
Menu entry URL	Specify the underlying URL of the menu entry. {EVENT.TAGS.<tag name>} macro is supported. This field is only mandatory if *Include event menu entry* is marked.

See common media type parameters for details on how to configure default messages and alert processing options.

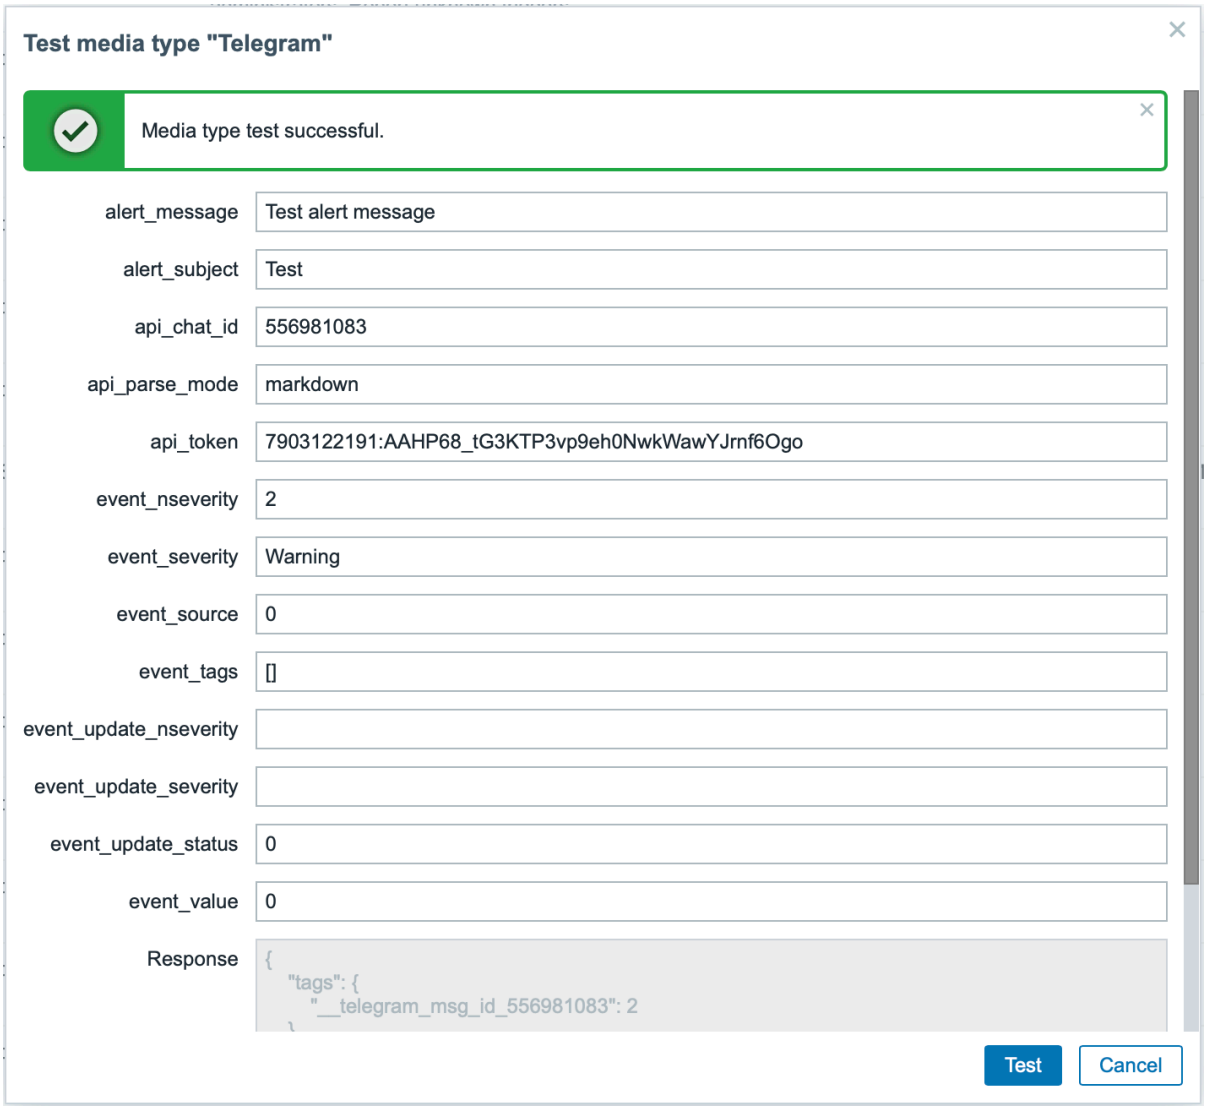
Warning:
Even if a webhook doesn't use default messages, message templates for operation types used by this webhook must still be defined.

Testing

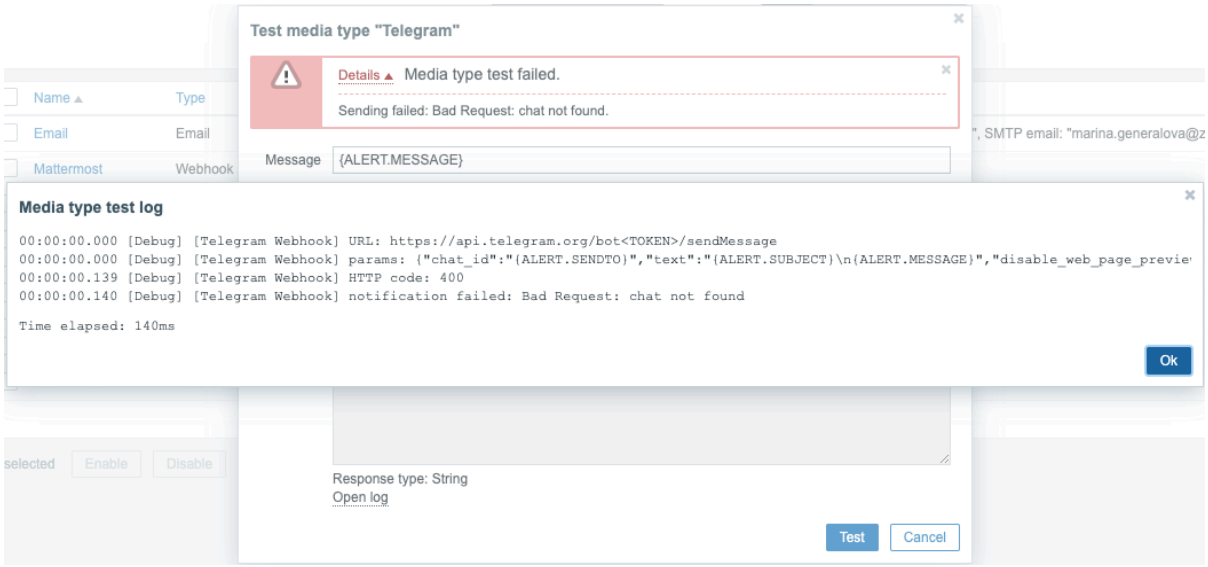
To test a configured webhook media type:

1. Locate the relevant webhook in the list of media types.
2. Click *Test* in the last column of the list (a testing window will open).
3. Edit the webhook parameter values as needed. Replace macros with example values; otherwise, macros will not resolve, and the test will fail.
4. Click *Test*.

Replacing or deleting values in the testing window affects the test procedure only, the actual webhook attribute values will remain unchanged.



To view media type test log entries without leaving the test window, click *Open log* (a new pop-up window will open).



If the webhook test is successful:

- *"Media type test successful."* message is displayed.
- Server response appears in the gray *Response* field.
- Response type (JSON or String) is specified below the *Response* field.

If the webhook test fails:

- *"Media type test failed."* message is displayed, followed by additional failure details.

User media

Once the media type is configured, go to the *Users* > *Users* section and assign the webhook media to an existing user or create a new user to represent the webhook. Steps for setting up user media for an existing user, being common for all media types, are described on the Media types page.

If a webhook uses tags to store ticket\message ID, avoid assigning the same webhook as a media to different users as doing so may cause webhook errors (applies to the majority of webhooks that utilize *Include event menu entry* option). In this case, the best practice is to create a dedicated user to represent the webhook:

1. After configuring the webhook media type, go to the *Users* > *Users* section and create a dedicated Zabbix user to represent the webhook - for example, with a username *Slack* for the Slack webhook. All settings, except media, can be left at their defaults as this user will not be logging into Zabbix.
2. In the user profile, go to a tab *Media* and add a webhook with the required contact information. If the webhook does not use a *Send to* field, enter any combination of supported characters to bypass validation requirements.
3. Grant this user at least read permissions to all hosts for which it should send the alerts.

When configuring alert action, add this user in the *Send to users* field in Operation details - this will tell Zabbix to use the webhook for notifications from this action.

Configuring alert actions

Actions determine which notifications should be sent via the webhook. Steps for configuring actions involving webhooks are the same as for all other media types with these exceptions:

- If a webhook uses webhook tags to store ticket\message ID and handle update\resolve operations, avoid using the same webhook in multiple alert actions for a single problem event. If {EVENT.TAGS.<tag name>} exists and gets updated in the webhook, its resulting value will be undefined. To avoid this, use a new tag name in the webhook for storing updated values. This applies to Jira, Jira Service Desk, Mattermost, Opsgenie, OTRS, Redmine, ServiceNow, Slack, Zammad, and Zendesk webhooks provided by Zabbix and to most webhooks utilizing the *Include event menu entry* option. Note, however, that a single webhook can be used in multiple operations or escalation steps of the same action, as well as in different actions that will not be triggered by the same problem event due to different conditions.
- When using a webhook in actions for internal events, ensure to mark the *Custom message* checkbox and define a custom message in the action operation configuration. Otherwise, a notification will not be sent.

Webhook script examples

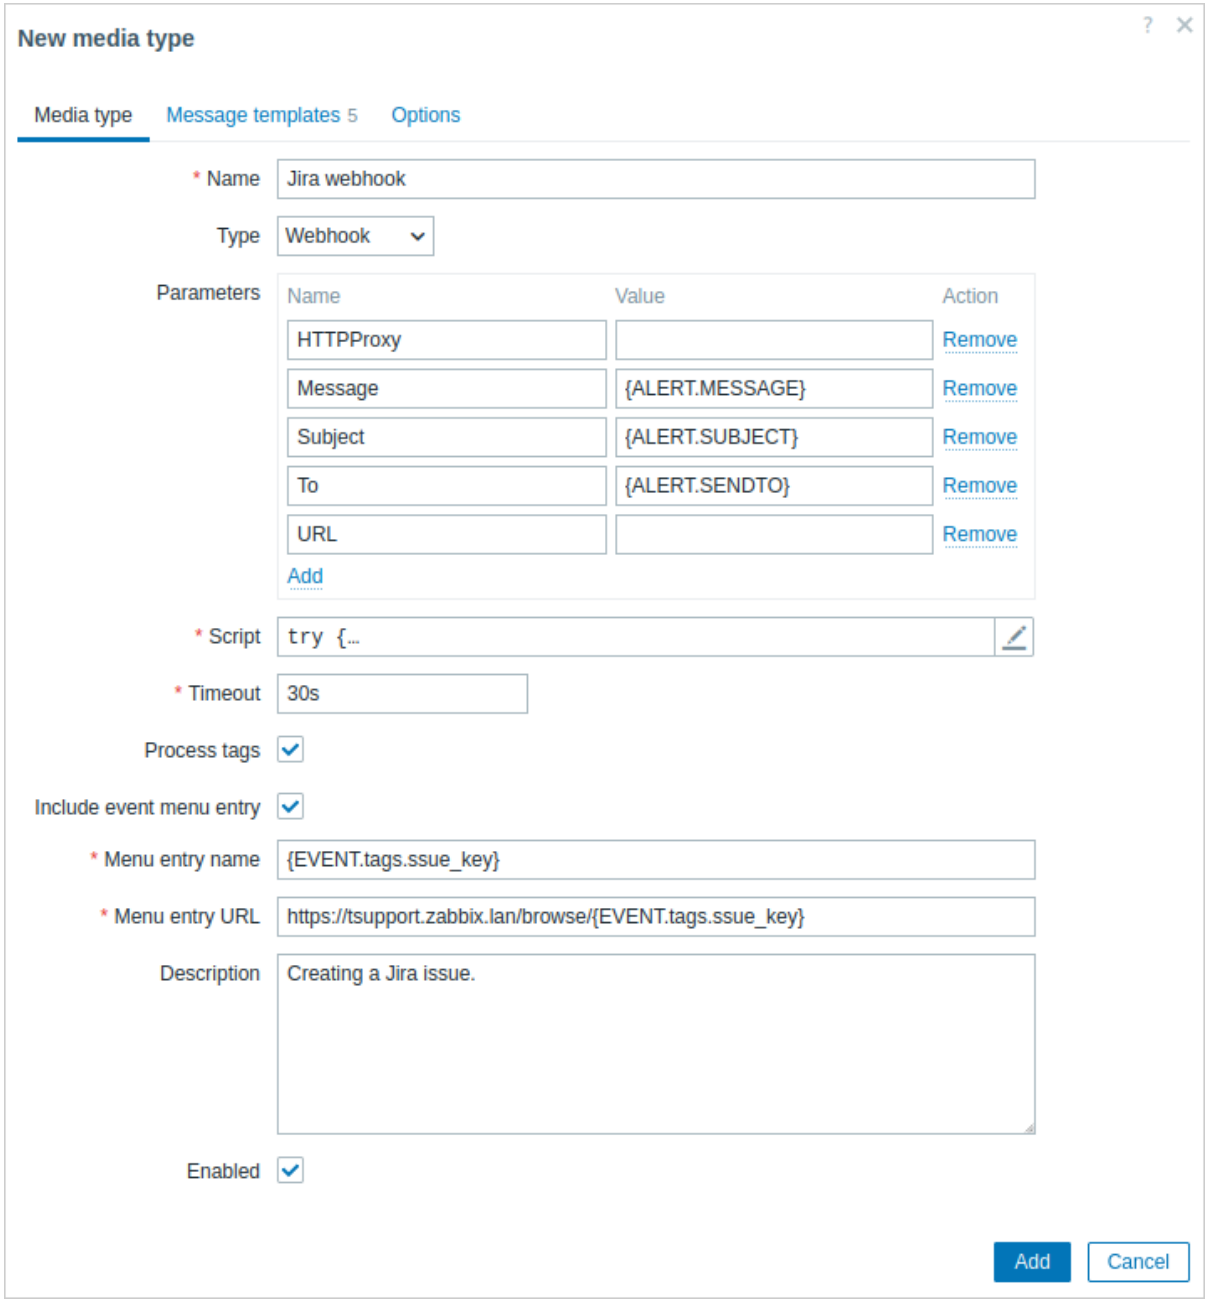
Overview

Though Zabbix offers a large number of webhook integrations available out-of-the-box, you may want to create your own webhooks instead. This section provides examples of custom webhook scripts (used in the *Script* parameter). See webhook section for description of other webhook parameters.

Attention:
Do not use undeclared assignments in preprocessing JavaScript. Use `var` to declare local variables.

Recovery events (whether generated automatically or following a manual close) include resolved event tags (including tags inherited from templates, hosts and triggers). Webhook scripts are executed after the alert is created; tags returned by a webhook script are therefore applied after the initial alert creation and may not be present in the alert that triggered the webhook. If an integration requires tags produced by a webhook to be present in the initial notification, fetch the event tags from the server (for example, using the Event API) or store webhook-produced tags in an external persistent store and correlate there.

Jira webhook (custom)



This script will create a JIRA issue and return some info on the created issue.

```
try {
    Zabbix.log(4, '[ Jira webhook ] Started with params: ' + value);

    var result = {
            'tags': {
                'endpoint': 'jira'
```

```
            }
        },
        params = JSON.parse(value),
        req = new HttpRequest(),
        fields = {},
        resp;

    if (params.HTTPProxy) {
        req.setProxy(params.HTTPProxy);
    }

    req.addHeader('Content-Type: application/json');
    req.addHeader('Authorization: Basic ' + params.authentication);

    fields.summary = params.summary;
    fields.description = params.description;
    fields.project = {key: params.project_key};
    fields.issuetype = {id: params.issue_id};

    resp = req.post('https://jira.example.com/rest/api/2/issue/',
        JSON.stringify({"fields": fields})
    );

    if (req.getStatus() != 201) {
        throw 'Response code: ' + req.getStatus();
    }

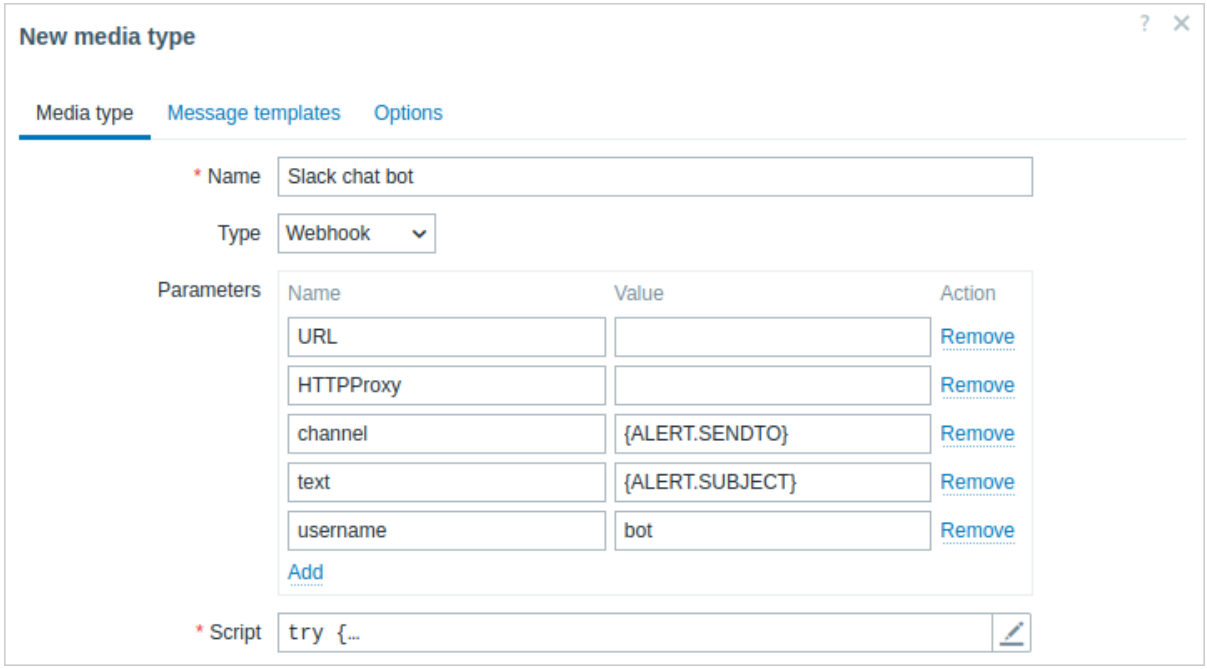
    resp = JSON.parse(resp);
    result.tags.issue_id = resp.id;
    result.tags.issue_key = resp.key;

    return JSON.stringify(result);
}
catch (error) {
    Zabbix.log(4, '[ Jira webhook ] Issue creation failed json : ' + JSON.stringify({"fields": fields}));
    Zabbix.log(3, '[ Jira webhook ] issue creation failed : ' + error);

    throw 'Failed with error: ' + error;
}
```

Slack webhook (custom)

This webhook will forward notifications from Zabbix to a Slack channel.




```
try {
    var params = JSON.parse(value),
        req = new HttpRequest(),
        response;

    if (params.HTTPProxy) {
        req.setProxy(params.HTTPProxy);
    }

    req.addHeader('Content-Type: application/x-www-form-urlencoded');

    Zabbix.log(4, '[ Slack webhook ] Webhook request with value=' + value);

    response = req.post(params.URL, 'payload=' + encodeURIComponent(value));
    Zabbix.log(4, '[ Slack webhook ] Responded with code: ' + req.getStatus() + '. Response: ' + response)

    try {
        response = JSON.parse(response);
    }
    catch (error) {
        if (req.getStatus() < 200 || req.getStatus() >= 300) {
            throw 'Request failed with status code ' + req.getStatus();
        }
        else {
            throw 'Request success, but response parsing failed.';
        }
    }

    if (req.getStatus() !== 200 || !response.ok || response.ok === 'false') {
        throw response.error;
    }

    return 'OK';
}
catch (error) {
    Zabbix.log(3, '[ Slack webhook ] Sending failed. Error: ' + error);

    throw 'Failed with error: ' + error;
}
```

Actions

Overview

If you want some operations taking place as a result of events (for example, notifications sent), you need to configure actions.

Actions can be defined in response to events of all supported types:

- Trigger actions - for events when trigger state changes between `OK` and `Problem`
- Service actions - for events when service state changes between `OK` and `Problem`
- Discovery actions - for events when network discovery takes place
- Autoregistration actions - for events when new active agents auto-register (or host metadata changes for registered ones)
- Internal actions - for events when items become unsupported or triggers go into an unknown state

The key differences of service actions are:

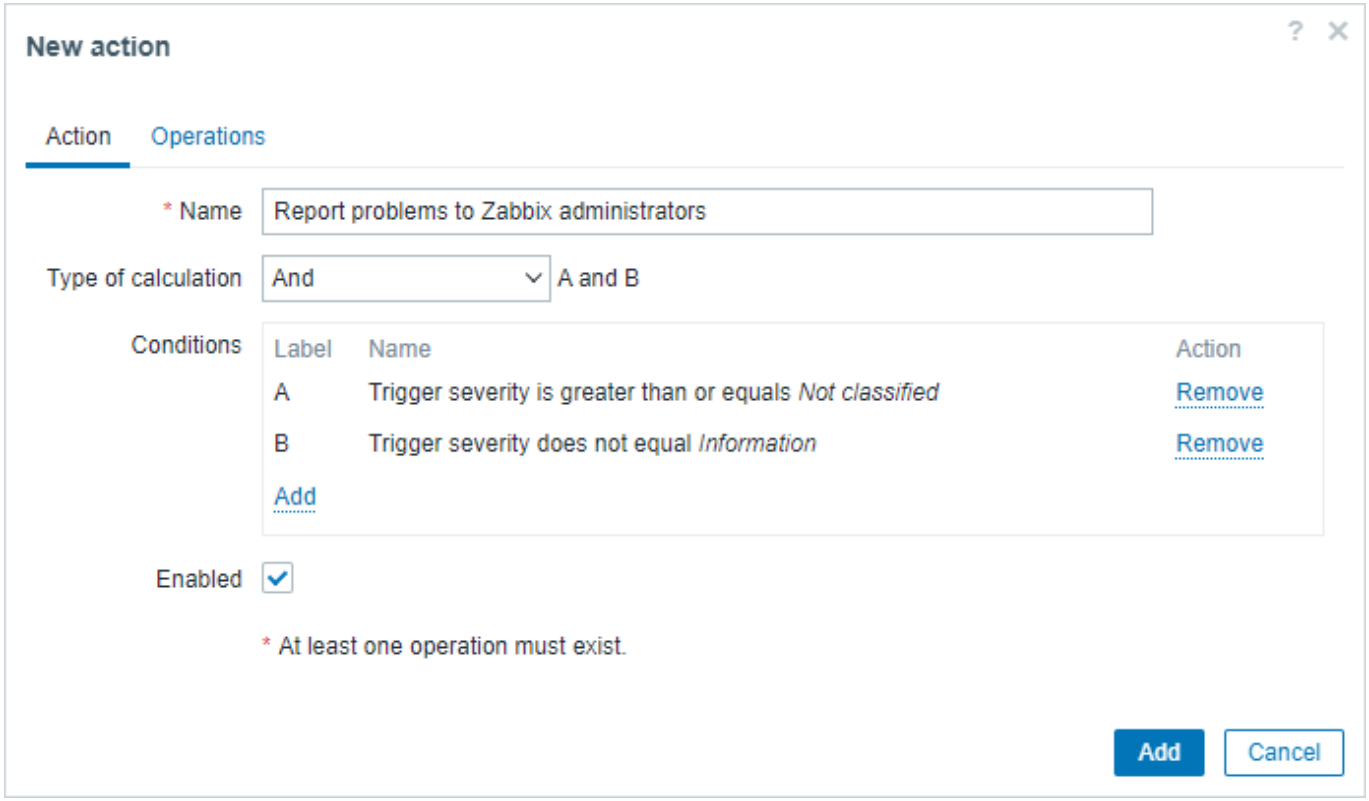
- User access to service actions depends on access rights to services granted by user's role
- Service actions support different set of conditions

Configuring an action

To configure an action, do the following:

1. Go to *Alerts* > *Actions* and select the required action type from the submenu (you can switch to another type later, using the title dropdown).
2. Click *Create action*.
3. Name the action.
4. Choose conditions upon which operations are carried out.
5. Choose the operations to carry out.

General action attributes:



All mandatory input fields are marked with a red asterisk.

Parameter	Description
Name	Unique action name.
Type of calculation	Select the evaluation option for action conditions (with more than one condition): **And** - all conditions must be met. **Or** - enough if one condition is met. **And/Or** - combination of the two: AND with different condition types and OR with the same condition type. **Custom expression** - a user-defined calculation formula for evaluating action conditions.

Parameter	Description
Conditions	List of action conditions. Click *Add* to add a new condition. If no conditions are configured, the action will run for every event that corresponds to the action type being configured.
Enabled	Mark the checkbox to enable the action. Otherwise, it will be disabled.

Conditions

Overview

It is possible to define that an action is executed only if the event matches a defined set of conditions. Conditions are set when configuring the action.

Condition matching is case-sensitive.

Trigger actions

The following conditions can be used in trigger-based actions:

Condition type	Supported operators	Description
Host group	equals does not equal	Specify host groups or host groups to exclude. **equals** - event belongs to this host group. **does not equal** - event does not belong to this host group. Specifying a parent host group implicitly selects all nested host groups. To specify the parent group only, all nested groups have to be additionally set with the **does not equal** operator.
Template	equals does not equal	Specify templates or templates to exclude. **equals** - event belongs to a trigger inherited from this template. **does not equal** - event does not belong to a trigger inherited from this template.
Host	equals does not equal	Specify hosts or hosts to exclude. **equals** - event belongs to this host. **does not equal** - event does not belong to this host.
Tag name	equals does not equal contains does not contain	Specify event tag or event tag to exclude. **equals** - event has this tag. **does not equal** - event does not have this tag. **contains** - event has a tag containing this string. **does not contain** - event does not have a tag containing this string.
Tag value	equals does not equal contains does not contain	Specify event tag and value combination or tag and value combination to exclude. **equals** - event has this tag and value. **does not equal** - event does not have this tag and value. **contains** - event has a tag and value containing these strings. **does not contain** - event does not have a tag and value containing these strings.
Trigger	equals does not equal	Specify triggers or triggers to exclude. **equals** - event is generated by this trigger. **does not equal** - event is generated by any other trigger, except this one.
Event name	contains does not contain	Specify a string in the name of the event generated by the trigger or a string to exclude. By default, the event name matches the trigger name unless a custom event name is specified in trigger configuration. **contains** - event name contains this string. **does not contain** - this string is excluded from the event name. Note: Entered value will be compared to event name with all macros expanded.
Trigger severity	equals does not equal is greater than or equals is less than or equals	Specify trigger severity. **equals** - equal to trigger severity. **does not equal** - not equal to trigger severity. **is greater than or equals** - more or equal to trigger severity. **is less than or equals** - less or equal to trigger severity.

Condition type	Supported operators	Description
Time period	in not in	Specify a time period or a time period to exclude. **in** - event time is within the time period. **not in** - event time is not within the time period. See the time period specification page for description of the format. User macros are supported.
Problem is suppressed	no yes	Specify if the problem is suppressed (not shown) because of host maintenance. **no** - problem is not suppressed. **yes** - problem is suppressed.

Service actions

The following conditions can be used in service actions:

Condition type	Supported operators	Description
Service	equals does not equal	Specify a service or a service to exclude. **equals** - event belongs to this service. **does not equal** - event does not belong to this service. Specifying a parent service implicitly selects all child services. To specify the parent service only, all nested services have to be additionally set with the **does not equal** operator.
Service name	contains does not contain	Specify a string in the service name or a string to exclude. **contains** - event is generated by a service, containing this string in the name. **does not contain** - this string cannot be found in the service name.
Service tag name	equals does not equal contains does not contain	Specify an event tag or an event tag to exclude. Service event tags can be defined in the service configuration section *Tags*. **equals** - event has this tag. **does not equal** - event does not have this tag. **contains** - event has a tag containing this string. **does not contain** - event does not have a tag containing this string.
Service tag value	equals does not equal contains does not contain	Specify an event tag and value combination or a tag and value combination to exclude. Service event tags can be defined in the service configuration section *Tags*. **equals** - event has this tag and value. **does not equal** - event does not have this tag and value. **contains** - event has a tag and value containing these strings. **does not contain** - event does not have a tag and value containing these strings.

Attention:
Make sure to define message templates for Service actions in the *Alerts* > *Media types* menu. Otherwise, the notifications will not be sent.

Discovery actions

The following conditions can be used in discovery-based events:

Condition type	Supported operators	Description
Host IP	equals does not equal	Specify an IP address range or a range to exclude for a discovered host. **equals** - host IP is in the range. **does not equal** - host IP is not in the range. It may have the following formats: Single IP: 192.168.1.33 Range of IP addresses: 192.168.1-10.1-254 IP mask: 192.168.4.0/24 List: 192.168.1.1-254, 192.168.2.1-100, 192.168.2.200, 192.168.4.0/24 Spaces in the list format are supported.
Service type	equals does not equal	Specify a service type of a discovered service or a service type to exclude. **equals** - matches the discovered service. **does not equal** - does not match the discovered service. Available service types: SSH, LDAP, SMTP, FTP, HTTP, HTTPS, POP, NNTP, IMAP, TCP, Zabbix agent, SNMPv1 agent, SNMPv2 agent, SNMPv3 agent, ICMP ping, telnet.
Service port	equals does not equal	Specify a TCP port range of a discovered service or a range to exclude. **equals** - service port is in the range. **does not equal** - service port is not in the range.
Discovery rule	equals does not equal	Specify a discovery rule or a discovery rule to exclude. **equals** - using this discovery rule. **does not equal** - using any other discovery rule, except this one.
Discovery check	equals does not equal	Specify a discovery check or a discovery check to exclude. **equals** - using this discovery check. **does not equal** - using any other discovery check, except this one.
Discovery object	equals	Specify the discovered object. **equals** - equal to discovered object (a device or a service).
Discovery status	equals	**Up** - matches 'Host Up' and 'Service Up' events. **Down** - matches 'Host Down' and 'Service Down' events. **Discovered** - matches 'Host Discovered' and 'Service Discovered' events. **Lost** - matches 'Host Lost' and 'Service Lost' events.
Uptime/Downtime	is greater than or equals is less than or equals	Uptime for 'Host Up' and 'Service Up' events. Downtime for 'Host Down' and 'Service Down' events. **is greater than or equals** - is more or equal to. Parameter is given in seconds. **is less than or equals** - is less or equal to. Parameter is given in seconds.
Received value	equals does not equal is greater than or equals is less than or equals contains does not contain	Specify the value received from an agent (Zabbix, SNMP) check in a discovery rule. String comparison. If several Zabbix agent or SNMP checks are configured for a rule, received values for each of them are checked (each check generates a new event which is matched against all conditions). **equals** - equal to the value. **does not equal** - not equal to the value. **is greater than or equals** - more or equal to the value. **is less than or equals** - less or equal to the value. **contains** - contains the substring. Parameter is given as a string. **does not contain** - does not contain the substring. Parameter is given as a string.
Proxy	equals does not equal	Specify a proxy or a proxy to exclude. **equals** - using this proxy. **does not equal** - using any other proxy except this one.
Proxy group	equals does not equal	Specify a proxy group or a proxy group to exclude. **equals** - using this proxy group. **does not equal** - using any other proxy group except this one.

Note:
Service checks in a discovery rule, which result in discovery events, do not take place simultaneously. Therefore, if **multiple** values are configured for `Service type`, `Service port` or `Received value` conditions in the action, they will be compared to one discovery event at a time, but **not** to several events simultaneously. As a result, actions with multiple values for the same check types may not be executed correctly.

Autoregistration actions

The following conditions can be used in actions based on active agent autoregistration:

Condition type	Supported operators	Description
Host metadata	contains does not contain matches does not match	Specify host metadata or host metadata to exclude. **contains** - host metadata contains the string. **does not contain** - host metadata does not contain the string. Host metadata can be specified in an agent configuration file. **matches** - host metadata matches regular expression. **does not match** - host metadata does not match regular expression.
Host name	contains does not contain matches does not match	Specify a host name or a host name to exclude. **contains** - host name contains the string. **does not contain** - host name does not contain the string. **matches** - host name matches regular expression. **does not match** - host name does not match regular expression.
Proxy	equals does not equal	Specify a proxy or a proxy to exclude. **equals** - using this proxy. **does not equal** - using any other proxy (or server) except this proxy.
Proxy group	equals does not equal	Specify a proxy group or a proxy group to exclude. **equals** - using this proxy group. **does not equal** - using any other proxy group except this one.

Internal event actions

The following conditions can be set for actions based on internal events:

Condition type	Supported operators	Description
Event type	equals	**Item in "not supported" state** - matches events where an item goes from a 'normal' to 'not supported' state. **Low-level discovery rule in "not supported" state** - matches events where a low-level discovery rule goes from a 'normal' to 'not supported' state. **Trigger in "unknown" state** - matches events where a trigger goes from a 'normal' to 'unknown' state.
Host group	equals does not equal	Specify host groups or host groups to exclude. **equals** - event belongs to this host group. **does not equal** - event does not belong to this host group.
Tag name	equals does not equal contains does not contain	Specify event tag or event tag to exclude. **equals** - event has this tag. **does not equal** - event does not have this tag. **contains** - event has a tag containing this string. **does not contain** - event does not have a tag containing this string.
Tag value	equals does not equal contains does not contain	Specify event tag and value combination or tag and value combination to exclude. **equals** - event has this tag and value. **does not equal** - event does not have this tag and value. **contains** - event has a tag and value containing these strings. **does not contain** - event does not have a tag and value containing these strings.
Template	equals does not equal	Specify templates or templates to exclude. **equals** - event belongs to an item/trigger/low-level discovery rule inherited from this template. **does not equal** - event does not belong to an item/trigger/low-level discovery rule inherited from this template.

Condition type	Supported operators	Description
Host	equals does not equal	Specify hosts or hosts to exclude. **equals** - event belongs to this host. **does not equal** - event does not belong to this host.

Type of calculation

The following options of calculating conditions are available:

- **And** - all conditions must be met

Note that using "And" calculation is disallowed between several triggers when they are selected as a `Trigger=` condition. Actions can only be executed based on the event of one trigger.

- **Or** - enough if one condition is met
- **And/Or** - combination of the two: AND with different condition types and OR with the same condition type, for example:

Host group equals Oracle servers
 Host group equals MySQL servers
 Event name contains 'Database is down'
 Event name contains 'Database is unavailable'

is evaluated as

(Host group equals Oracle servers **or** Host group equals MySQL servers**) and (**Event name contains 'Database is down' **or** Event name contains 'Database is unavailable'**)**

- **Custom expression** - a user-defined calculation formula for evaluating action conditions. It must include all conditions (represented as uppercase letters A, B, C, ...) and may include spaces, tabs, brackets ( ), **and** (case sensitive), **or** (case sensitive), **not** (case sensitive).

While the previous example with `And/Or` would be represented as (A or B) and (C or D), in a custom expression you may as well have multiple other ways of calculation:

(A and B) and (C or D)
 (A and B) or (C and D)
 ((A or B) and C) or D
 (not (A or B) and C) or not D
 etc.

Deleting objects used in actions

When attempting to delete objects that are used in an action condition/operation, the user will receive a corresponding error message.

Actions are not disabled if action conditions or operations:

- have deleted objects;
- have objects that are no longer discovered by low-level discovery and have been removed (after the *Delete lost resources* period).

Action conditions/operations are not removed if they have missing objects. Instead, deleted objects are displayed as *Deleted host group*, *Deleted host*, *Deleted trigger*, etc. A warning icon with message "This action has conditions or operations referencing deleted object(s)." is displayed in the "Info" column of the action list.

Operations

Overview

You can define the following operations for all events:

- Send a message
- Execute a remote command

Attention:
Zabbix server does not create alerts if access to the host is explicitly "denied" for the user defined as action operation recipient or if the user has no rights defined to the host at all.

For discovery and autoregistration events, additional operations are available:

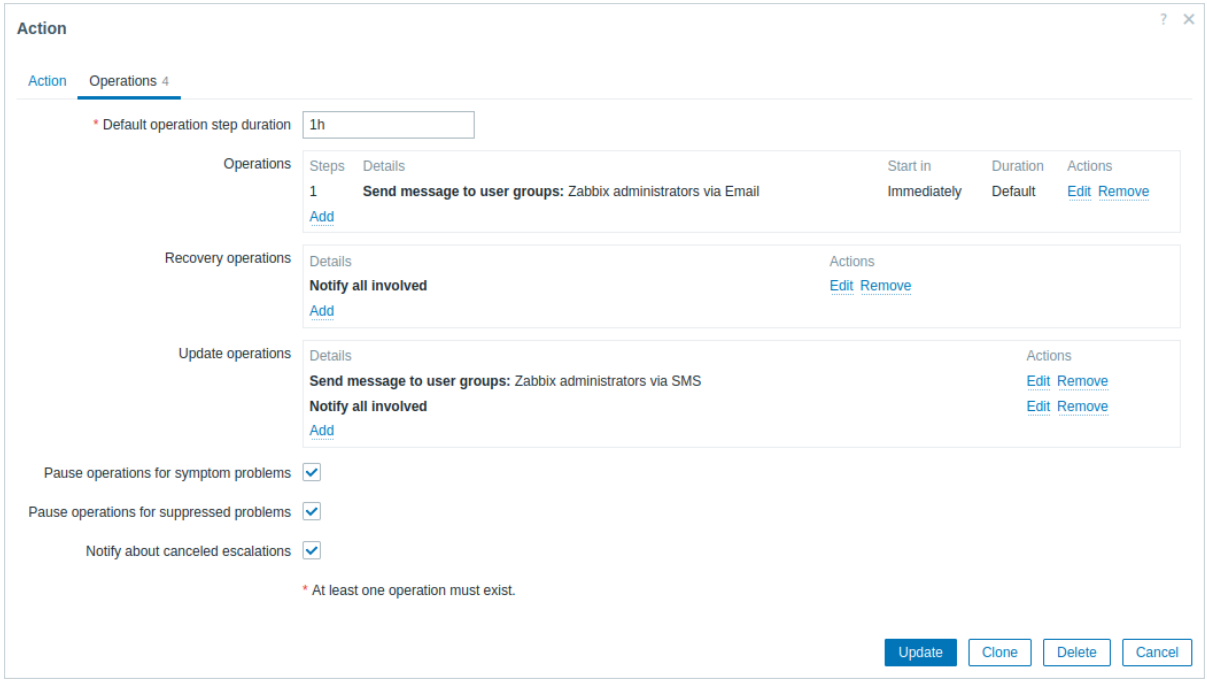
- Add host
- Remove host
- Enable host
- Disable host
- Add to host group
- Remove from host group

- Add host tags
- Remove host tags
- Link template
- Unlink template
- Set host inventory mode

Configuring an operation

To configure an operation:

1. Go to the *Operations* tab in the action configuration.
2. Configure the operation details.



General operation attributes:

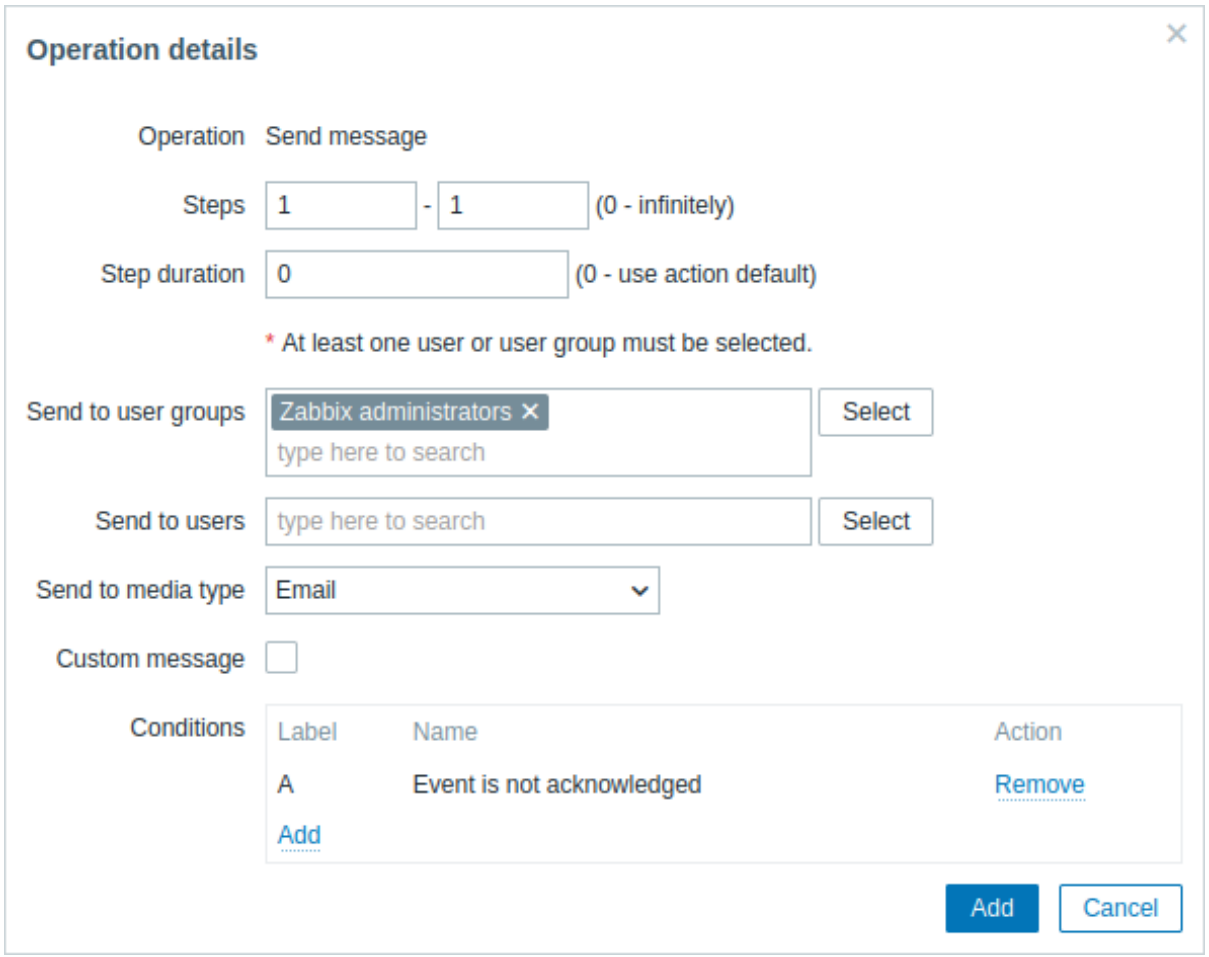
Parameter	Description
Default operation step duration	Duration of one operation step by default (60 seconds to 1 week). For example, an hour-long step duration means that if an operation is carried out, an hour will pass before the next step. Time suffixes are supported, e.g., 60s, 1m, 2h, 1d. User macros are supported.
Operations	Action operations (if any) are displayed, with these details: **Steps** - escalation step(s) to which the operation is assigned. **Details** - type of operation and its recipient/target. The operation list also displays the media type (email, SMS or script) used as well as the name and surname (in parentheses after the username) of a notification recipient. **Start in** - how long after an event the operation is performed. **Duration (sec)** - step duration is displayed. *Default* is displayed if the step uses default duration, and a time is displayed if custom duration is used. **Actions** - links for editing and removing an operation are displayed.
Recovery operations	Action operations (if any) are displayed, with these details: **Details** - type of operation and its recipient/target. The operation list also displays the media type (email, SMS or script) used as well as the name and surname (in parentheses after the username) of a notification recipient. **Actions** - links for editing and removing an operation are displayed.
Update operations	Action operations (if any) are displayed, with these details: **Details** - type of operation and its recipient/target. The operation list also displays the media type (email, SMS or script) used as well as the name and surname (in parentheses after the username) of a notification recipient. **Actions** - links for editing and removing an operation are displayed.

Parameter	Description
Pause operations for symptom problems	Mark this checkbox to pause operations (after the first operation) for symptom problems. Note that this setting affects only problem escalations; recovery and update operations will not be affected. This option is available for *Trigger actions* only.
Pause operations for suppressed problems	Mark this checkbox to delay the start of operations for the duration of a maintenance period. When operations are started, after the maintenance, all operations are performed including those for the events during the maintenance. Note that this setting affects only problem escalations; recovery and update operations will not be affected. If you unmark this checkbox, operations will be executed without delay even during a maintenance period. This option is not available for *Service actions*.
Notify about canceled escalations	Unmark this checkbox to disable notifications about canceled escalations (when host, item, trigger or action is disabled).

All mandatory input fields are marked with a red asterisk.

To configure details of a new operation, click Add in the *Operations* block. To edit an existing operation, click Edit next to the operation. A pop-up window will open where you can edit the operation step details.

Operation details



Parameter	Description
Operation	Select the operation: **Send message** - send message to user. **<remote command name>** - execute a remote command. Commands are available for execution if previously defined in global scripts with *Action operation* selected as its scope. More operations are available for discovery and autoregistration based events (see above).
Steps	Select the step(s) to assign the operation to in an escalation schedule: **From** - execute starting with this step. **To** - execute until this step (0=infinity, execution will not be limited).
Step duration	Custom duration for these steps (0=use default step duration). Time suffixes are supported, e.g., 60s, 1m, 2h, 1d. User macros are supported. Several operations can be assigned to the same step. If these operations have different step duration defined, the shortest one is taken into account and applied to the step.
Operation type: send message	
Send to user groups	Select user groups to send the message to. The user group must have at least "read" permissions to the host in order to be notified.
Send to users	Select users to send the message to. The user must have at least "read" permissions to the host in order to be notified.
Send to media type	Send message to all available (configured and enabled) media types or a specific one only.
Custom message	If selected, the custom message can be configured. For notifications about internal events via webhooks, custom message is mandatory.
Subject	Subject of the custom message. The subject may contain macros. It is limited to 255 characters.
Message	The custom message. The message may contain macros. It is limited to certain amount of characters depending on the type of database (see Sending message for more information).
Operation type: remote command	

Parameter	Description
Target list	Select targets to execute the command on: **Current host** - command is executed on the host of the trigger that caused the problem event. This option will not work if there are multiple hosts in the trigger. **Host** - select host(s) to execute the command on. **Host group** - select host group(s) to execute the command on. Specifying a parent host group implicitly selects all nested host groups. Thus the remote command will also be executed on hosts from nested groups. A command on a host is executed only once, even if the host matches more than once (e.g., from several host groups; individually and from a host group). The target list is meaningless if a custom script is executed on Zabbix server. Selecting more targets in this case only results in the script being executed on the server more times. Note that for global scripts, the target selection also depends on the *Host group* setting in the global script configuration. *Target list* option is not available for *Service actions* because in this case remote commands are always executed on Zabbix server.
Conditions	Condition for performing the operation: **Event is not acknowledged** - only when the event is unacknowledged. **Event is acknowledged** - only when the event is acknowledged. *Conditions* option is only available for *Trigger actions*.

When done, click *Add* to add the operation to the list of *Operations*.

Sending message

Overview

Sending a message is one of the best ways of notifying people about a problem. That is why it is one of the primary actions offered by Zabbix.

Configuration

To be able to send and receive notifications from Zabbix you have to:

- Define the media to send a message to.

If the operation takes place outside of the When active time period defined for the selected media in the user configuration, the message will not be sent.

The default trigger severity ('Not classified') **must be** checked in the user media configuration if you want to receive notifications for non-trigger events such as discovery, active agent autoregistration or internal events.

- Configure an action operation that sends a message to one of the defined media.

Attention:
Zabbix sends notifications only to those users that have at least 'read' permissions to the host that generated the event. At least one host of a trigger expression must be accessible.

You can configure custom scenarios for sending messages using escalations.

To successfully receive and read emails from Zabbix, email servers/clients must support standard 'SMTP/MIME email' format since Zabbix sends UTF-8 data (If the subject contains ASCII characters only, it is not UTF-8 encoded.). The subject and the body of the message are base64-encoded to follow 'SMTP/MIME email' format standard.

Message limit after all macros expansion is the same as message limit for Remote commands.

Tracking messages

You can view the status of messages sent in *Monitoring* > *Problems*.

In the *Actions* column you can see summarized information about actions taken. In there green numbers represent messages sent, red ones - failed messages. *In progress* indicates that an action is initiated. *Failed* informs that no action has executed successfully.

If you click the event time to view event details, you will be able to see details of messages sent (or not sent) due to the event in the *Actions* block.

In *Reports* > *Action log* you will see details of all actions taken for those events that have an action configured.

Remote commands

Overview

With remote commands you can define that a certain pre-defined command is automatically executed on the monitored host upon some condition.

Thus remote commands are a powerful mechanism for smart pro-active monitoring.

In the most obvious uses of the feature you can try to:

- Automatically restart some application (web server, middleware, CRM) if it does not respond
- Use IPMI 'reboot' command to reboot some remote server if it does not answer requests
- Automatically free disk space (removing older files, cleaning /tmp) if running out of disk space
- Migrate a VM from one physical box to another depending on the CPU load
- Add new nodes to a cloud environment upon insufficient CPU (disk, memory, whatever) resources

Configuring an action for remote commands is similar to configuring an action for sending a message; the only difference is that Zabbix executes a command instead of sending a message.

Remote commands can be executed by Zabbix server, Zabbix proxy, or Zabbix agent. Note the following:

- Remote commands on Zabbix agent can be executed directly by Zabbix server or through Zabbix proxy.
- Remote commands executed by Zabbix server are run as described in Command execution, including exit code checking.
- Remote commands are executed even if the target host is in maintenance mode.

Remote commands are disabled by default on both the Zabbix agent and Zabbix proxy. They can be enabled as follows:

- For remote commands executed by Zabbix agent, add an `AllowKey=system.run[<command>,*]` or `AllowKeyRegexp=^system\.r` parameter for each allowed command in the agent configuration, where `*` matches either execution mode (`wait` or `nowait`).
- For remote commands executed by Zabbix proxy, set the `EnableRemoteCommands` parameter to `1` in the proxy configuration. Note that remote commands do not need to be enabled on Zabbix proxy when they are executed by Zabbix agent on a host monitored by the proxy.

Attention:
Do not use `AllowKey=system.run[*]` or `AllowKeyRegexp=^system\.run\[.*\]$` unless necessary, as it allows the execution of any command on the host. For security reasons, it is recommended to allow only exact commands, while `system.run[*]` is denied by default (`DenyKey=system.run[*]`).

See also Restricting agent checks.

Remote command limit

Remote command limit after resolving all macros depends on the type of database and character set (non-ASCII characters require more than one byte to be stored):

Database	Limit in characters	Limit in bytes
MySQL	65535	65535
PostgreSQL	65535	not limited
SQLite (only Zabbix proxy)	65535	not limited

Remote command execution output (return value) is limited to 16MB (including trailing whitespace that is truncated). IPMI remote command limit is based on the installed IPMI library. Note that database limits apply to all remote commands.

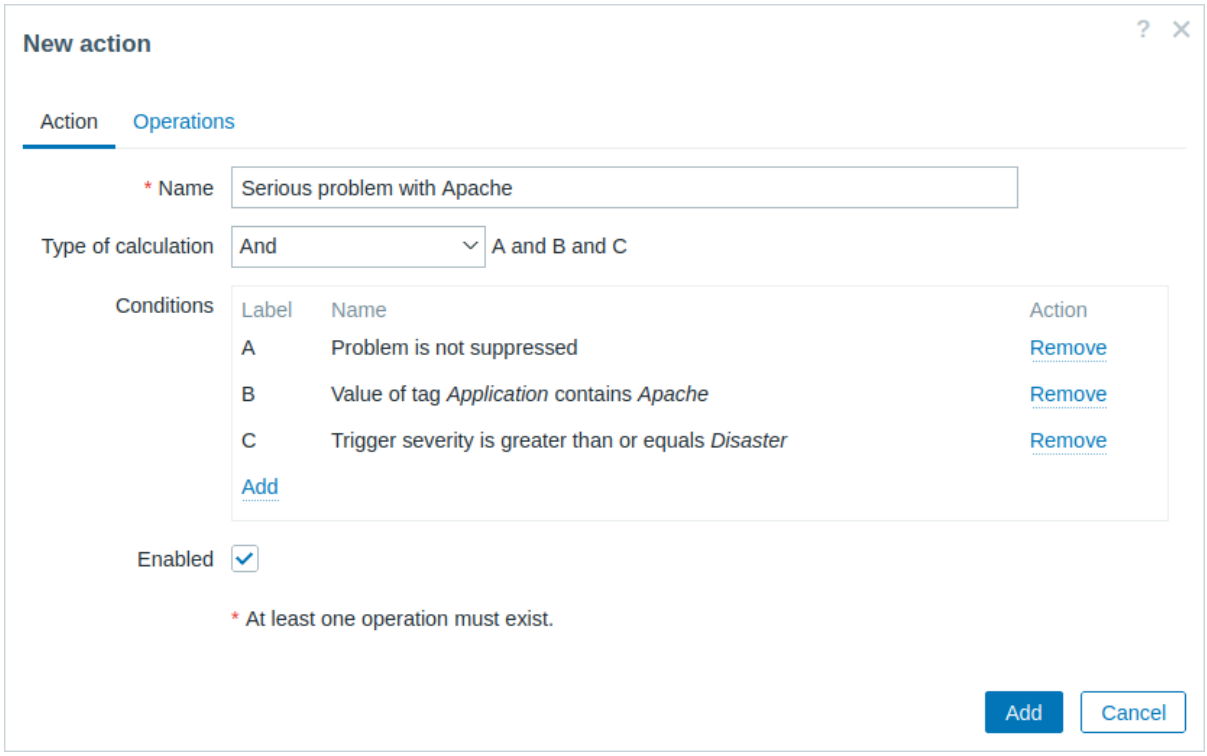
Configuration

Remote commands executed by Zabbix agent must first be allowed in the agent configuration using the `AllowKey` or `AllowKeyRegexp` parameter.

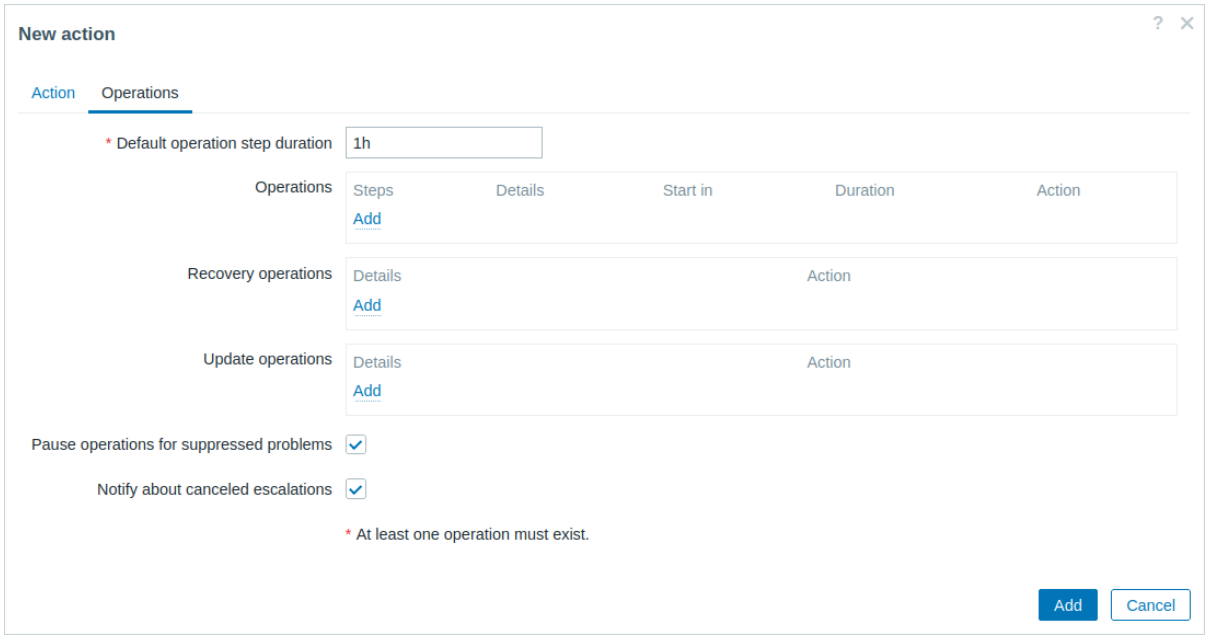
Make sure that the `AllowKey=system.run[<command>,*]` or `AllowKeyRegexp=^system\.run\[<command>,.*\]$` parameter is added to the agent configuration for each allowed command to allow its execution in `nowait` mode. Restart the agent daemon after changing the parameter.

Then, when configuring a new action in *Alerts* > *Actions* > *Trigger actions*:

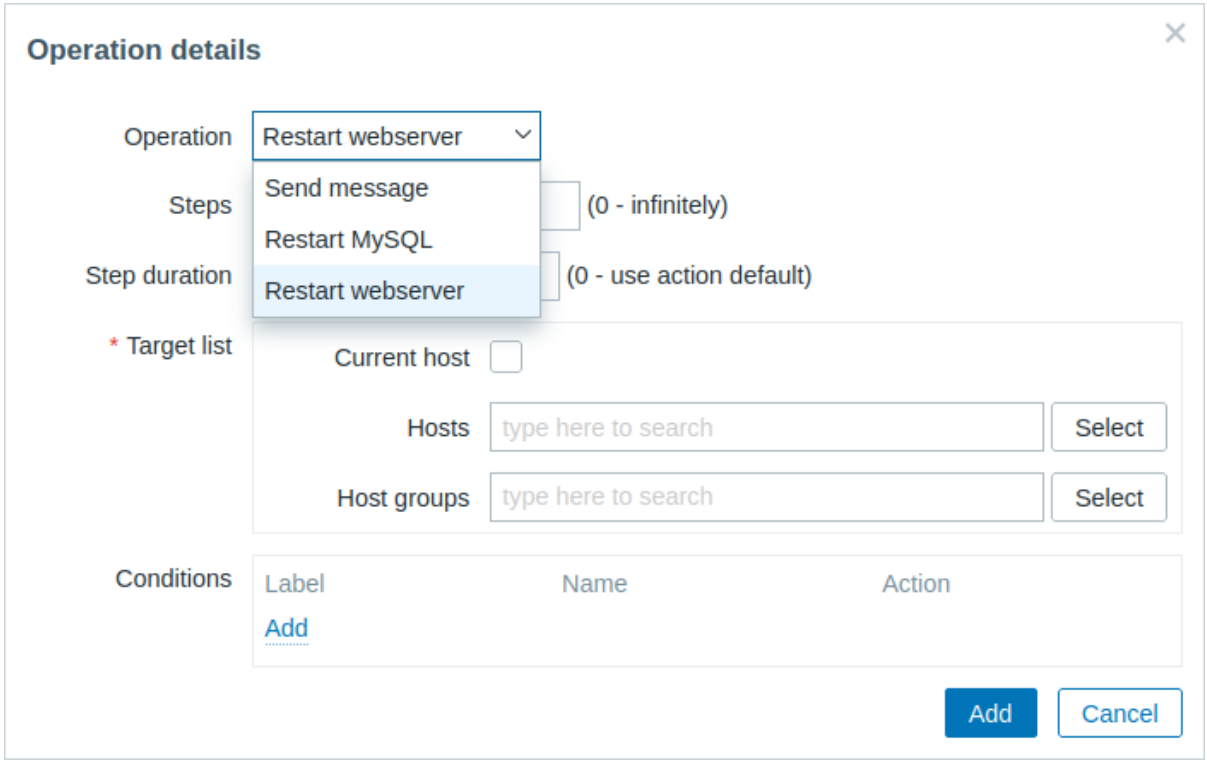
1. Define the appropriate conditions, for example, set that the action is activated upon any disaster problems with one of Apache applications.



2. In the *Operations* tab, click *Add* in the *Operations*, *Recovery operations*, or *Update operations* section.



3. Select one of the predefined scripts from the *Operation* dropdown list and set the *Target list* for the script.



Predefined scripts

Scripts that are available for action operations (webhook, script, SSH, Telnet, IPMI) are defined in global scripts with *Action operation* selected as its scope.

For example:

```
sudo /etc/init.d/apache restart
```

In this case, Zabbix will try to restart an Apache process. With this command, make sure that the command is executed on Zabbix agent (click the *Zabbix agent* button against *Execute on*).

Attention:
Note the use of sudo - Zabbix user does not have permissions to restart system services by default. See below for hints on how to configure sudo.

Note:
Starting with Zabbix agent 7.0, remote commands can also be executed on an agent that is operating in active mode. Zabbix agent - whether active or passive - should run on the remote host, and executes the commands in background.

Remote commands on Zabbix agent are executed without timeout by the `system.run[,nowait]` key and are not checked for execution results. On Zabbix server and Zabbix proxy, remote commands are executed with timeout as set in the `TrapperTimeout` parameter of `zabbix_server.conf` or `zabbix_proxy.conf` file and are checked for execution results. For additional information, see *Script timeout*.

Access permissions

Make sure that the 'zabbix' user has execute permissions for configured commands. One may be interested in using sudo to give access to privileged commands. To configure access, execute as root:

```
visudo
```

Example lines that could be used in sudoers file:

```
# allows 'zabbix' user to run all commands without password.
zabbix ALL=NOPASSWD: ALL

# allows 'zabbix' user to restart apache without password.
zabbix ALL=NOPASSWD: /etc/init.d/apache restart
```

Note:
On some systems `sudoers` file will prevent non-local users from executing commands. To change this, comment out `requiretty` option in `/etc/sudoers`.

Remote commands with multiple interfaces

If the target system has multiple interfaces of the selected type (Zabbix agent or IPMI), remote commands will be executed on the default interface.

It is possible to execute remote commands via SSH and Telnet using another interface than the Zabbix agent one. The available interface to use is selected in the following order:

- Zabbix agent default interface
- SNMP default interface
- JMX default interface
- IPMI default interface

IPMI remote commands

For IPMI remote commands the following syntax should be used:

```
<command> [<value>]
```

where

- <command> - one of IPMI commands without spaces
- <value> - 'on', 'off' or any unsigned integer. <value> is an optional parameter.

Examples

Examples of global scripts that may be used as remote commands in action operations.

Example 1

Restart of Windows on certain condition.

In order to automatically restart Windows upon a problem detected by Zabbix, define the following script:

Script parameter	Value
Scope	'Action operation'
Type	'Script'
Command	`c:\windows\system32\shutdown.exe -r -f`

Example 2

Restart the host by using IPMI control.

Script parameter	Value
Scope	'Action operation'
Type	'IPMI'
Command	reset

Example 3

Power off the host by using IPMI control.

Script parameter	Value
Scope	'Action operation'
Type	'IPMI'
Command	power off

Additional operations

Overview

In this section you may find some details of additional operations for discovery/autoregistration events.

Adding host

Hosts are added during the discovery process, as soon as a host is discovered, rather than at the end of the discovery process.

> **Note:**
> As network discovery can take some time, due to many unavailable hosts/services, having patience and using reasonable IP ranges is recommended.

When adding a host, its name is decided by the standard **gethostbyname** function. If the host can be resolved, resolved name is used. If not, the IP address is used. Besides, if IPv6 address must be used for a host name, then all ":" (colons) are replaced by "_" (underscores), since colons are not allowed in host names.

> **Attention:**
> If performing discovery by a proxy, currently hostname lookup still takes place on Zabbix server.

> **Attention:**
> If a host already exists in Zabbix configuration with the same name as a newly discovered one, Zabbix will add **_N** to the hostname, where **N** is increasing number, starting with 2.

Using macros in messages

Overview

In message subjects and message text you can use macros for more efficient problem reporting.

In addition to a number of built-in macros, user macros and expression macros are also supported. A full list of macros supported by Zabbix is available.

Examples

Examples here illustrate how you can use macros in messages.

Example 1

Message subject:

```
Problem: {TRIGGER.NAME}
```

When you receive the message, the message subject will be replaced by something like:

```
Problem: Processor load is too high on Zabbix server
```

Example 2

Message:

```
Processor load is: {?last(/zabbix.zabbix.com/system.cpu.load[,avg1])}
```

When you receive the message, the message will be replaced by something like:

```
Processor load is: 1.45
```

Example 3

Message:

```
Latest value: {?last(/{HOST.HOST}/{ITEM.KEY})}
MAX for 15 minutes: {?max(/{HOST.HOST}/{ITEM.KEY},15m)}
MIN for 15 minutes: {?min(/{HOST.HOST}/{ITEM.KEY},15m)}
```

When you receive the message, the message will be replaced by something like:

```
Latest value: 1.45
MAX for 15 minutes: 2.33
MIN for 15 minutes: 1.01
```

Example 4

Message:

```
http://<server_ip_or_name>/zabbix/tr_events.php?triggerid={TRIGGER.ID}&eventid={EVENT.ID}
```

When you receive the message, it will contain a link to the *Event details* page, which provides information about the event, its trigger, and a list of latest events generated by the same trigger.

Example 5

Informing about values from several hosts in a trigger expression.

Message:

```
Problem name: {TRIGGER.NAME}
Trigger expression: {TRIGGER.EXPRESSION}

1. Item value on {HOST.NAME1}: {ITEM.VALUE1} ({ITEM.NAME1})
2. Item value on {HOST.NAME2}: {ITEM.VALUE2} ({ITEM.NAME2})
```

When you receive the message, the message will be replaced by something like:

```
Problem name: Processor load is too high on a local host
Trigger expression: last(/Myhost/system.cpu.load[percpu,avg1])>5 or last(/Myotherhost/system.cpu.load[perc

1. Item value on Myhost: 0.83 (Processor load (1 min average per core))
2. Item value on Myotherhost: 5.125 (Processor load (1 min average per core))
```

Example 6

Receiving details of both the problem event and recovery event in a recovery message:

Message:

```
Problem:

Event ID: {EVENT.ID}
Event value: {EVENT.VALUE}
Event status: {EVENT.STATUS}
Event time: {EVENT.TIME}
Event date: {EVENT.DATE}
Event age: {EVENT.AGE}
Event acknowledgment: {EVENT.ACK.STATUS}
Event update history: {EVENT.UPDATE.HISTORY}

Recovery:

Event ID: {EVENT.RECOVERY.ID}
Event value: {EVENT.RECOVERY.VALUE}
Event status: {EVENT.RECOVERY.STATUS}
Event time: {EVENT.RECOVERY.TIME}
Event date: {EVENT.RECOVERY.DATE}
Operational data: {EVENT.OPDATA}
```

When you receive the message, the macros will be replaced by something like:

```
Problem:

Event ID: 21874
Event value: 1
Event status: PROBLEM
Event time: 13:04:30
Event date: 2018.01.02
Event age: 5m 0s
Event acknowledgment: Yes
Event update history: 2018.01.02 13:05:51 "John Smith (Admin)"
Actions: acknowledged.

Recovery:

Event ID: 21896
Event value: 0
```

```
Event status: OK
Event time: 13:10:07
Event date: 2018.01.02
Operational data: Current value is 0.83
```

Recovery operations

Overview

Recovery operations allow you to be notified when problems are resolved.

Both messages and remote commands are supported in recovery operations. While several operations can be added, escalation is not supported - all operations are assigned to a single step and therefore will be performed simultaneously.

Use cases

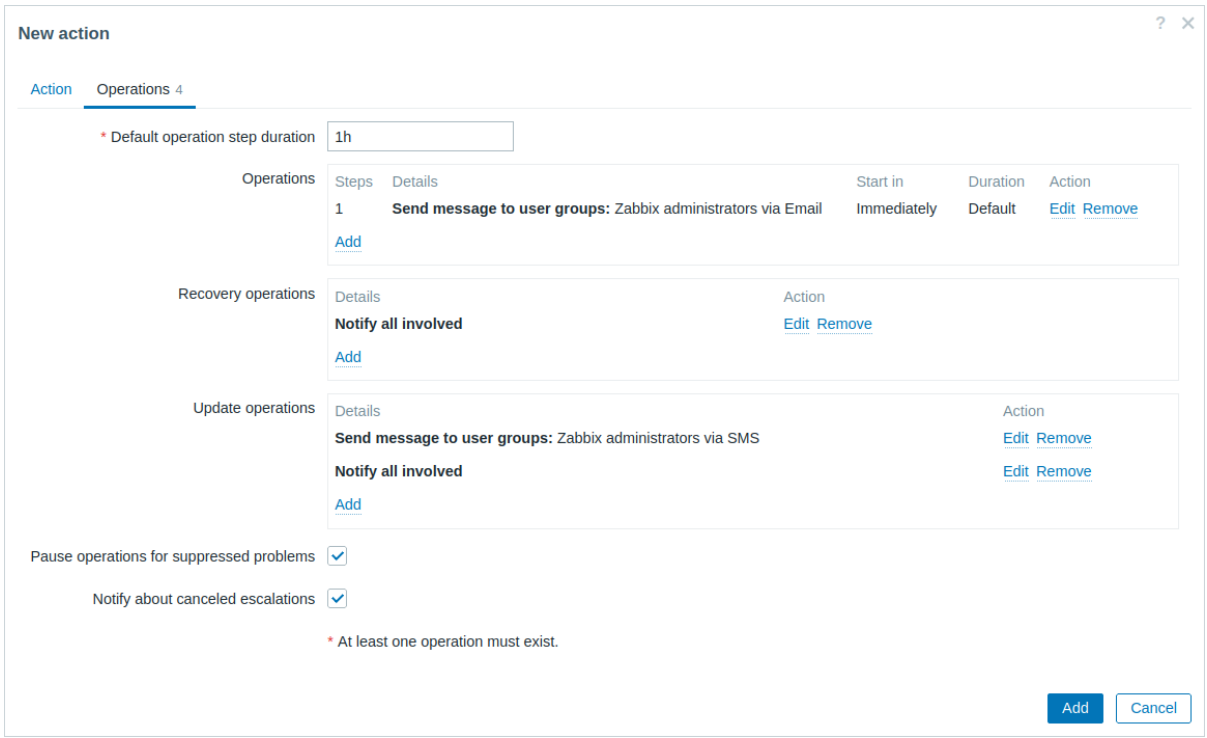
Some use cases for recovery operations are as follows:

1. Notify on a recovery all users that were notified on the problem:
   - Select *Notify all involved* as operation type.
2. Have multiple operations upon recovery: send a notification and execute a remote command:
   - Add operation types for sending a message and executing a command.
3. Open a ticket in external helpdesk/ticketing system and close it when the problem is resolved:
   - Create an external script that communicates with the helpdesk system.
   - Create an action having operation that executes this script and thus opens a ticket.
   - Have a recovery operation that executes this script with other parameters and closes the ticket.
   - Use the {EVENT.ID} macro to reference the original problem.

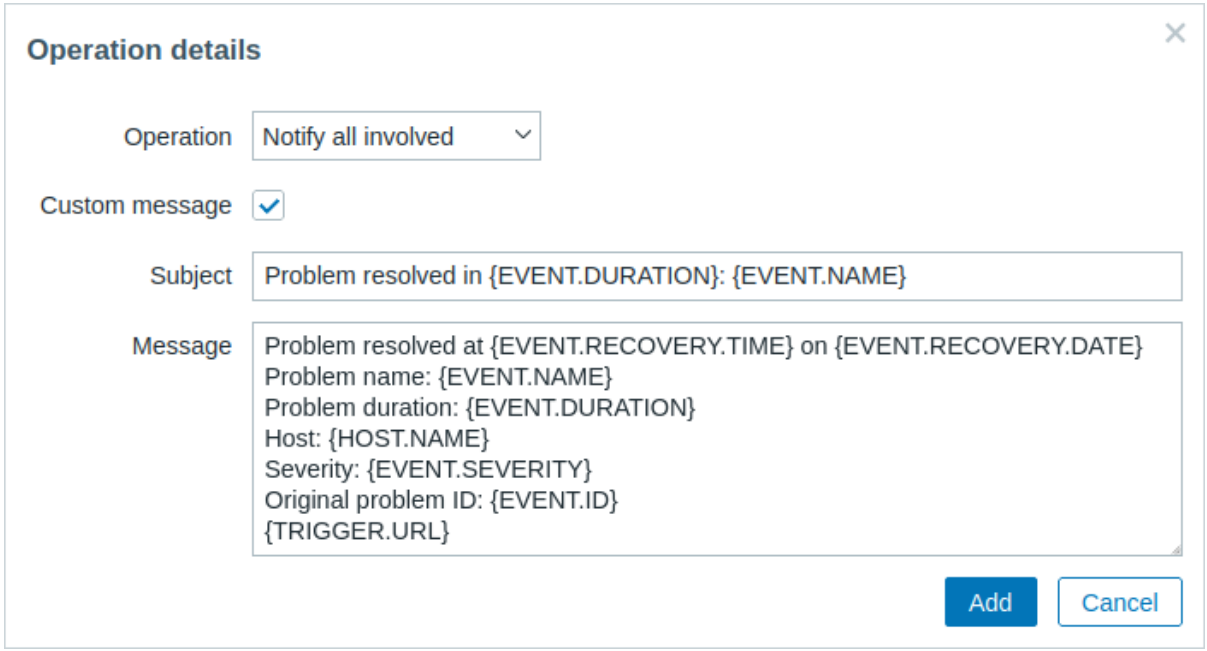
Configuring a recovery operation

To configure a recovery operation, follow these steps:

1. Go to the *Operations* tab in the action configuration form.
2. To configure a new recovery operation, click *Add* in the *Recovery operations* section. To edit an existing operation, click *Edit* next to the operation.
3. Configure the operation details.



Recovery operation details



Three operation types are available for recovery events:

- **Send message** - send recovery message to specified user;
- **Notify all involved** - send recovery message to all users who were notified on the problem event;
- **<remote command name>** - execute a remote command. Commands are available for execution if previously defined in global scripts with *Action operation* selected as its scope.

Parameters for each operation type are described below. All mandatory input fields are marked with a red asterisk. When done, click *Add* to add operation to the list of *Recovery operations*.

Note:
Note that if the same recipient is defined in several operation types without specified *Custom message*, duplicate notifications are not sent.

Operation type: send message

Parameter	Description
Send to user groups	Select user groups to send the recovery message to. The user group must have at least "read" permissions to the host in order to be notified.
Send to users	Select users to send the recovery message to. The user must have at least "read" permissions to the host in order to be notified.
Send to media type	Send default recovery message to all available (configured and enabled) media types or a specific one only.
Custom message	If selected, a custom message can be defined.
Subject	Subject of the custom message. The subject may contain macros.
Message	The custom message. The message may contain macros.

Operation type: remote command

Parameter	Description
Target list	Select targets to execute the command on: **Current host** - command is executed on the host of the trigger that caused the problem event. This option will not work if there are multiple hosts in the trigger. **Host** - select host(s) to execute the command on. **Host group** - select host group(s) to execute the command on. Specifying a parent host group implicitly selects all nested host groups. Thus the remote command will also be executed on hosts from nested groups. A command on a host is executed only once, even if the host matches more than once (e.g., from several host groups; individually and from a host group). The target list is meaningless if the command is executed on Zabbix server. Selecting more targets in this case only results in the command being executed on the server more times. Note that for global scripts, the target selection also depends on the *Host group* setting in the global script configuration.

Operation type: notify all involved

Parameter		Description
Custom message		If selected, a custom message can be defined.
	Subject	Subject of the custom message. The subject may contain macros.
	Message	The custom message. The message may contain macros.

Update operations

Overview

Update operations are available in actions with the following event sources:

- *Triggers* - when problems are updated by other users, i.e. commented upon, acknowledged, severity has been changed, closed (manually);
- *Services* - when the severity of a service has changed but the service is still not recovered.

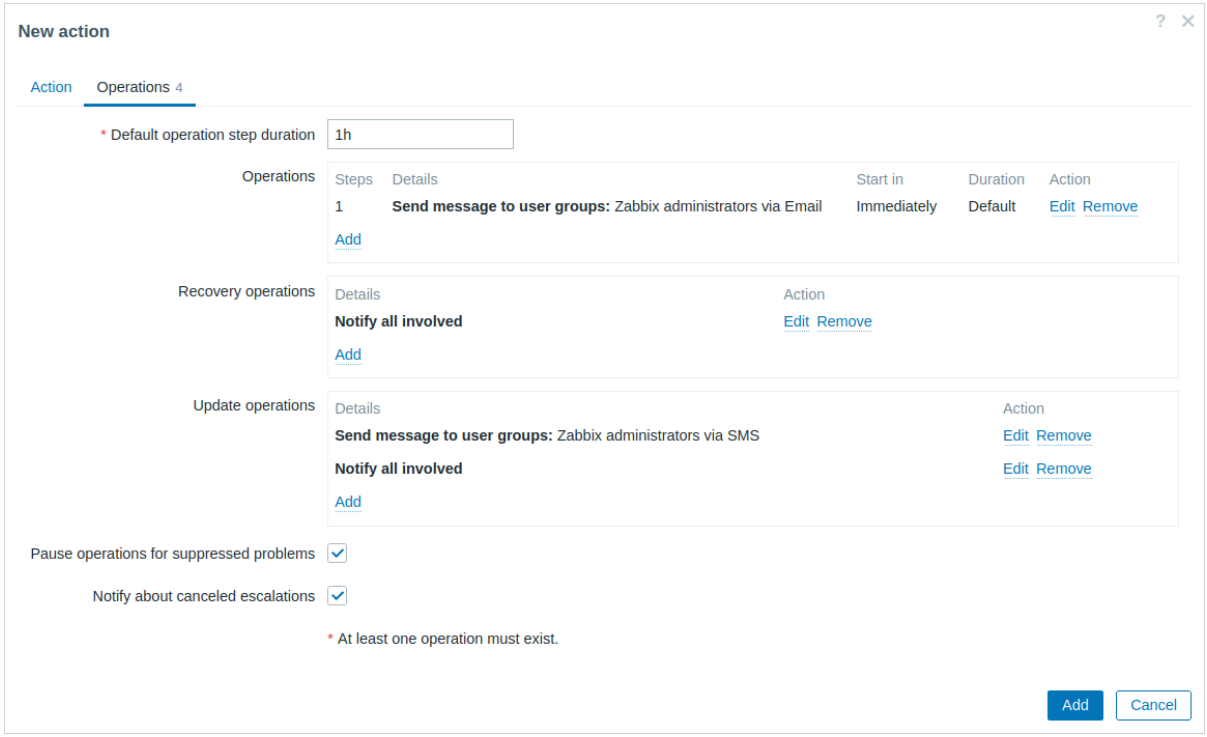
Please note that users do not receive notifications about their own updates.

Both messages and remote commands are supported in update operations. While several operations can be added, escalation is not supported - all operations are assigned to a single step and therefore will be performed simultaneously.

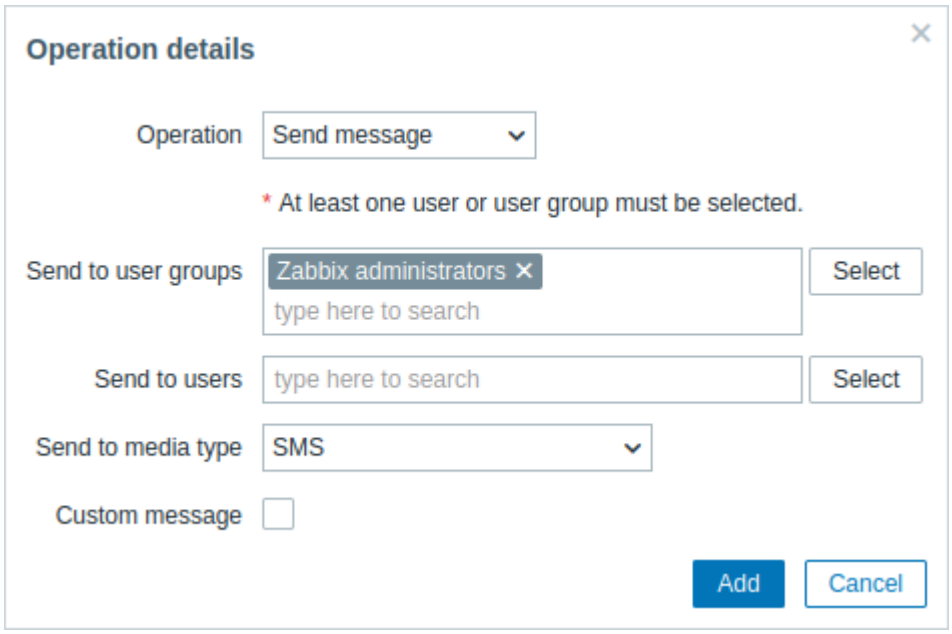
Configuring an update operation

To configure an update operation, follow these steps:

1. Go to the *Operations* tab in the action configuration form.
2. To configure a new update operation, click *Add* in the *Update operations* section. To edit an existing operation, click *Edit* next to the operation.
3. Configure the operation details.



Update operation details



Update operations offer the same set of parameters as Recovery operations.

Escalations

Overview

With escalations you can create custom scenarios for sending notifications or executing remote commands.

In practical terms it means that:

- Users can be informed about new problems immediately.
- Notifications can be repeated until the problem is resolved.
- Sending a notification can be delayed.
- Notifications can be escalated to another ”higher” user group.
- Remote commands can be executed immediately or when a problem is not resolved for a lengthy period.

Actions are escalated based on the **escalation step**. Each step has a duration in time.

You can define both the default duration and a custom duration of an individual step. The minimum duration of one escalation step is 60 seconds.

You can start actions, such as sending notifications or executing commands, from any step. Step one is for immediate actions. If you want to delay an action, you can assign it to a later step. For each step, several actions can be defined.

The number of escalation steps is not limited.

Escalations are defined when configuring an operation. Escalations are supported for problem operations only, not recovery.

Miscellaneous aspects of escalation behavior

Let's consider what happens in different circumstances if an action contains several escalation steps.

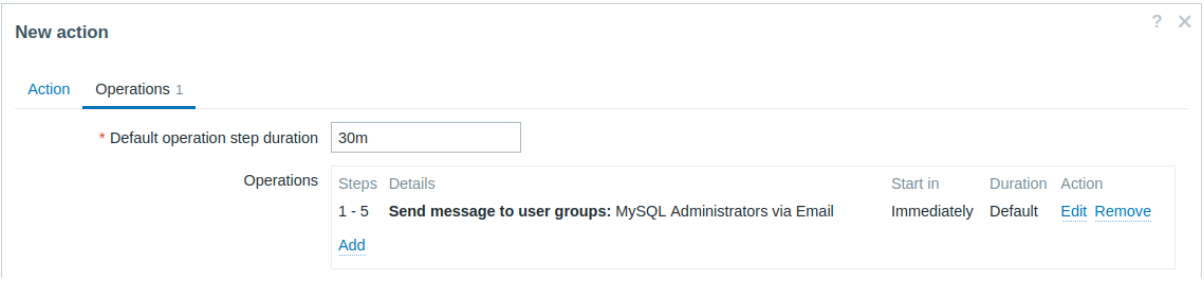
Situation	Behavior
The host in question goes into maintenance after the initial problem notification is sent	Depending on the *Pause operations for suppressed problems* setting in the action configuration, all remaining escalation steps are executed either with a delay caused by the maintenance period or without delay. A maintenance period does not cancel operations.
*The time period defined in the **Time period** action condition ends after the initial notification is sent*	All remaining escalation steps are executed. The *Time period* condition cannot stop operations; it has effect with regard to when actions are started/not started, not operations.
A problem starts during maintenance and continues (is not resolved) after maintenance ends	Depending on the *Pause operations for suppressed problems* setting in the action configuration, all escalation steps are executed either from the moment maintenance ends or immediately.
A problem starts during a no-data maintenance and continues (is not resolved) after maintenance ends	It must wait for the trigger to fire, before all escalation steps are executed.
Different escalations follow in close succession and overlap	The execution of each new escalation supersedes the previous escalation, but for at least one escalation step that is always executed on the previous escalation. This behavior is relevant in actions upon events that are created with EVERY problem evaluation of the trigger.
During an escalation in progress (like a message being sent), based on any type of event: - the action is disabled Based on trigger event: - the trigger is disabled - the host or item is disabled Based on internal event about triggers: - the trigger is disabled Based on internal event about items/low-level discovery rules: - the item is disabled - the host is disabled	The message in progress is sent and then one more message on the escalation is sent. The follow-up message will have the cancellation text at the beginning of the message body (*NOTE: Escalation canceled*) naming the reason (for example, *NOTE: Escalation canceled: action '<Action name>' disabled*). This way the recipient is informed that the escalation is canceled and no more steps will be executed. This message is sent to all who received the notifications before. The reason of cancellation is also logged to the server log file (starting from Debug Level 3=Warning). Note that the *Escalation canceled* message is also sent if operations are finished, but recovery operations are configured and are not executed yet.
During an escalation in progress (like a message being sent) the action is deleted	No more messages are sent. The information is logged to the server log file (starting from Debug Level 3=Warning), for example: `escalation canceled: action id:334 deleted`

Escalation examples

Example 1

Sending a repeated notification once every 30 minutes (5 times in total) to a "MySQL Administrators" group. To configure:

- In *Operations* tab, set the *Default operation step duration* to "30m" (30 minutes).
- Set the escalation *Steps* to be from "1" to "5".
- Select the "MySQL Administrators" group as the recipients of the message.



Notifications will be sent at 0:00, 0:30, 1:00, 1:30, 2:00 hours after the problem starts (unless, of course, the problem is resolved sooner).

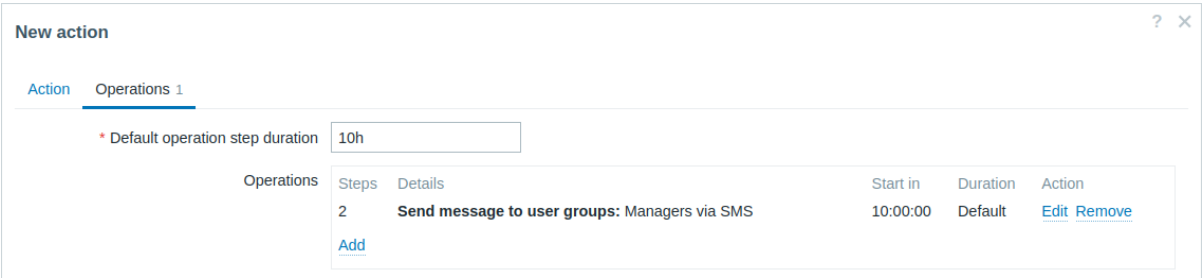
If the problem is resolved and a recovery message is configured, it will be sent to those who received at least one problem message within this escalation scenario.

Note:
If the trigger that generated an active escalation is disabled, Zabbix sends an informative message about it to all those that have already received notifications.

Example 2

Sending a delayed notification about a long-standing problem. To configure:

- In *Operations* tab, set the *Default operation step duration* to "10h" (10 hours).
- Set the escalation *Steps* to be from "2" to "2".



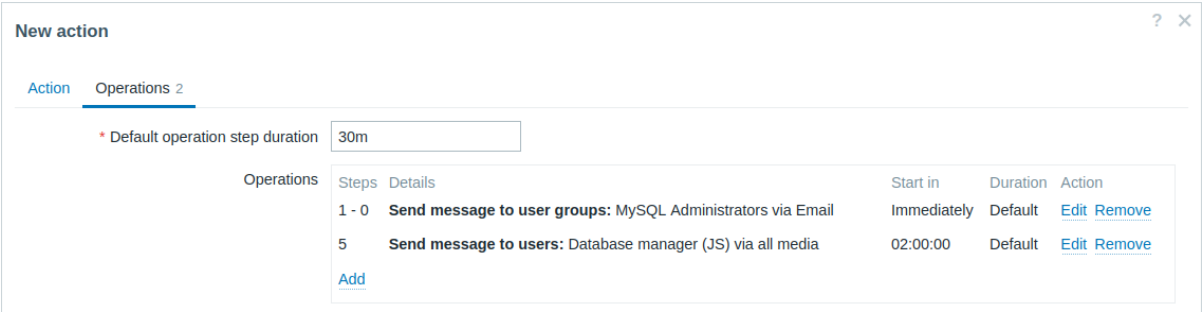
A notification will only be sent at Step 2 of the escalation scenario, or 10 hours after the problem starts.

You can customize the message text to something like "The problem is more than 10 hours old".

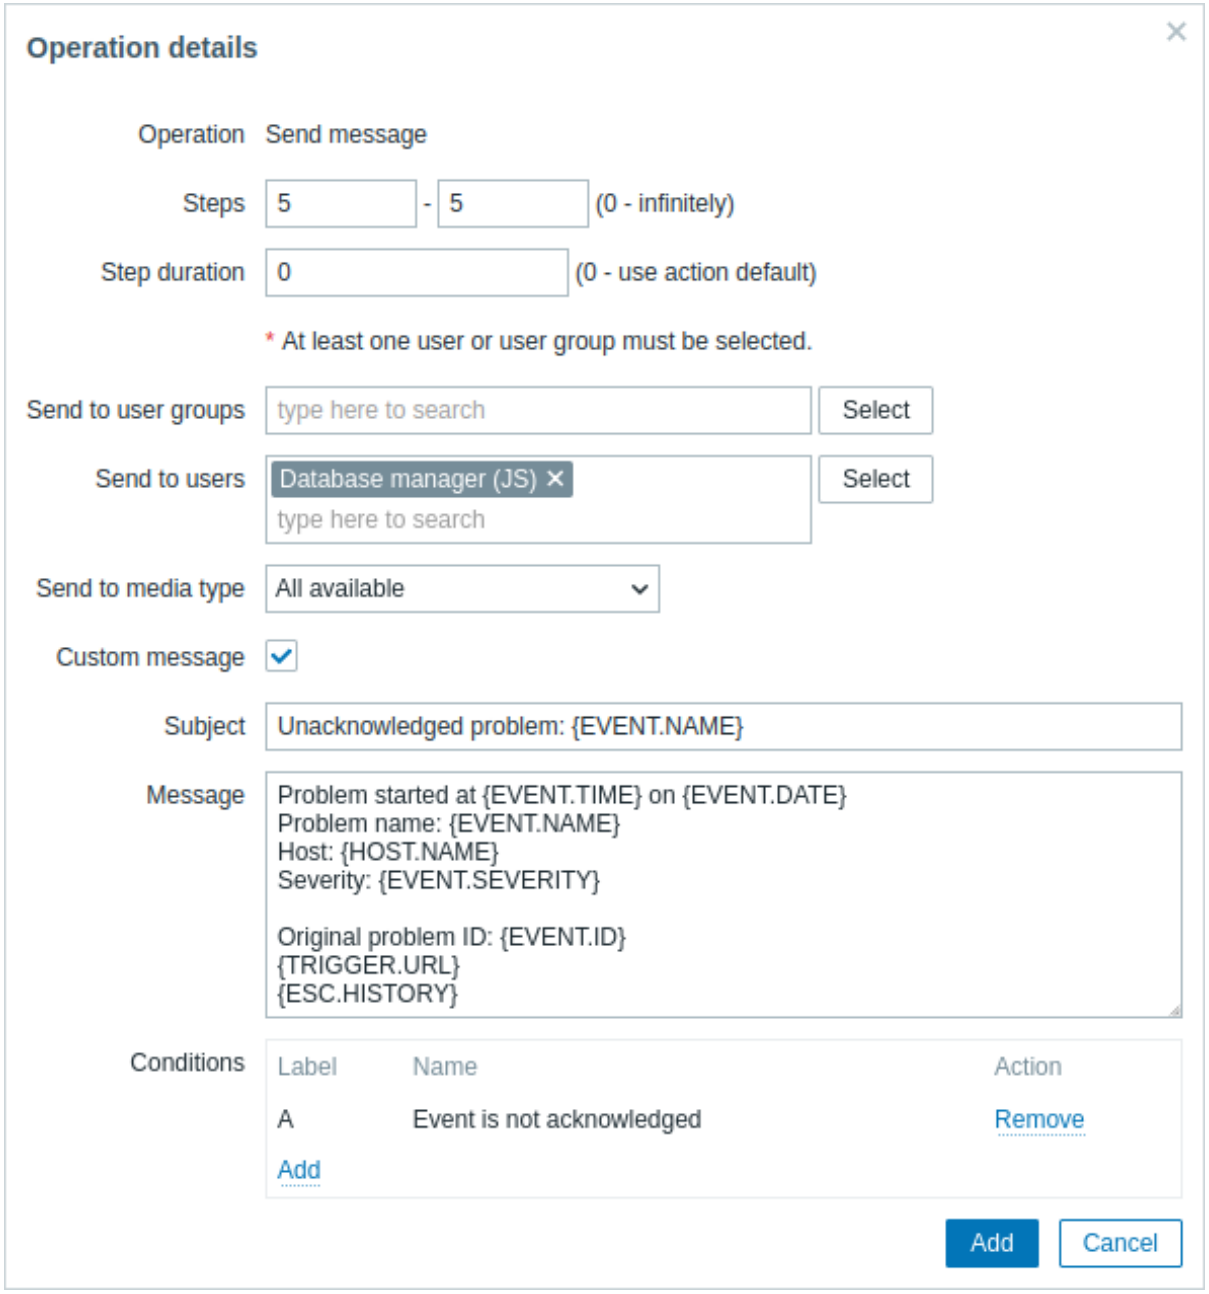
Example 3

Escalating the problem to the Boss.

In the first example above we configured periodical sending of messages to MySQL administrators. In this case, the administrators will get four messages before the problem will be escalated to the Database manager. Note that the manager will get a message only in case the problem is not acknowledged yet, supposedly no one is working on it.



Details of Operation 2:



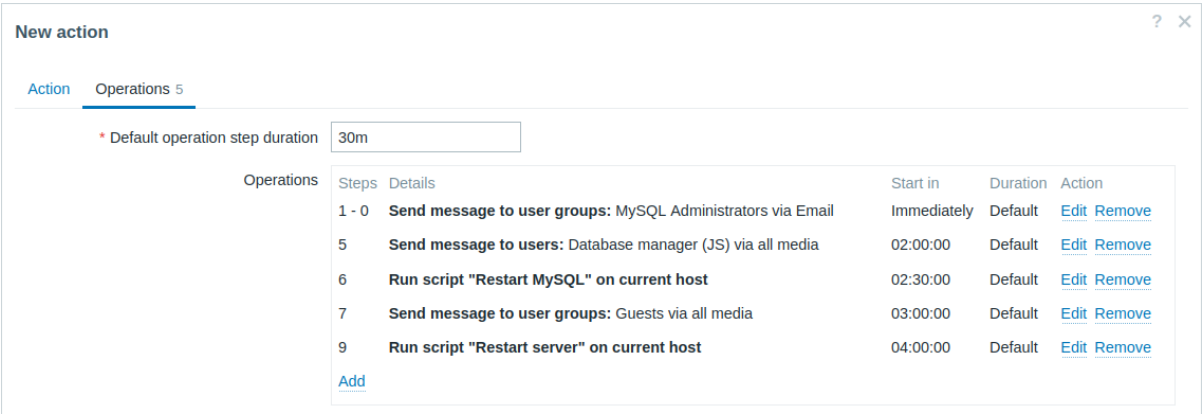
Note the use of {ESC.HISTORY} macro in the customized message. The macro will contain information about all previously executed steps on this escalation, such as notifications sent and commands executed.

Example 4

A more complex scenario. After multiple messages to MySQL administrators and escalation to the manager, Zabbix will try to restart the MySQL database. It will happen if the problem exists for 2:30 hours and it hasn't been acknowledged.

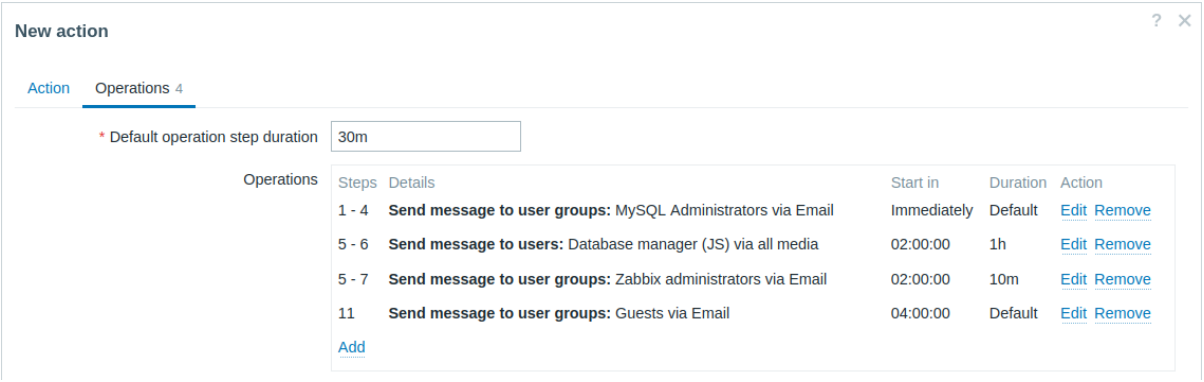
If the problem still exists, after another 30 minutes Zabbix will send a message to all guest users.

If this does not help, after another hour Zabbix will reboot server with the MySQL database (second remote command) using IPMI commands.



Example 5

An escalation with several operations that have overlapping step ranges and custom intervals. The default operation step duration is 30 minutes.



Notifications will be sent as follows:

- To MySQL administrators at 0:00, 0:30, 1:00, and 1:30 after the problem starts.
- To Database manager at 2:00 and 2:10 (the shorter custom step duration of 10 minutes defined in the subsequent operation overrides the longer step duration of 1 hour configured here, as outlined in Operation details for *Step duration* when steps overlap).
- To Zabbix administrators at 2:00, 2:10, and 2:20 after the problem starts (the custom step duration of 10 minutes is applied).
- To guest users at 4:00 after the problem starts (the default step duration of 30 minutes takes effect between steps 8 and 11).

Receiving notification on unsupported items

Overview

It is possible to receive notifications on unsupported items in Zabbix.

It is part of the concept of internal events in Zabbix, allowing users to be notified on these occasions. Internal events reflect a change of state:

- when items go from 'normal' to 'unsupported' (and back);
- when triggers go from 'normal' to 'unknown' (and back);
- when low-level discovery rules go from 'normal' to 'unsupported' (and back).

This section presents a how-to for **receiving notification** when an item turns unsupported.

Configuration

Overall, the process of setting up the notification should feel familiar to those who have set up alerts in Zabbix before.

Step 1

Configure some media, such as email, SMS, or script to use for the notifications. Refer to the corresponding sections of the manual to perform this task.

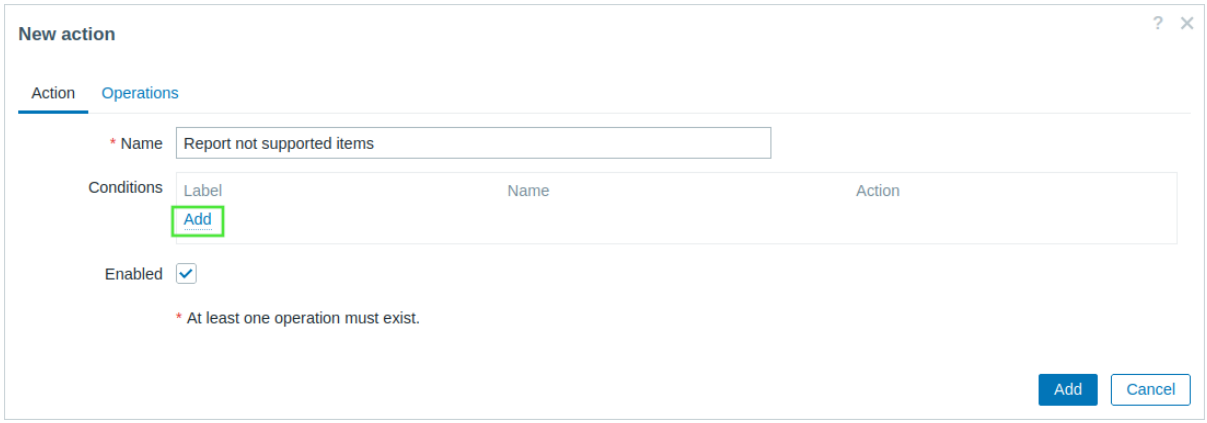
Attention:
For notifying on internal events the default severity ('Not classified') is used, so leave it checked when configuring user media if you want to receive notifications for internal events.

Step 2

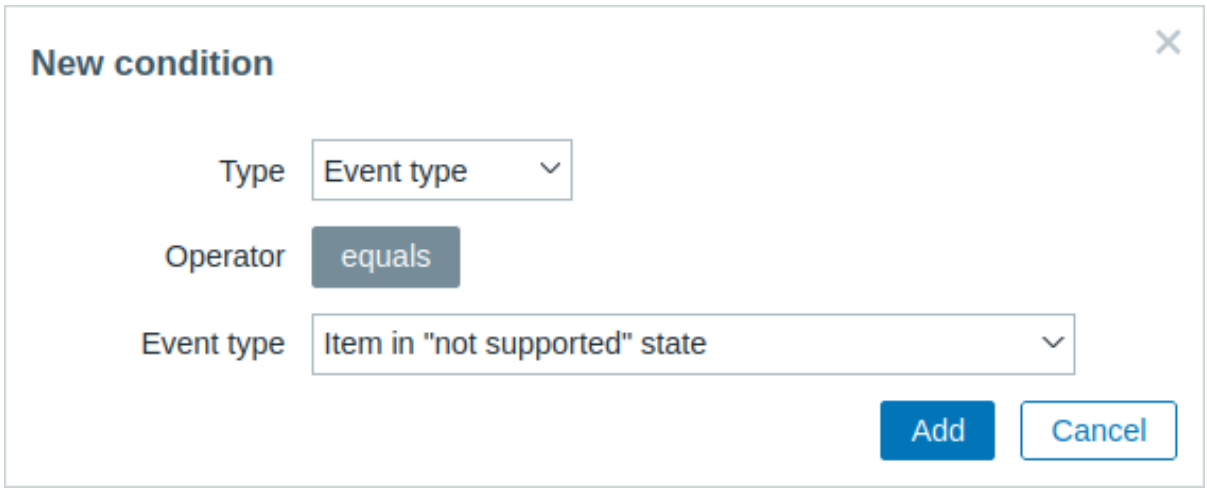
1. Go to *Alerts* > *Actions* > *Internal actions*.
2. Click *Create action* at the upper-right corner of the page to open an action configuration form.

Step 3

1. In the *Action* tab enter a name for the action.
2. Click *Add* in the *Conditions* block to add a new condition.



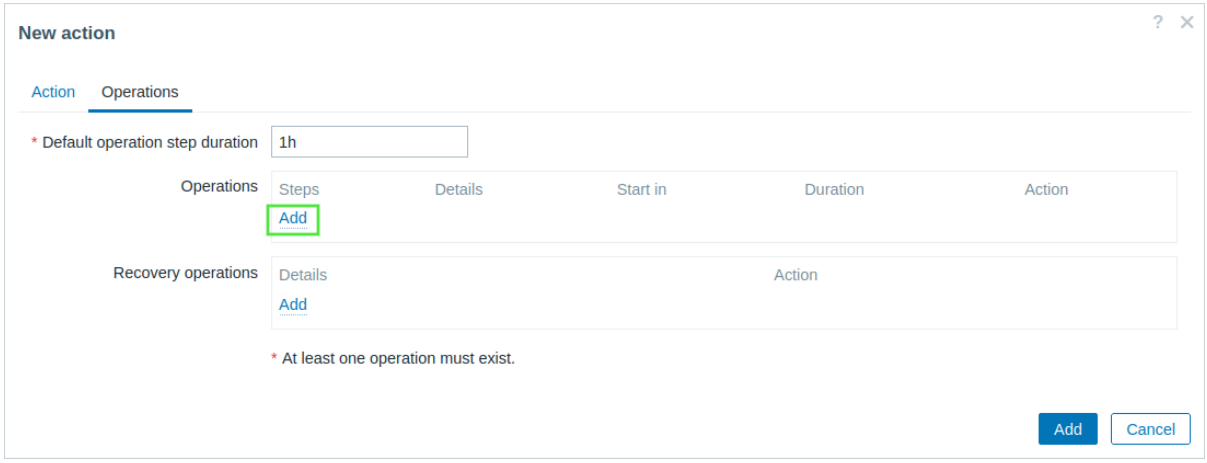
3. In the *New condition* pop-up window select "Event type" as the condition type and then select "Item in 'not supported' state" as the event type.



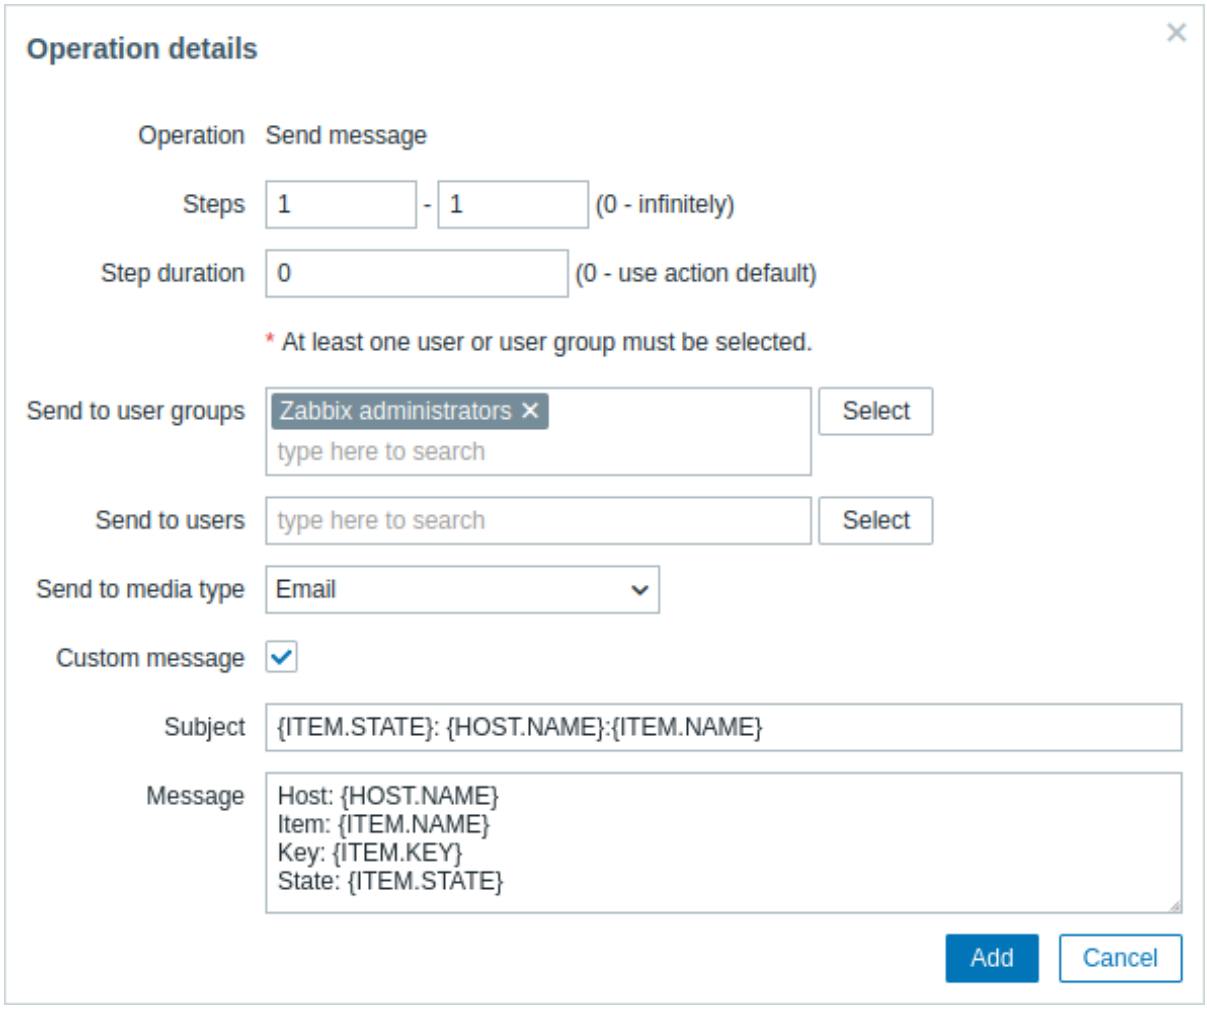
4. Click *Add* to actually list the condition in the *Conditions* block.

Step 4

1. In the *Operations* tab, click *Add* in the *Operations* block to add a new operation.



2. Select some recipients of the message (user groups/users) and the media type (or "All") to use for delivery.
3. Check the *Custom message* checkbox if you wish to enter the custom subject/content of the problem message.



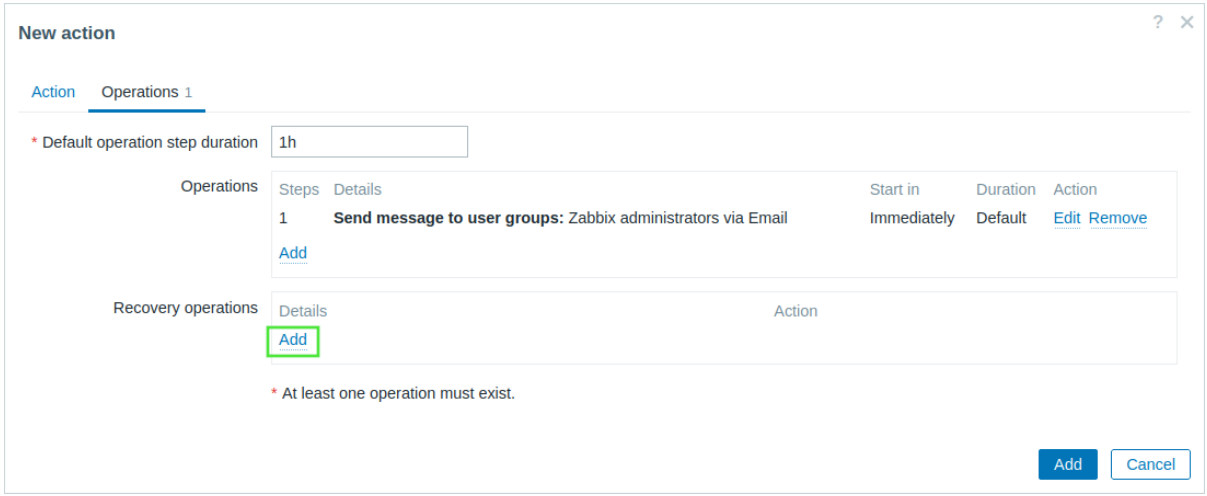
4. Click *Add* to actually list the operation in the *Operations* block.

If you wish to receive more than one notification, set the operation step duration (interval between messages sent) and add another step.

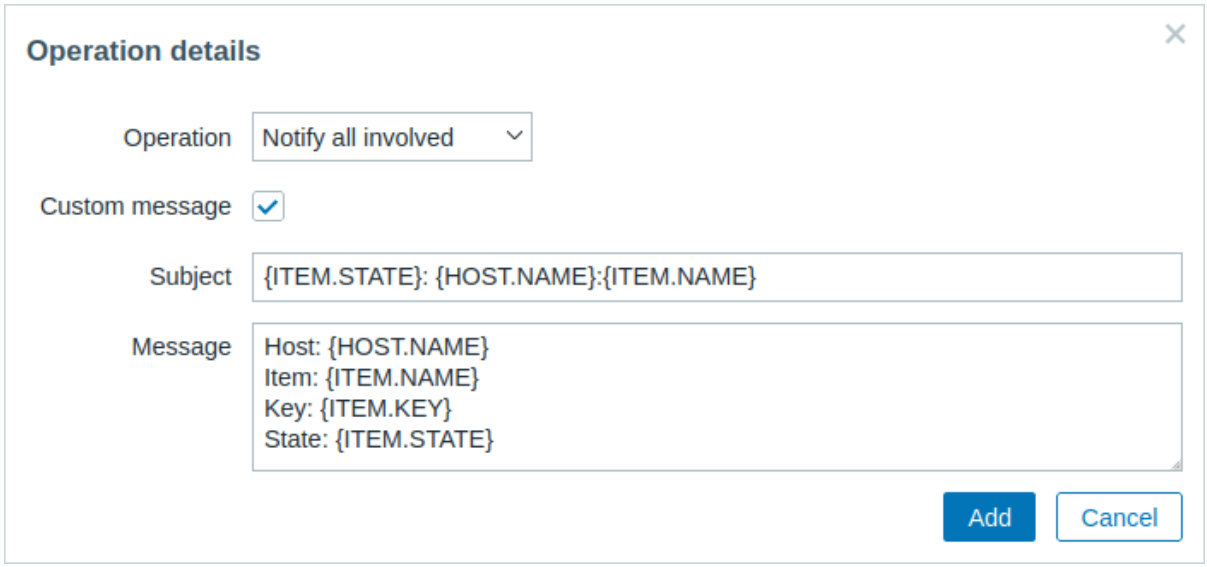
Step 5

The *Recovery operations* block allows you to configure a recovery notification when an item goes back to the normal state.

1. Click *Add* in the *Recovery operations* block to add a new recovery operation.



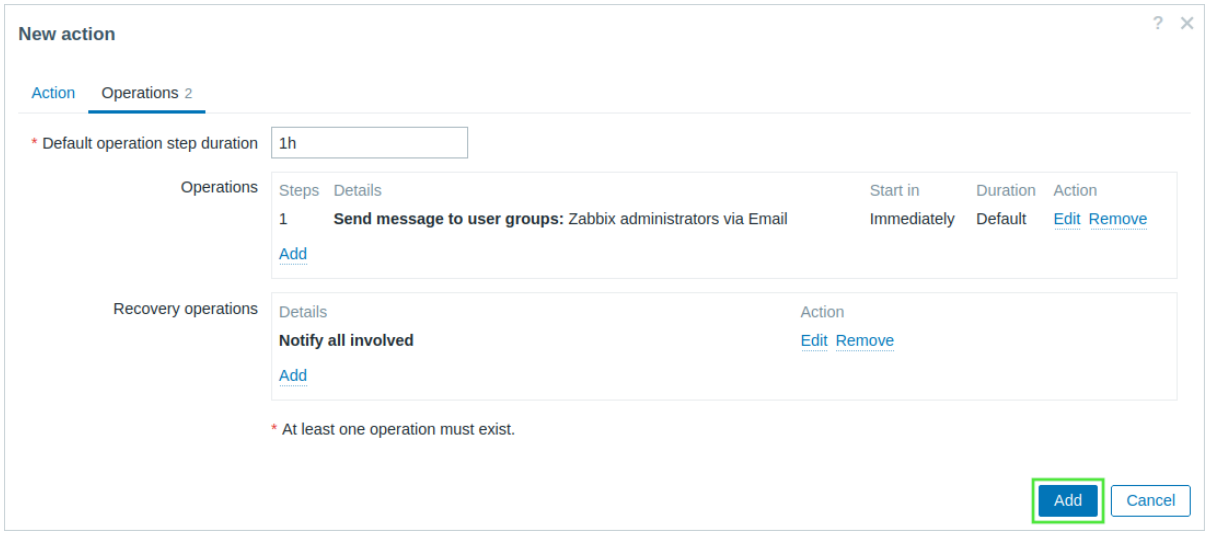
2. Select the operation type ”Notify all involved”.
3. Select *Custom message* checkbox if you wish to enter the custom subject/content of the recovery message.



4. Click *Add* in the *Operation details* pop-up window to actually list the operation in the *Recovery operations* block.

Step 6

When finished, click the *Add* button at the bottom of the form.



And that’s it, you’re done! Now you can look forward to receiving your first notification from Zabbix if some item turns unsupported.

Macros

Overview

Zabbix supports a number of built-in macros which may be used in various situations. These macros are variables, identified by a specific syntax:

`{MACRO}`

Macros resolve to a specific value depending on the context.

Effective use of macros allows you to save time and make Zabbix configuration more transparent.

In one of typical uses, a macro may be used in a template. Thus a trigger on a template may be named "Processor load is too high on {HOST.NAME}". When the template is applied to the host, such as Zabbix server, the name will resolve to "Processor load is too high on Zabbix server" when the trigger is displayed in the Monitoring section.

Macros may be used in item key parameters. A macro may be used for only a part of the parameter, for example, `item.key[server_{HOST.HOST}_local]`. Double-quoting the parameter is not necessary as Zabbix will take care of any ambiguous special symbols, if present in the resolved macro.

There are other types of macros in Zabbix.

Zabbix supports the following macros:

- `{MACRO}` - built-in macro (see full list)
- `{<macro>.<func>(<params>)}` - macro functions
- `{$MACRO}` - user-defined macro, optionally with context
- `{#MACRO}` - macro for low-level discovery
- `{?EXPRESSION}` - expression macro

Macro functions

Overview

Macro functions offer the ability to customize macro values (for example, shorten or extract specific substrings), making them easier to work with.

The syntax of a macro function is:

```
{macro.func(params)}
```

where

- **macro** - the macro to customize;
- **func** - the function to apply (see supported functions);
- **params** - a comma-delimited list of function parameters, which must be **double-quoted** if:
  - start with a space or double quotes;
  - contain closing parentheses or a comma.

For example:

```
{{TIME}.fmttime(format,time_shift)}
{{ITEM.VALUE}.regsub(pattern, output)}
{{$USERMACRO}.regsub(pattern, output)}
{{#LLDMACRO}.regsub(pattern, output)}
```

Macro functions are supported for

- Built-in macros
- User macros
- Low-level discovery macros
- Expression macros

Macro functions can be used in all locations supporting the listed macros. This applies unless explicitly stated that only a macro is expected (for example, when configuring host macros or low-level discovery rule filters).

A single function per macro is supported; multiple macro functions in chain are not supported.

Note:
Please see escaping examples for cases when macro functions are used inside other contexts (function, item key, another macro, etc).

Supported functions

The functions are listed without additional information. Click the function to see the full details.

Function	Description
btoa	Encoding macro value into Base64 encoding.
fmtnum	Number formatting to control the number of digits printed after the decimal point.
fmttime	Time formatting.
htmldecode	Decoding macro value from HTML encoding.
htmlencode	Encoding macro value into HTML encoding.
iregsub	Substring extraction by a regular expression match (case-insensitive).
jsonpath	Return the JSONPath result.
lowercase	Transformation of macro value characters into lowercase.
regrepl	Replacement of character/substring in macro value.
regsub	Substring extraction by a regular expression match (case-sensitive).
tr	Transliteration of macro value characters.
uppercase	Transformation of macro value characters into uppercase.
urldecode	Decoding macro value from URL encoding.
urlencode	Encoding macro value into URL encoding.
xmlxpath	Return the XML XPath result.

Function details

Optional function parameters are indicated by < >.

btoa

Encoding a macro value into Base64 encoding. Base64 encoding is a method for representing binary data as text, useful for storing and secure transmission of binary content over text-based protocols.

Example:

```
{{ITEM.VALUE}.btoa()} - will Base64-encode a value like "zabbix" into "emFiYml4"
```

fmtnum(digits)

Number formatting to control the number of digits printed after the decimal point.

Parameters:

- **digits** - the number of digits after decimal point. Valid range: 0-20. No trailing zeros will be produced.

Examples:

```
{{ITEM.VALUE}.fmtnum(2)} - will return "24.35" from a received value of "24.3483523"
{{ITEM.VALUE}.fmtnum(0)} - will return "24" from a received value of "24.3483523"
```

fmttime(format,<time_shift>)

Time formatting.
 Note that this function can be used with macros that resolve to a value in one of the following time formats:

- `hh:mm:ss`
- `yyyy-mm-ddThh:mm:ss[tz]` (ISO8601 standard)
- UNIX timestamp

Parameters:

- **format** - mandatory format string, compatible with `strftime` function formatting;
- **time_shift** (optional) - the time shift applied to the time before formatting; should start with `-<N><time_unit>` or `+<N><time_unit>`, where:
  - `N` - the number of time units to add or subtract;
  - `time_unit` - h (hour), d (day), w (week), M (month) or y (year).

Comments:

- The `time_shift` parameter supports multistep time operations and may include `/<time_unit>` for shifting to the beginning of the time unit (`/d` - midnight, `/w` - 1st day of the week (Monday), `/M` - 1st day of the month, etc.). Examples: `-1w` - exactly 7 days back; `-1w/w` - Monday of the previous week; `-1w/w+1d` - Tuesday of the previous week.
- Time operations are calculated from left to right without priorities. For example, `-1M/d+1h/w` will be parsed as `((-1M/d)+1h)/w`.

Examples:

```
{{TIMESTAMP}.fmttime(%B)} - will return "October" from a received value of "1633098961"
{{TIMESTAMP}.fmttime(%d %B,-1M/M)} - will return "1 September" from a received value of "1633098961"
{{EVENT.TIMESTAMP}.fmttime("%x %X")} - will return the locale-specific representation of the date and time
{{EVENT.TIME}.fmttime("%x %X")} - will return the locale-specific representation of the date and time; for
```

htmldecode

Decoding a macro value from HTML encoding.

The following characters are supported:

Value	Decoded value
&	&
<	<
>	>
"	"
'	'
'	'

Example:

```
{{ITEM.VALUE}.htmldecode()} - will HTML-decode a value like "&lt;" into "<"
```

htmlencode

Encoding a macro value into HTML encoding.

The following characters are supported:

Value	Encoded value
&	&
<	<
>	>
"	"
'	'

Example:

```
{{ITEM.VALUE}.htmlencode()} - will HTML-encode a character like "<" into "&lt;"
```

iregsub(pattern,output)

Substring extraction by a regular expression match (case-insensitive).

Parameters:

- **pattern** - the regular expression to match;
- **output** - the output options. **\1 - \9** placeholders are supported to capture groups. **\0** returns the matched text.

Comments:

- If there is no match for the regular expression, the function returns an empty string.
- If the function pattern is an incorrect regular expression, then the macro evaluates to 'UNKNOWN' (except for low-level discovery macros, in which case the function will be ignored, and the macro will remain unresolved).
- References to non-existent capture groups in the replacement string are replaced with an empty string.

Example:

```
{{ITEM.VALUE}.iregsub("fail|error|fault|problem","ERROR")} - will resolve to "ERROR" if "fail", "error", "
```

jsonpath(pattern,<default>)

Return the JSONPath result.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **pattern** - the required pattern;
- **default** - the optional fallback value if the JSONPath query returns no data. Note that failure is returned on other errors.

Comments:

- The function fails with an error if:
  - The JSON is invalid.
  - The JSONPath query fails.
  - The JSONPath result contains no data and the `default` parameter is not provided.
- The function is not supported in Zabbix frontend. Wherever macros are resolved by the frontend (for example, trigger operational data in *Monitoring* > *Problems*), the macro resolves to *UNKNOWN*.

Example:

```
{{ITEM.VALUE}.jsonpath("$.columns.servicename", "NA")}
```

lowercase

Transformation of all macro value characters into lowercase. Works with single-byte character sets (such as ASCII) and does not support UTF-8.

Example:

```
{{ITEM.VALUE}.lowercase()} - will transform a value like "Zabbix SERVER" into "zabbix server" (lowercase)
```

regrepl(pattern,replacement,<pattern2>,<replacement2>,...)

Replacement of character/substring in macro value.

Parameters:

- **pattern** - the regular expression to match;
- **replacement** - the replacement string. **\1 - \9** placeholders are supported in replacement strings to capture groups.

Comments:

- All occurrences matching the pattern are replaced.
- The patterns and replacements are processed sequentially, with each subsequent pair being applied in accordance with the outcome of the previous replacement.
- References to non-existent capture groups in the replacement string are replaced with an empty string.

Examples:

```
{{ITEM.VALUE}.regrepl("oldParam", "newParam")} - will replace "oldParam" with "newParam"
{{ITEM.VALUE}.regrepl("([^a-z])","\\\1")} - all non-letter characters will be escaped with a backslash
{$THRESHOLD:"{{#FSNAME}.regrepl(\"\\$\",\"\")}"} - will remove a trailing backslash (for example, to repla
{{ITEM.VALUE}.regrepl("_v1\.0", "_v2.0", "\(final\)", "")} - will replace multiple parts in item value
```

regsub(pattern,output)

Substring extraction by a regular expression match (case-sensitive).

Parameters:

- **pattern** - the regular expression to match;
- **output** - the output options. **\1 - \9** placeholders are supported to capture groups. **\0** returns the matched text.

Comments:

- If there is no match for the regular expression, the function returns an empty string.
- If the function pattern is an incorrect regular expression, then the macro evaluates to 'UNKNOWN' (except for low-level discovery macros, in which case the function will be ignored, and the macro will remain unresolved).
- References to non-existent capture groups in the replacement string are replaced with an empty string.

Examples:

```
{{ITEM.VALUE}.regsub("^([0-9]+)", Problem ID: \1)} - will resolve to "Problem ID: 123" if a value like "12
{{ITEM.VALUE}.regsub("fail|error|fault|problem","ERROR")} - will resolve to "ERROR" if "fail", "error", "f
```

See more examples.

tr(characters,replacement)

Transliteration of macro value characters.

- **characters** - the set of characters to replace;
- **replacement** - the set of positionally corresponding replacement characters.

Examples:

```
{{ITEM.VALUE}.tr(abc, xyz)} - will replace all occurrences of "a" with "x", "b" with "y", "c" with "z"
{{ITEM.VALUE}.tr(abc, xyzq)} - will replace all occurrences of "a" with "x", "b" with "y", "c" with "z" ("
{{ITEM.VALUE}.tr(abcde, xyz)} - will replace all occurrences of "a" with "x", "b" with "y", "c" with "z",
{{ITEM.VALUE}.tr("\\\'", "\/\"")} - will replace all occurrences of backslash with forward slash, single q
{{ITEM.VALUE}.tr(A-Z,a-z)} - will convert all letters to lowercase
{{ITEM.VALUE}.tr(0-9a-z,*)} - will replace all numbers and lowercase letters with "*"
{{ITEM.VALUE}.tr(0-9,ab)} - will replace all occurrences of 0 with "a", and replace all occurrences of 1,
{{ITEM.VALUE}.tr(0-9abcA-L,*)} - will replace all numbers, "abc" characters, and A-L range with "*"
{{ITEM.VALUE}.tr("\n","*")} - will replace end-of-line occurrences with *
{{ITEM.VALUE}.tr("e", "\n")} - will replace all "e" with end-of-line
```

To include literal characters:

```
backslash - must be escaped as \\
single quote - must be escaped as \'
double quote - must be escaped as \"
```

Supported escape sequences with backslash:

```
\\\\ => \\ - double backslash to single backslash
\\a  => \a - alert
\\b  => \b - backspace
\\f  => \f - form feed
\\n  => \n - newline
\\r  => \r - return
\\t  => \t - horizontal tab
\\v  => \v - vertical tab
```

uppercase

Transformation of all macro value characters into uppercase. Works with single-byte character sets (such as ASCII) and does not support UTF-8.

Example:

```
{{ITEM.VALUE}.uppercase()} - will transform a value like "Zabbix Server" into "ZABBIX SERVER" (uppercase)
```

urldecode

Decoding a macro value from URL encoding.

Example:

```
{{ITEM.VALUE}.urldecode()} - will URL-decode a value like "%2F" into "/"
```

urlencode

Encoding a macro value into URL encoding.

Example:

```
{{ITEM.VALUE}.urlencode()} - will URL-encode a character like "/" into "%2F"
```

xmlxpath(pattern,<default>)

Return the XML XPath result.
 Supported value types: *String*, *Text*, *Log*.

Parameters:

- **pattern** - the required pattern;
- **default** - the optional fallback value if the XPath query returns no data. Note that failure is returned on other errors.

Comments:

- The function fails with an error if:
  - The XML is invalid.
  - The XPath query fails.
  - The XPath result contains no data and the `default` parameter is not provided.
  - Zabbix was compiled without XML support.

Example:

```
{{ITEM.VALUE}.xmlxpath("/response/error/status","unknown")}
```

Additional examples

The table below shows more examples of using macro functions.

`{#IFALIAS}` is an LLD macros and is only defined in low-level discovery contexts (discovery rules, prototypes and the items/triggers created from them). Using it outside LLD will leave the token unexpanded.

Macro function	Received value	Output
`{{ITEM.VALUE}.regsub(^[0-9]+, Problem)}`	123Log line	Problem
`{{ITEM.VALUE}.regsub("^([0-9]+)", "Problem")}`	123 Log line	Problem
`{{ITEM.VALUE}.regsub(".*", "Problem ID: \1")}`	Log line	Problem ID:
`{{ITEM.VALUE}.regsub("^(\w+).*?([0-9]+)", " Problem ID: \1_\2 ")}`	MySQL crashed errno 123	Problem ID: MySQL_123
`{{ITEM.VALUE}.regsub("([1-9]+", "Problem ID: \1")}`	123 Log line	*UNKNOWN* (invalid regular expression)
`{{#IFALIAS}.regsub("(.*)_([0-9]+)", \1)}`	customername_1	customername
`{{#IFALIAS}.regsub("(.*)_([0-9]+)", \2)}`	customername_1	1
`{{#IFALIAS}.regsub("(.*)_([0-9]+", \1)}`	customername_1	{{#IFALIAS}.regsub("(.*)_([0-9]+", \1)} (invalid regular expression)
`{$MACRO:"{{#IFALIAS}.regsub(\"(.*)_([0-9]+\", \1)}"}`	customername_1	{$MACRO:"customername"}
`{$MACRO:"{{#IFALIAS}.regsub(\"(.*)_([0-9]+\", \2)}"}`	customername_1	{$MACRO:"1"}
`{$MACRO:"{{#IFALIAS}.regsub(\"(.*)_([0-9]+\", \1)}"}`	customername_1	{$MACRO:"{{#IFALIAS}.regsub(\"(.*)_([0-9]+\", \1)}"} (invalid regular expression)
`"{$MACRO:"\{{#IFALIAS}.regsub("(.*)_([0-9]+", \1)}\"}"`	customername_1	"{$MACRO:"\customername\"}"
`"{$MACRO:"\{{#IFALIAS}.regsub("(.*)_([0-9]+", \2)}\"}"`	customername_1	"{$MACRO:"\1\"}"
`"{$MACRO:\"{{#IFALIAS}.regsub(\\"(.*)_([0-9]+\\", \1)}\"}"`	customername_1	"{$MACRO:\"{{#IFALIAS}.regsub(\\"(.*)_([0-9]+\\", \1)}\"}" (invalid regular expression)

Seeing full item values

Long values of resolved {ITEM.VALUE} and {ITEM.LASTVALUE} macros for text/log items are truncated to 20 characters in some frontend locations. To see the full values of these macros you may use macro functions, e.g.:

```
{{ITEM.VALUE}.regsub("(.*)", \1)}
{{ITEM.LASTVALUE}.regsub("(.*)", \1)}
```

See also: {ITEM.VALUE} and {ITEM.LASTVALUE} macro details.

User macros

Overview

User macros are supported in Zabbix for greater flexibility, in addition to the macros supported out-of-the-box.

User macros can be defined on global, template and host level. These macros have a special syntax:

`{$MACRO}`

Zabbix resolves macros according to the following precedence:

1. Host level macros (checked first).
2. Macros defined for first level templates of the host (i.e., templates linked directly to the host), sorted by template ID.
3. Macros defined for second level templates of the host, sorted by template ID.
4. Macros defined for third level templates of the host, sorted by template ID, etc.
5. Global macros (checked last).

In other words, if a macro does not exist for a host, Zabbix will try to find it in the host templates of increasing depth. If still not found, a global macro will be used, if exists.

Warning:
If a macro with the **same name** exists on multiple linked templates of the same level, the macro from the template with the lowest ID will be used. Thus having macros with the same name in multiple templates is a configuration risk.

If Zabbix is unable to find a macro, the macro will not be resolved.

Attention:
Macros (including user macros) are left unresolved in the Configuration section (for example, in the trigger list) by design to make complex configuration more transparent.

User macros can be used in:

- item name
- item key parameter
- item update intervals and flexible intervals
- trigger name and description
- trigger expression parameters and constants (see examples)
- many other locations - see the full list

Common use cases of global and host macros

- use a global macro in several locations; then change the macro value and apply configuration changes to all locations with one click
- take advantage of templates with host-specific attributes: passwords, port numbers, file names, regular expressions, etc.

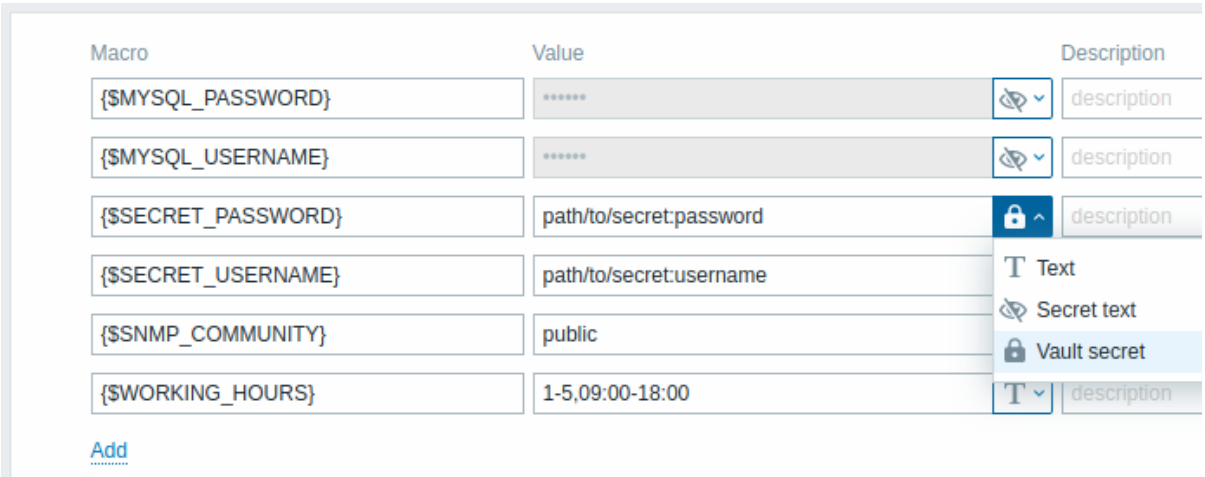
Note:
It is recommended to use host macros instead of global macros because adding, updating or deleting global macros forces incremental configuration update for all hosts. For more information, see Passive and active agent checks.

Configuration

To define user macros, go to the corresponding location in the frontend:

- For global macros, visit *Administration* > *Macros*.
- For host and template level macros, open host or template properties and look for the *Macros* tab.

A user macro has the following attributes:



Parameter	Description
Macro	Macro name. The name must be wrapped in curly brackets and start with a dollar sign. Example: `{$FRONTEND_URL}`. The following characters are allowed in the macro names: **A-Z** (uppercase only) , **0-9** , _ , **.**
Value	Macro value. Three value types are supported: **Text** (default) - plain-text value **Secret text** - the value is masked with asterisks **Vault secret** - the value contains a path/query to a vault secret. To change the value type click the button at the end of the value input field. Maximum length of a user macro value is 2048 characters.
Description	Text field used to provide more information about this macro.

When configuring user macros, **note** the following context-specific behaviors:

- if user macros are used in template items or triggers, consider adding those macros to the template as well (even if they are defined globally); that way, *Text* type macros will work as expected after exporting the template to XML and importing into another system (secret macro values are not exported)
- if user macros are used in trigger expressions, those macros will resolve only if referencing a parameter or constant; they will NOT resolve if referencing a host, item key, function, operator, or another trigger expression (secret macros cannot be used in trigger expressions)
- if user macros are used on a host that has a low-level discovery rule with host prototypes, discovered hosts will inherit all user macros defined on that host

Examples

Example 1

Use of host-level macro in the "Status of SSH daemon" item key:

```
net.tcp.service[ssh,,{$SSH_PORT}]
```

This item can be assigned to multiple hosts, providing that the value of `{$SSH_PORT}` is defined on those hosts.

Example 2

Use of host-level macro in the "CPU load is too high" trigger:

```
last(/ca_001/system.cpu.load[,avg1])>{$MAX_CPULOAD}
```

Such a trigger would be created on the template, not edited in individual hosts.

Note:
If you want to use the amount of values as the function parameter (for example, `max(/host/key,#3)`), include hash mark in the macro definition like this: `SOME_PERIOD => #3`.

Example 3

Use of two macros in the "CPU load is too high" trigger:

```
min(/ca_001/system.cpu.load[,avg1],{$CPULOAD_PERIOD})>{$MAX_CPULOAD}
```

Note that a macro can be used as a parameter of trigger function, in this example function `min()`.

Example 4

Synchronize the agent unavailability condition with the item update interval:

- define `{$INTERVAL}` macro and use it in the item update interval;
- use `{$INTERVAL}` as parameter of the agent unavailability trigger:

```
nodata(/ca_001/agent.ping,{$INTERVAL})=1
```

Example 5

Centralize configuration of working hours:

- create a global `{$WORKING_HOURS}` macro equal to `1-5,09:00-18:00`;
- use it in the *Working time* field in *Administration > General > GUI*;
- use it in the *When active* field in *Users > Users*, *Media* tab of a user;

- use it to set up more frequent item polling during working hours:

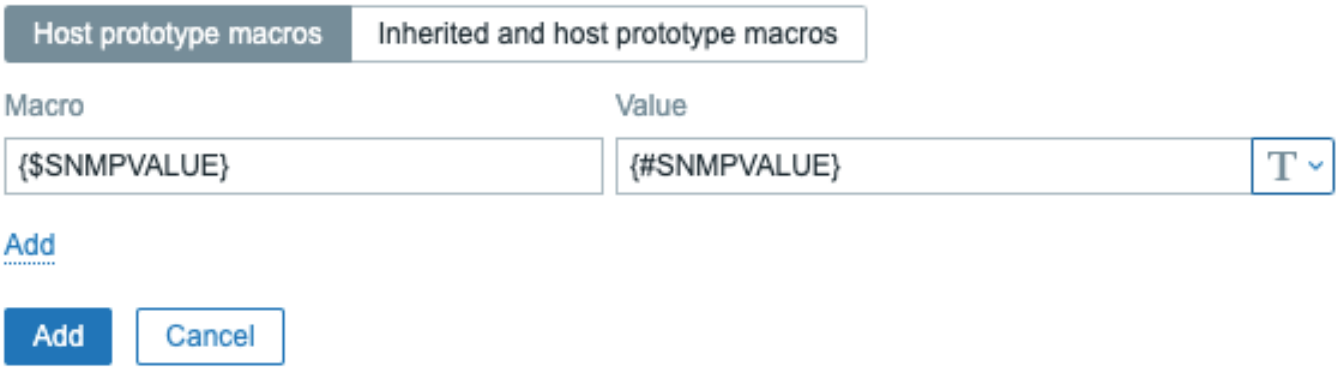


- use it in the *Time period* action condition;
- adjust the working time in *Administration* > *Macros*, if needed.

Example 6

Use a host prototype macro to configure items for discovered hosts:

1. On a host prototype define user macro `{$SNMPVALUE}` with `{#SNMPVALUE}` low-level discovery macro as a value:



2. Assign the *Generic SNMPv2* template to the host prototype.
3. Use `{$SNMPVALUE}` in the *SNMP OID* field of *Generic SNMPv2* template items.

User macro context

See user macros with context.

User macros with context

Overview

An optional context can be used in user macros, allowing the default value to be overridden with a context-specific one.

The context is appended to the macro name; the syntax depends on whether the context is a static text value:

```
{$MACRO:"static text"}
```

or a regular expression:

```
{$MACRO:regex:"regular expression"}
```

Note that a macro with regular expression context can only be defined in user macro configuration. If the `regex:` prefix is used elsewhere as user macro context, like in a trigger expression, it will be treated as static context.

Context quoting is optional (see also important notes).

Macro context examples:

Example	Description
`{$LOW_SPACE_LIMIT}`	User macro without context.
`{$LOW_SPACE_LIMIT:/tmp}`	User macro with context (static string).
`{$LOW_SPACE_LIMIT:regex:"^/tmp$"}`	User macro with context (regular expression). Same as `{$LOW_SPACE_LIMIT:/tmp}`.
`{$LOW_SPACE_LIMIT:regex:"^/var/log/.*$"}`	User macro with context (regular expression). Matches all strings prefixed with /var/log/.

Use cases

User macros with context can be defined to accomplish more flexible thresholds in trigger expressions (based on the values retrieved by low-level discovery). For example, you may define the following macros:

- {$LOW_SPACE_LIMIT} = 10
- {$LOW_SPACE_LIMIT:/home} = 20
- {$LOW_SPACE_LIMIT:regex:"^/[a-z]+$"} = 30

Then a low-level discovery macro may be used as macro context in a trigger prototype for mounted file system discovery:

```
last(/host/vfs.fs.size[{#FSNAME},pfree])<{$LOW_SPACE_LIMIT:"{#FSNAME}"}
```

After the discovery different low-space thresholds will apply in triggers depending on the discovered mount points or file system types. Problem events will be generated if:

- /home folder has less than 20% of free disk space
- folders that match the regexp pattern (like /etc, /tmp or /var) have less than 30% of free disk space
- folders that don't match the regexp pattern and are not /home have less than 10% of free disk space

Important notes

- If more than one user macro with context exists, Zabbix will try to match the simple context macros first and then context macros with regular expressions in an undefined order.

Warning:
Do not create different context macros matching the same string to avoid undefined behavior.

- If a macro with its context is not found on host, linked templates or globally, then the macro without context is searched for.
- Only low-level discovery macros are supported in the context. Any other macros are ignored and treated as plain text.

Technically, macro context is specified using rules similar to item key parameters, except macro context is not parsed as several parameters if there is a , character:

- Macro context must be quoted with " if the context contains a } character or starts with a " character. Quotes inside quoted context must be escaped with the \ character.
- The \ character itself is not escaped, which means it's impossible to have a quoted context ending with the \ character - the macro {$MACRO:"a:\b\c\"} is invalid.
- The leading spaces in context are ignored, the trailing spaces are not:
  - For example, {$MACRO:A} is the same as {$MACRO: A}, but not {$MACRO:A }.
- All spaces before leading quotes and after trailing quotes are ignored, but all spaces inside quotes are not:
  - Macros {$MACRO:"A"}, {$MACRO: "A"}, {$MACRO:"A" } and {$MACRO: "A" } are the same, but macros {$MACRO:"A"} and {$MACRO:" A "} are not.

The following macros are all equivalent, because they have the same context: {$MACRO:A}, {$MACRO: A} and {$MACRO:"A"}. This is in contrast with item keys, where 'key[a]', 'key[a]' and 'key["a"]' are the same semantically, but different for uniqueness purposes.

Secret user macros

Overview

Zabbix provides two options for protecting sensitive information in user macro values:

- Secret text
- Vault secret

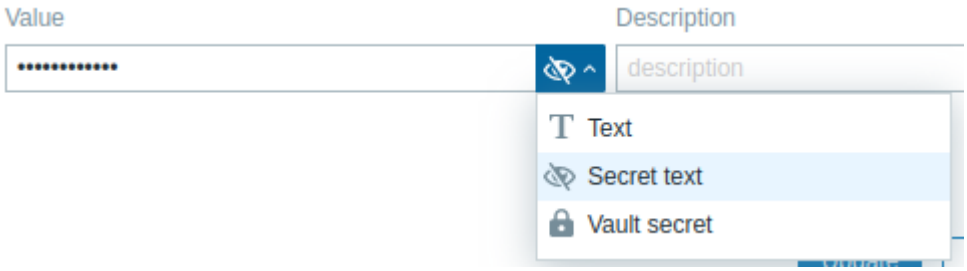
Note:
While the value of a secret macro is hidden, it can be revealed through use in items. For example, in an external script, an `echo` statement referencing a secret macro may be used to reveal the macro value to the frontend, because Zabbix server has access to the real macro value. See locations where secret macro values are unmasked.

Secret macros cannot be used in trigger expressions.

Secret text

With Secret text macros, the macro value is masked with asterisks.

To make a macro value secret, click the button at the end of the *Value* field and select the *Secret text* option:



Once the configuration is saved, it will no longer be possible to view the value.

To change the macro value, hover over the *Value* field and click the *Set new value* button (appears on hover):



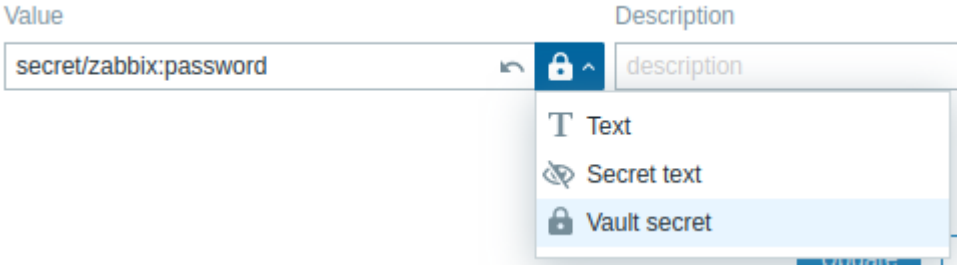
When you click the *Set new value* button (or change the macro value type), the current value will be erased. You can restore the original value by clicking the arrow at the end of the *Value* field (only available before saving the new configuration). Note that restoring the original value will not expose it.

Note:
URLs that contain a secret macro will not work, as the macro in them will be resolved as "******".

Vault secret

With Vault secret macros, the macro value is stored in an external secret management software (vault).

To configure a Vault secret macro, click the button at the end of the *Value* field and select the *Vault secret* option:



The macro value must point to a vault secret. The input format depends on the vault provider. For provider-specific configuration examples, see:

- HashiCorp
- CyberArk

Vault secret macro values are retrieved from the vault by Zabbix server (and Zabbix proxy, if *Resolve secret vault macros by* is set to *Zabbix server and proxy*) on every refresh of configuration data and then stored in the configuration cache. Zabbix server and Zabbix proxy may use different vaults.

If *Resolve secret vault macros by* is set to *Zabbix server*, then vault secrets are retrieved by server only and Zabbix proxy receives values of Vault secret macros from Zabbix server on each configuration sync and stores them in its own configuration cache. That means a Zabbix proxy cannot start data collection after a restart until it receives the configuration update from Zabbix server.

To manually refresh secret values from the vault, use the `secrets_reload` runtime control option (server only).

Encryption must be enabled between Zabbix server and proxy; otherwise, a server warning message is logged.

Warning:
If a macro value cannot be retrieved successfully, the corresponding item using the value will turn unsupported.

Unmasked locations

This list provides locations of parameters where secret macro values are unmasked.

Note:
Secret macro values will remain masked in the locations below if referenced indirectly. For example, {ITEM.KEY}, {ITEM.KEY<1-9>}, {LLDRULE.KEY} built-in macros used in media types (Script or Webhook parameters) will resolve to item keys containing masked secret macros, such as `net.tcp.port[******,******]` instead of `net.tcp.port[192.0.2.1,80]`.

Context		Parameter
Items, item prototypes, LLD rules		
	Item	*Item key parameters*
	Item prototype	*Item prototype key parameters*
	Low-level discovery rule	*Discovery item key parameters*
	SNMP agent	*SNMP community*
		Context name (SNMPv3)
		Security name (SNMPv3)
		Authentication passphrase (SNMPv3)
		Privacy passphrase (SNMPv3)
	HTTP agent	*URL*
		Query fields
		Request body
		Headers
		User name
		Password
		SSL key password
	Script	*Parameters*
		Script
	Browser	*Parameters*
		Script
	Database monitor	*SQL query*
	TELNET agent	*Script*
		User name
		Password
	SSH agent	*Script*
		User name
		Password
	Simple check	*User name*
		Password
	JMX agent	*User name*
		Password
Item value preprocessing		
	JavaScript preprocessing step	*Script*
Web scenarios		
	Web scenario	*Variable value*
		Header value
		URL
		Query field value
		Post field value
		Raw post
	Web scenario authentication	*User*
		Password
		SSL key password
Connectors		
	Connector	*URL*
		Username
		Password
		Token
		HTTP proxy
		SSL certificate file
		SSL key file
		SSL key password
Network discovery		
	SNMP	*SNMP community*

Context		Parameter
		Context name (SNMPv3)
		Security name (SNMPv3)
		Authentication passphrase (SNMPv3)
		Privacy passphrase (SNMPv3)
Global scripts		
	Webhook	*JavaScript script*
		JavaScript script parameter value
	Telnet	*Username*
		Password
	SSH	*Username*
		Password
	Script	*Script*
Media types		
	Script	*Script parameters*
	Webhook	*Parameters*
IPMI management		
	Host	*Username*
		Password

Low-level discovery macros

Overview

There is a type of macro used within the low-level discovery (LLD) function:

```
{#MACRO}
```

It is a macro that is used in an LLD rule and returns real values of the file system name, network interface, SNMP OID, etc.

These macros can be used to create entity prototypes (item, trigger, graph, LLD rule, host, and host group prototypes). During discovery of real file systems, network interfaces, virtual machines, etc., these macros are substituted with real values and are the basis for creating real entities.

Some low-level discovery macros come "pre-packaged" with the LLD function in Zabbix - {#FSNAME}, {#FSTYPE}, {#IFNAME}, {#SNMPINDEX}, {#SNMPVALUE}. However, adhering to these names is not compulsory when creating a custom low-level discovery rule. Then you may use any other LLD macro name and refer to that name.

Supported data types

When defining custom discovery rules, property values returned in JSON objects for LLD macros must be one of the following primitive types:

- string
- number
- boolean

Arrays, objects, and null values are not supported. Any LLD macro referring to such a value will remain unexpanded and will appear literally (e.g., `'{#MY_MACRO}'`) during item preprocessing and creation.

Supported locations

LLD macros can be used:

- in the low-level discovery rule filter
- in nested low-level discovery rules, in
  - JSONPath preprocessing parameters
  - JSONPath field for custom LLD macros
- for item prototypes and discovery prototypes in
  - name
  - key parameters
  - unit
  - update interval[1]
  - timeout[1]
  - history storage period[1]
  - trend storage period[1]
  - item value preprocessing steps

  - SNMP OID
  - IPMI sensor field
  - calculated/aggregate item expression, in:
    * expression constants and function parameters
    * item key parameters
  - aggregate item filter conditions (host group name and tag name)
  - SSH script and Telnet script
  - database monitoring SQL query
  - JMX item endpoint field
  - description
  - HTTP agent URL field
  - HTTP agent HTTP query fields field
  - HTTP agent request body field
  - HTTP agent required status codes field
  - HTTP agent headers field key and value
  - HTTP agent HTTP authentication username field
  - HTTP agent HTTP authentication password field
  - HTTP agent HTTP proxy field
  - HTTP agent HTTP SSL certificate file field
  - HTTP agent HTTP SSL key file field
  - HTTP agent HTTP SSL key password field
  - tags
- for trigger prototypes in
  - name
  - operational data
  - expression (only in constants and function parameters)
  - URL
  - description
  - tags
- for graph prototypes in
  - name
- for host prototypes in
  - name
  - visible name
  - custom interface fields: IP, DNS, port, SNMPv1/SNMPv2 community, SNMPv2/SNMPv3 max repetition count, SNMPv1/SNMPv2/SNMPv3 retries, SNMPv3 context name, SNMPv3 security name, SNMPv3 authentication passphrase, SNMPv3 privacy passphrase
  - host group prototype name
  - host tag value
  - host macro value
  - (see the full list)

In all those places, except the low-level discovery rule filter, LLD macros can be used inside static user macro context.

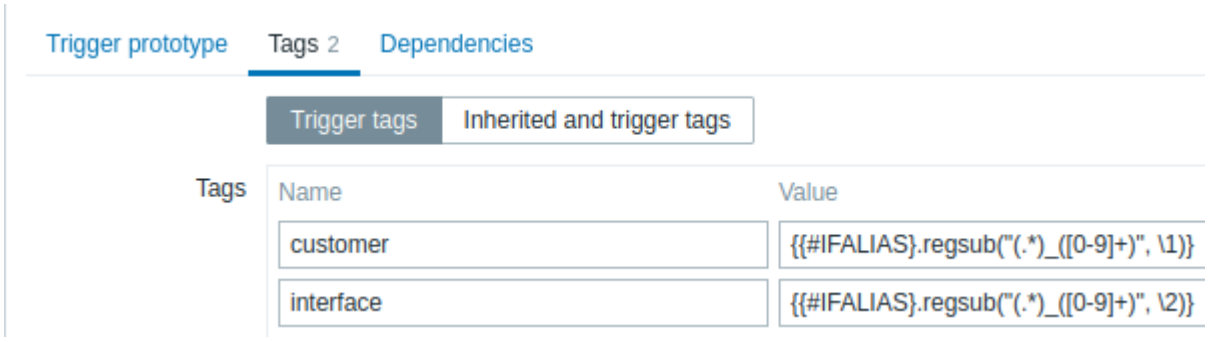
Using macro functions

Macro functions are supported with low-level discovery macros (except in low-level discovery rule filter), allowing a certain part of the macro value to be extracted using a regular expression.

For example, you may want to extract the customer name and interface number from the following LLD macro for the purposes of event tagging:

```
{#IFALIAS}=customername_1
```

To do so, the `regsub` macro function can be used with the macro in the event tag value field of a trigger prototype:



Note that commas are not allowed in unquoted item key parameters, so the parameter containing a macro function has to be quoted. The backslash (\) character should be used to escape double quotes inside the parameter. Example:

```
net.if.in["{{#IFALIAS}.regsub(\"(.*)_([0-9]+)\", \1)}",bytes]
```

For more information on macro function syntax, see: Macro functions

Macro functions are supported in low-level discovery macros since Zabbix 4.0.

Item prototypes with no key parameters

When item key parameters are not used, place the LLD macro inside the item key parameter brackets `[...]` so it is treated as a parameter and substituted during discovery:

```
v_[{#MACRO}]
```

Footnotes

1 In the fields marked with [1] a single macro has to fill the whole field. Multiple macros in a field or macros mixed with text are not supported.

Expression macros

Overview

Expression macros let you perform calculations in fields.

Their value is calculated by first resolving any inner macros and then evaluating the resulting expression.

Syntax:

```
{?EXPRESSION}
```

EXPRESSION uses the same syntax and supports the same functions as trigger expressions.

Example:

```
{?trendavg(/host/item1,1M:now/M)/trendavg(/host/item1,1M:now/M-1y)*100}
```

For locations that support expression macros, search for "{?EXPRESSION}" in the supported macro table.

Notes on usage:

- Use {FUNCTION.*} macros to reference function values of trigger expressions/recovery expressions.
- Use {HOST.HOST<1-9>} and {ITEM.KEY<1-9>} macros to reference hosts and items.
- In templates, use {HOST.HOST<1-9>} macros or omit the host altogether for the first host—for example, `{?avg(//item1,1h)}`—instead of template names, as template names are not replaced with host names during template linking.

```
{?{FUNCTION.VALUE2} - {FUNCTION.VALUE3}}
{?max(/{HOST.HOST}/{ITEM.KEY},3h)}
```

See also Trigger expression examples for an example of using an expression macro in an event name.

Users and user groups

Overview

All users in Zabbix access the Zabbix application through the web-based frontend. Each user is assigned a unique login name and a password.

All user passwords are encrypted and stored in the Zabbix database. Users cannot use their user id and password to log directly into the UNIX server unless they have also been set up accordingly to UNIX. Communication between the web server and the user browser can be protected using SSL.

With a flexible user permission schema, you can restrict and differentiate rights to:

- Access administrative Zabbix frontend functions.
- Perform certain actions in the frontend.
- Access monitored hosts in host groups.
- Use specific API methods.

Configuring a user

Overview

The initial Zabbix installation has two predefined users:

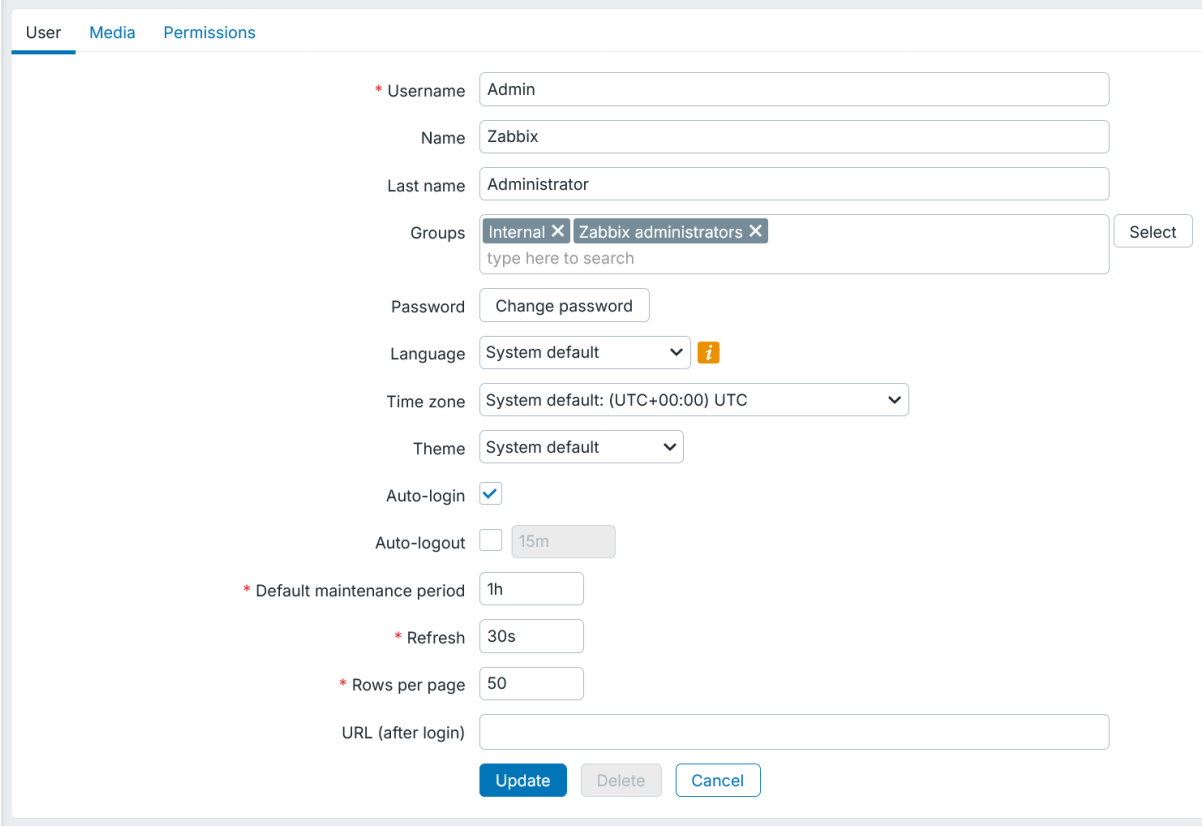
- *Admin* - a Zabbix superuser with full permissions.
- *guest* - a special Zabbix user. The 'guest' user is disabled by default. To enable the user, remove the predefined *Disabled* group from the user groups. You may log in with this user and access monitoring pages in Zabbix. Note that by default, 'guest' has no permissions on Zabbix objects.

To configure a user:

1. Go to *Users*> *Users*.
2. Click *Create user* (or on a user name to edit an existing user).
3. Edit user attributes in the form.

General attributes

The *User* tab contains general user attributes:



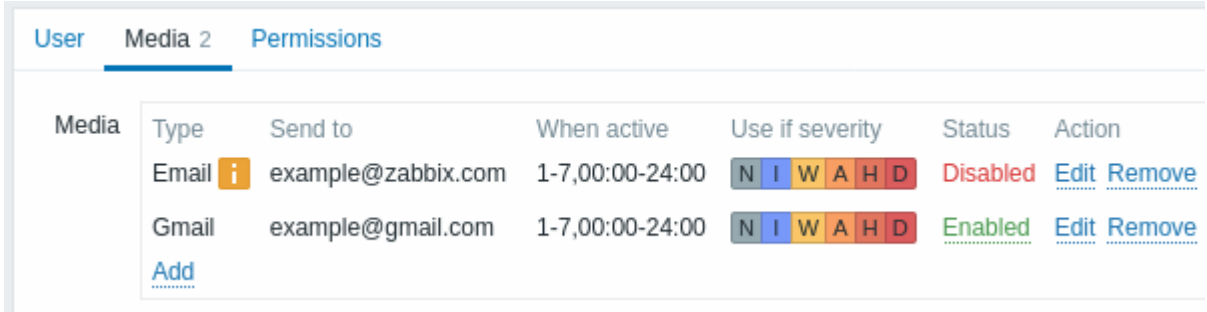
All mandatory input fields are marked with a red asterisk.

Parameter	Description
Username	Unique username, used as the login name.
Name	User first name (optional). If not empty, visible in acknowledgment information and notification recipient information.
Last name	User last name (optional). If not empty, visible in acknowledgment information and notification recipient information.
Groups	Select user groups the user belongs to. This field is auto-complete so starting to type the name of a user group will offer a dropdown of matching groups. Scroll down to select. Alternatively, click *Select* to add groups. Click 'x' to remove the selected. Adherence to user groups determines what host groups and hosts the user will have access to.

Parameter	Description
Password	Two fields for entering the user password, or a *Change password* button if the user already exists. Clicking the *Change password* button opens two fields for entering a new password. For the user with the *Super admin role* changing own password, clicking the *Change password* button opens an additional field for entering the current (old) password. On a successful password change, the user for which the password was changed will be logged out of all active sessions. Note that the password can only be changed for users using Zabbix internal authentication.
Language	Language of the Zabbix frontend. The php gettext extension is required for the translations to work.
Time zone	Select the time zone to override global time zone on user level or select **System default** to use global time zone settings.
Theme	Defines how the frontend looks like: **System default** - use default system settings **Blue** - standard blue theme **Blue (classic)** - standard blue theme using legacy fonts **Dark** - alternative dark theme **Dark blue** - dark blue theme **Dark (classic)** - alternative dark theme using legacy fonts **High-contrast light** - light theme with high contrast **High-contrast dark** - dark theme with high contrast
Auto-login	Mark this checkbox to make Zabbix remember the user and log the user in automatically for 30 days. When logging in with *Remember for 30 days*: - Auto-logout is reset (session persists for 30 days). - Auto-login is enabled for seamless re-authentication. When logging in without *Remember for 30 days*: - Auto-login is enabled for seamless re-authentication. - Auto-logout remains governed by the standard timeout setting. Browser cookies are used for this.
Auto-logout	With this checkbox marked the user will be logged out automatically, after the set amount of seconds (minimum 90 seconds, maximum 1 day). Note, that this setting is overridden if *Remember for 30 days* is enabled, as the session is extended for the full period. Time suffixes are supported, e.g., 90s, 5m, 2h, 1d. Note that this option will not work: * If the ”Show warning if Zabbix server is down” global configuration option is enabled and Zabbix frontend is kept open. * When Monitoring menu pages perform background information refreshes. * If logging in with the *Remember me for 30 days* option checked.
Default maintenance period	Set the default maintenance period for ad-hoc maintenance.
Refresh	Set the refresh rate used for graphs, plain text data, etc. Can be set to 0 to disable. Time suffixes are supported, e.g., 90s, 5m, 1h.
Rows per page	You can determine how many rows per page will be displayed in lists.
URL (after login)	You can make Zabbix transfer the user to a specific URL after successful login, for example, to *Problems* page.

User media

The *Media* tab contains a listing of all media defined for the user. Media are used for sending notifications.



Click *Add* to assign media to the user.

If the media type has been disabled:

- A yellow info icon is displayed after the name.
- *Disabled* is displayed in the Status column.

> **Note:**
> User permissions to change media details for themselves can be granted/revoked based on their user role (see *Create and edit own media* option). Super admin user permissions to change media details for others can also be granted/revoked based on their user role (see *Create and edit user media* option)

See the Media types section for details on configuring user media.

Permissions

The *Permissions* tab contains information on the following elements:

- User role (mandatory for any newly created user) that can only be changed by a *Super admin* type user.

> **Attention:**
> Users cannot be created without a user role (except with Zabbix User API). Previously created users which do not have a role may still be edited without assigning a role to them. However, once a role is assigned, it can only be changed, not removed.

 Note that users without a role can log into Zabbix only using LDAP or SAML authentication, provided their LDAP/SAML information matches the user group mappings configured in Zabbix.

- User type (*User*, *Admin*, *Super admin*) that is defined in the user role configuration.
- Host and template groups that the user has access to.
  - *User* and *Admin* type users, by default, do not have access to any groups, templates, and hosts. To grant such access, users must be included in user groups configured with permissions to the relevant entities.
- Access rights to sections and elements of Zabbix frontend, modules, and API methods.
  - Elements with allowed access are displayed in green, while those with denied access - in light gray.
  - If access to API is enabled, the API methods of the user role are listed under *Allowed methods* or *Denied methods*, depending on whether *Allow list* or *Deny list* is selected in the user role configuration. If no API methods are specified, *None* is displayed. *None* under *Allowed methods* means that no API methods are allowed, while *None* under *Denied methods* means that all API methods available for the user type are allowed.
- Rights to perform specific actions.
  - Actions that the user is allowed to perform are displayed in green, while those that are denied - in light gray.

See the Permissions page for details.

Permissions

Overview

Permissions in Zabbix depend on the user type, customized user roles and access to hosts, which is specified based on the user group.

User types

Permissions in Zabbix depend, primarily, on the user type:

- *User* - has limited access rights to menu sections (see below) and no access to any resources by default. Any permissions to host or template groups must be explicitly assigned.
- *Admin* - has incomplete access rights to menu sections (see below). The user has no access to any host groups by default. Any permissions to host or template groups must be explicitly given.
- *Super admin* - has access to all menu sections. The user has a read-write access to all host and template groups. Permissions cannot be revoked by denying access to specific groups.

Menu access

> **Note:**
> Restricted access to some UI elements only prevents opening that page—it does not remove the ability to access underlying data on other parts of the interface.

The following table illustrates access to Zabbix menu sections per user type:

Menu section		User	Admin	Super admin
Dashboards		+	+	+
Monitoring		+	+	+
	Problems	+	+	+
	Hosts	+	+	+
	Latest data	+	+	+
	Maps	+	+	+
	Discovery		+	+
APM		+	+	+
	Traces	+	+	+
	Metrics	+	+	+
	Logs	+	+	+
Services		+	+	+
	Services	+	+	+
	SLA		+	+
	SLA report	+	+	+
Inventory		+	+	+
	Overview	+	+	+
	Hosts	+	+	+
Reports		+	+	+
	System information		+	+
	Scheduled reports		+	+
	Availability report	+	+	+
	Top 100 triggers	+	+	+
	Audit log			+
	Action log			+
	Notifications		+	+
Data collection			+	+
	Template groups		+	+
	Host groups		+	+
	Templates		+	+
	Hosts		+	+
	Maintenance		+	+
	Event processing			+
	Discovery		+	+
Alerts			+	+
	Trigger actions		+	+
	Service actions		+	+
	Discovery actions		+	+
	Autoregistration actions		+	+
	Internal actions		+	+
	Media types			+
	Scripts			+
Users				+
	User groups			+
	User roles			+
	Users			+
	API tokens			+
	Authentication			+
	Devices			+
Administration				+
	Data source			+
	General			+
	Audit log			+
	Housekeeping			+
	Proxy groups			+
	Proxies			+
	Macros			+
	Queue			+

User roles

User roles allow making custom adjustments to the permissions defined by the user type. While no permissions can be added (that would exceed those of the user type), some permissions can be revoked.

Furthermore, a user role determines access not only to menu sections, but also to services, modules, API methods and various actions in the frontend.

User roles are configured in the *Users* > *User roles* section by Super admin users.

User roles are assigned to users in the user configuration form, *Permissions* tab, by Super admin users.

Users

User | Media | Permissions

* Role: Admin role ✕ | Select

User type: Admin

Permissions

User group permissions

Group	Type	Permissions
All groups	Hosts	None
All groups	Templates	None

Proxies: None

Proxy groups: None

Permissions can be assigned for user groups only.

Access to UI elements

Dashboards: Dashboards

Monitoring: Problems | Hosts | Latest data | Maps | Discovery

APM: Traces | Metrics | Logs

Services: Services | SLA | SLA report

Inventory: Overview | Hosts

Reports: System information | Scheduled reports | Availability report | Top 100 triggers | Notifications

Data collection: Template groups | Host groups | Templates | Hosts | Maintenance | Discovery

Alerts: Trigger actions | Service actions | Discovery actions | Autoregistration actions | Internal actions

User profile settings

Redirect URL after login

Access to services

Read-write access to services: None

Read-only access to services: All

Access to modules

Action log | Clock | Discovery status | Favorite graphs | Favorite maps | Gauge | Geomap | Graph | Graph (classic) | Graph prototype | Honeycomb | Host availability | Host card | Host navigator | Item card | Item history | Item navigator | Item value | Map | Map navigation tree | Pie chart | Problem hosts | Problems | Problems by severity | Scatter plot | SLA report | System information | Top hosts | Top items | Top triggers | Trigger overview | URL | Web monitoring

Access to API

Disabled

Devices

Enabled | Manage own devices

Access to actions

Create and edit dashboards | Create and edit maps | Create and edit maintenance | Add problem comments | Change severity | Acknowledge problems | Suppress problems | Close problems | Execute scripts | Manage API tokens | Manage scheduled reports | Manage SLA | Invoke "Execute now" on read-only hosts | Change problem ranking | Create and edit own media | Select "Server" for monitoring and discovery

Add | Cancel

Access to hosts

Access to any host and template data in Zabbix is granted to user groups on the host/template group level only.

That means that an individual user cannot be directly granted access to a host (or host group). It can only be granted access to a host by being part of a user group that is granted access to the host group that contains the host.

Similarly, a user can only be granted access to a template by being part of a user group that is granted access to the template group that contains the template.

User groups

Overview

User groups allow you to group users both for organizational purposes and for assigning permissions to data. Permissions to viewing and configuring data of host groups and template groups are assigned to user groups, not individual users.

It may often make sense to separate what information is available for one group of users and what - for another. This can be accomplished by grouping users and then assigning varied permissions to host and template groups.

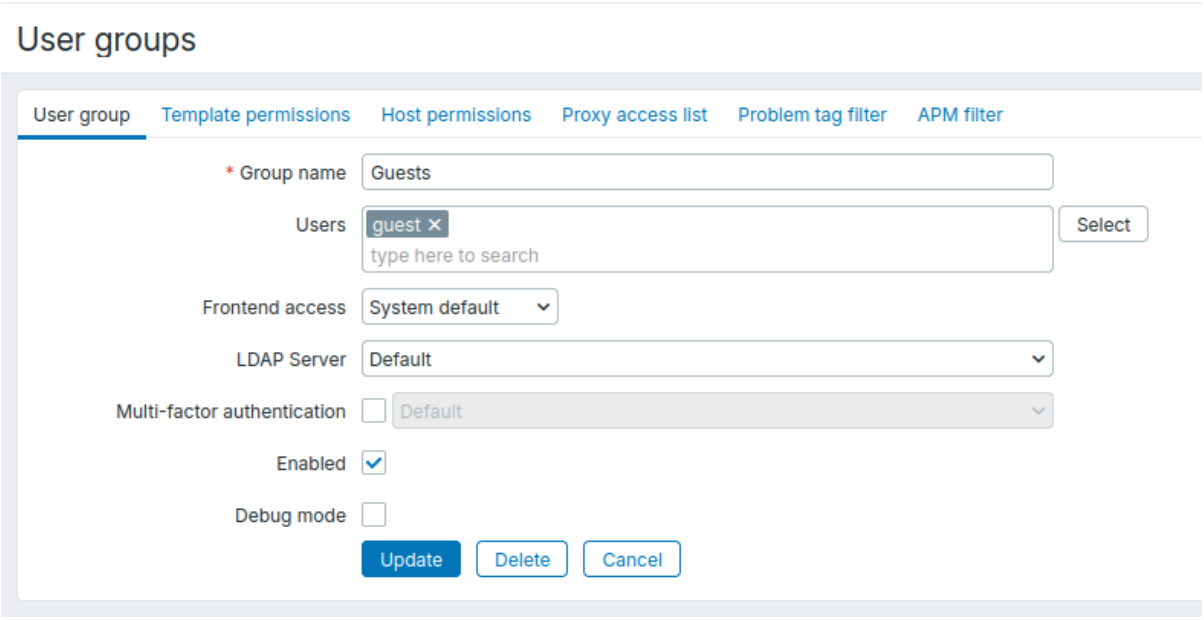
A user can belong to any number of groups.

Configuration

To configure a user group:

1. Go to *Users* > *User groups*.
2. Click *Create user group* (or the group name to edit an existing group).
3. Edit group attributes in the form.

The **User group** tab contains general group attributes:

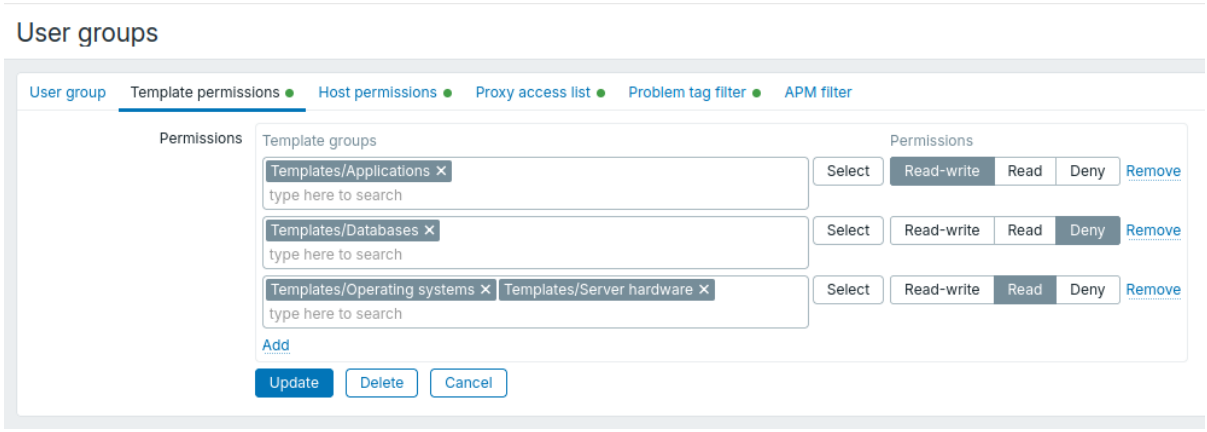


All mandatory input fields are marked with a red asterisk.

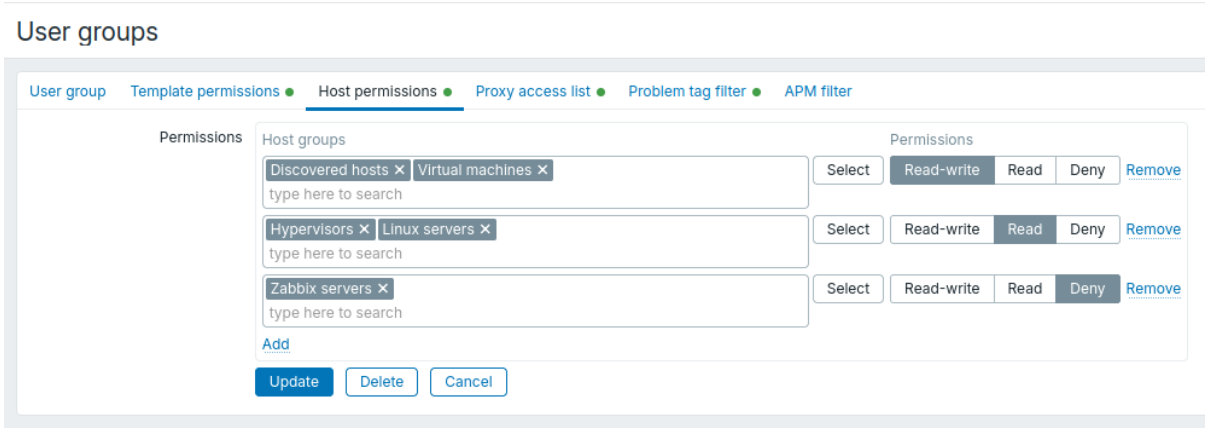
Parameter	Description
Group name	Unique group name.
Users	To add users to the group start typing the name of an existing user. When the dropdown with matching user names appears, scroll down to select. Alternatively you may click the *Select* button to select users in a popup.
Frontend access	How the users of the group are authenticated. **System default** - use default authentication method (set globally) **Internal** - use Zabbix internal authentication (even if LDAP authentication is used globally). Ignored if HTTP authentication is the global default. **LDAP** - use LDAP authentication (even if internal authentication is used globally). Ignored if HTTP authentication is the global default. **Disabled** - access to Zabbix frontend is forbidden for this group

Parameter	Description
LDAP server	Select which LDAP server to use to authenticate the user. This field is enabled only if *Frontend access* is set to LDAP or System default.
Multi-factor authentication	Select which multi-factor authentication method to use to authenticate the user: **Default** - use the method set as default in MFA configuration; this option is selected by default for new user groups if MFA is enabled; **<Method name>** - use selected method (for example, "Zabbix TOTP"); **Disabled** - MFA is disabled for this group; this option is selected by default for new user groups if MFA is disabled. Note that if a user belongs to multiple user groups with MFA enabled (or at least one group has MFA enabled), the following authentication rules apply: if any group uses the "Default" MFA method, it will authenticate the user; otherwise, the first method (ordered alphabetically) will be used for authentication.
Enabled	Status of user group and group members. *Checked* - user group and users are enabled *Unchecked* - user group and users are disabled
Debug mode	Mark this checkbox to activate debug mode for the users.

The **Template permissions** tab allows you to specify user group access to template groups and, through them, to template data:



The **Host permissions** tab allows you to specify user group access to host groups and, through them, to host data:



Click Add to choose the template/host groups (be it a parent or a nested group) and assign permissions to those. Start typing the group name (a dropdown of matching groups will appear) or click *Select* for a popup window listing all groups to be opened.

Then use the option buttons to assign permissions to the chosen groups. Possible permissions are the following:

- **Read-write** - read-write access to a group
- **Read** - read-only access to a group
- **Deny** - access to a group denied

If the same template/host group is added in several rows with different permissions set, the strictest permission will be applied.

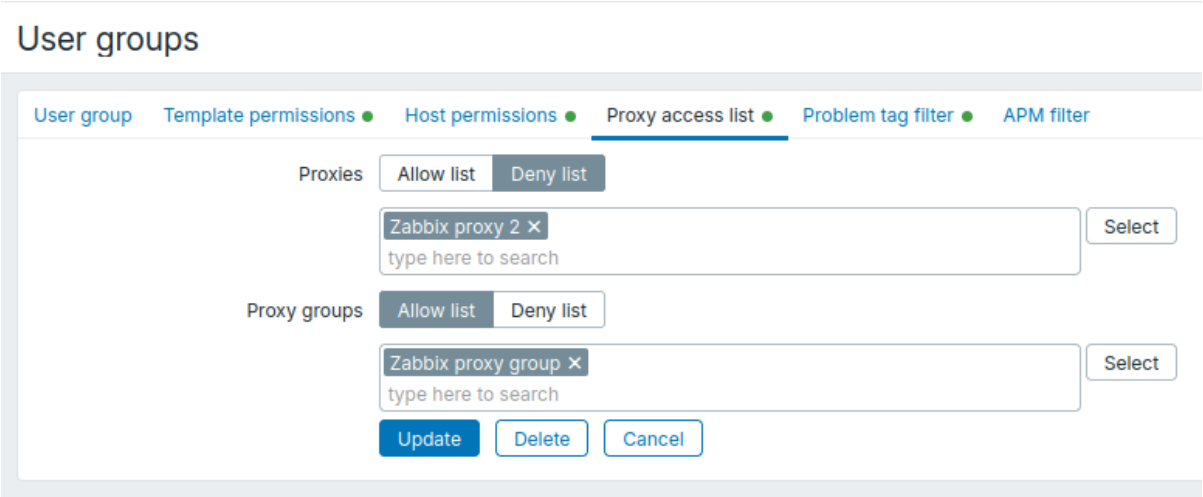
Note that a *Super admin* user can enforce nested groups to have the same level of permissions as the parent group; this can be done in the host/template group configuration form.

The **Template permissions** and **Host permissions** tabs support the same set of parameters.

Current permissions to groups are displayed in the *Permissions* block, and those can be modified or removed.

Note:
If a user group has **Read-write** permissions to a host and **Deny** or no permissions to a template linked to this host, the users of such group will not be able to edit templated items on the host, and template name will be displayed as *Inaccessible template*.

The **Proxy access list** tab allows you to specify user group access to proxies, either individually or through proxy groups:



Use the *Proxy groups* field to apply the permission to the selected groups. A proxy that belongs to a group can only be allowed or denied through that group. Use the *Proxies* field to add individual proxies that do not belong to any proxy group.

Proxies and *Proxy groups* support the same set of parameters.

Click the *Select* button next to the *Proxies* or *Proxy groups* field to select proxies or proxy groups, respectively. Alternatively, start typing the name of a proxy or proxy group in the corresponding field. A dropdown list of matching proxies or proxy groups appears.

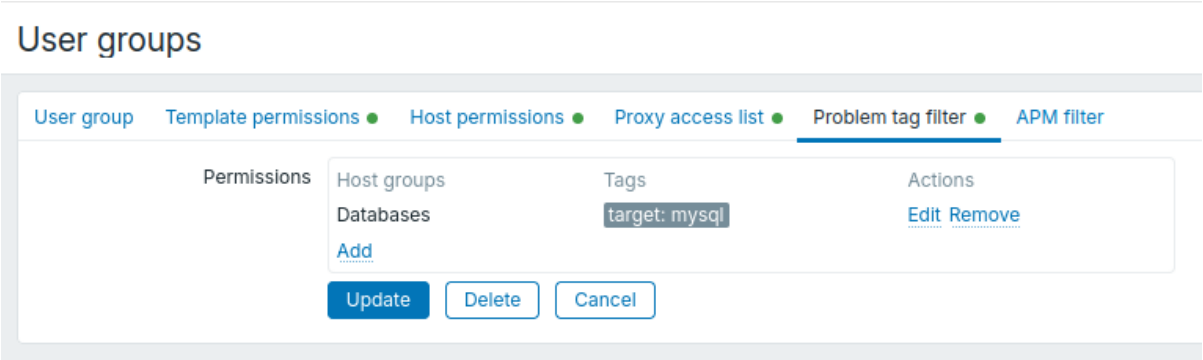
Select *Allow list* to allow the user group to access only the selected proxies and proxy groups.
 If the *Proxies* selection is empty for the *Allow list*, the user group has no access to any proxy that does not belong to a proxy group. If the *Proxy groups* selection is empty for the *Allow list*, the user group has no access to any proxy group.

Select *Deny list* to deny the user group access to the selected proxies and proxy groups. Access to all other proxies and proxy groups is allowed.
 If the *Proxies* selection is empty for the *Deny list*, the user group has full access to all proxies that does not belong to a proxy group. If the *Proxy groups* selection is empty for the *Deny list*, the user group has full access to all proxy groups.

Permission precedence rules for users who belong to multiple user groups:

- *Deny list* overrides *Allow list*. If any group denies access to a proxy/proxy group, the user cannot access it.
- *Allow list* permissions are cumulative. If multiple groups allow access, the user can access all proxies allowed by those groups.

The **Problem tag filter** tab allows setting tag-based permissions for user groups to see problems filtered by tag name and value:

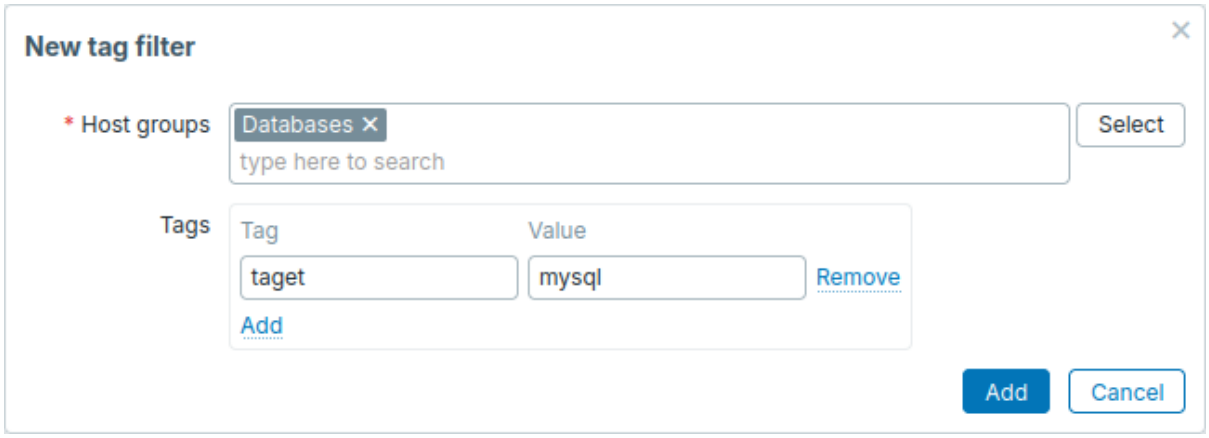


Click Add to choose the host groups. To select a host group to apply a tag filter for, click *Select* to get the complete list of existing host groups or start typing the name of a host group to get a dropdown of matching groups. Only host groups will be displayed, because problem tag filter cannot be applied to template groups.

Tag names without values can be added, but values without names cannot. Only the first three tags (with values, if any) are displayed in the *Permissions* block; if there are more, those can be seen by clicking or hovering over the ••• icon.

Tag filter allows separating the access to host group from the possibility to see problems.

For example, if a database administrator needs to see only "MySQL" database problems, it is required to create a user group for database administrators first, then specify "target" tag name and "mysql" value.



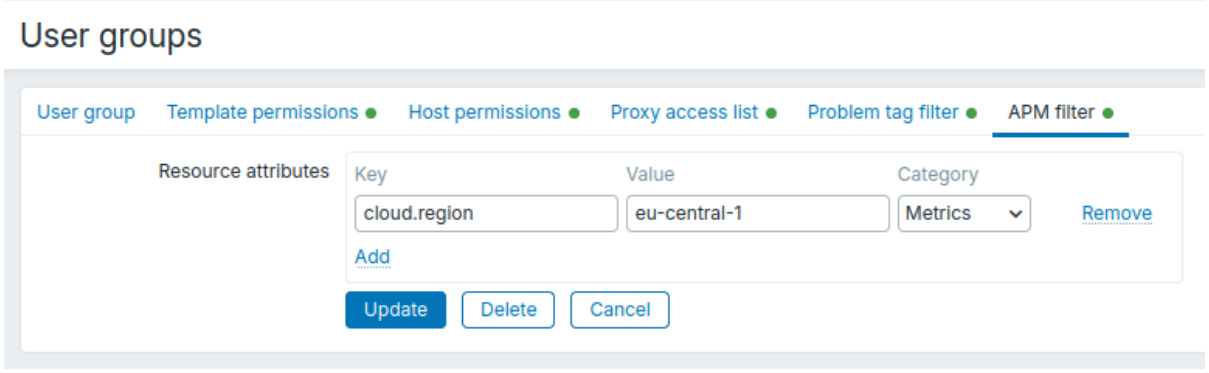
If "target" tag name is specified and value field is left blank, the user group will see all problems with tag name "target" for the selected host group.

Make sure tag name and tag value are correctly specified, otherwise, the user group will not see any problems.

Let's review an example when a user is a member of several user groups selected. Filtering in this case will use OR condition for tags.

User group A			**User group B**			**Visible result for a user (member) of both groups**
Tag filter						
Host group	*Tag name*	*Tag value*	*Host group*	*Tag name*	*Tag value*	
Databases	target	mysql	Databases	target	oracle	target:mysql or target:oracle problems visible
Not configured in the **Problem tag filter**			Databases	target	oracle	target:oracle problems visible

The **APM filter** tab allows you to specify user group access to APM telemetry data (traces, metrics, and logs).



Click Add to add resource attributes that define which incoming data the user group is allowed to access.

Specify the resource attribute:

Key	Enter the attribute key (e.g., `deployment.environment`, `cloud.region`).
Value	Enter the attribute value (e.g., `production`, `eu-central-1`).

Category	Select the category of incoming data to which the attribute applies: *Traces*, *Metrics*, *Logs*, or *All*.

If no filters are specified in the *APM filter* tab, members of the user group can access all incoming data.

If filters are specified, members of the user group can access only incoming data that contains at least one of the specified resource attributes for the corresponding data category or for *All*. An attribute matches a filter only if both the key and the value match.

Access is checked against the combined list of APM filters from all user groups to which the user belongs:

- Traces: a user can access a span if at least one of the span's resource attributes matches a filter for the *Traces* or *All* category.
- Metrics: a user can access a metric if at least one of the metric's resource attributes matches a filter for the *Metrics* or *All* category.
- Logs: a user can access a log record if at least one of the log record's resource attributes matches a filter for the *Logs* or *All* category.

Access from several user groups

A user may belong to any number of user groups. These groups may have different access permissions to hosts or templates.

Therefore, it is important to know what entities an unprivileged user will be able to access as a result. In the following example, consider how access to host **X** (in Hostgroup 1) will be affected in various situations for a user who is in user groups A and B.

- If Group A has only *Read* access to Hostgroup 1, but Group B *Read-write* access to Hostgroup 1, the user will get **Read-write** access to 'X'.

> **Attention:**
> "Read-write" permissions have precedence over "Read" permissions.

- In the same scenario as above, if 'X' is simultaneously also in Hostgroup 2 that is **denied** to Group A or B, access to 'X' will be **unavailable**, despite a *Read-write* access to Hostgroup 1.
- If Group A has no permissions defined and Group B has a *Read-write* access to Hostgroup 1, the user will get **Read-write** access to 'X'.
- If Group A has *Deny* access to Hostgroup 1 and Group B has a *Read-write* access to Hostgroup 1, the user will get access to 'X' **denied**.

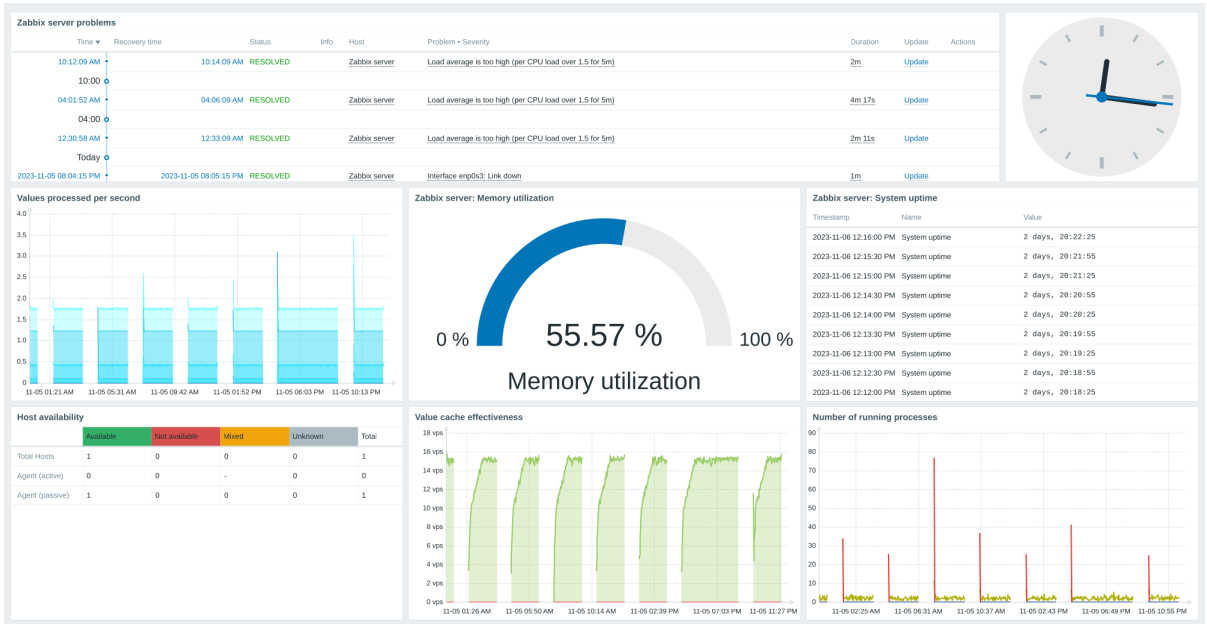
Other details

- An Admin level user with *Read-write* access to a host will not be able to link/unlink templates, if he has no access to the template group they belong to. With *Read* access to the template group he will be able to link/unlink templates to the host, however, will not see any templates in the template list and will not be able to operate with templates in other places.
- An Admin level user with *Read* access to a host will not see the host in the configuration section host list; however, the host triggers will be accessible in IT service configuration.
- Any non-Super Admin user (including 'guest') can see network maps as long as the map is empty or has only images. When hosts, host groups or triggers are added to the map, permissions are respected.
- Zabbix server will not send notifications to users defined as action operation recipients if access to the concerned host is explicitly "denied".

Scheduled reports

Overview

With the *Scheduled reports* feature, you can set up a PDF version of a given dashboard to be sent to specified recipients at recurring intervals.



Pre-requisites:

- Zabbix web service must be installed and configured correctly to enable scheduled report generation - see Setting up scheduled reports for instructions.
- A user must have a user role of type *Admin* or *Super admin* with the following permissions:
  - *Scheduled reports* in the *Access to UI elements* block (to view report settings)
  - *Manage scheduled reports* in the *Access to actions* block (to create/edit reports)

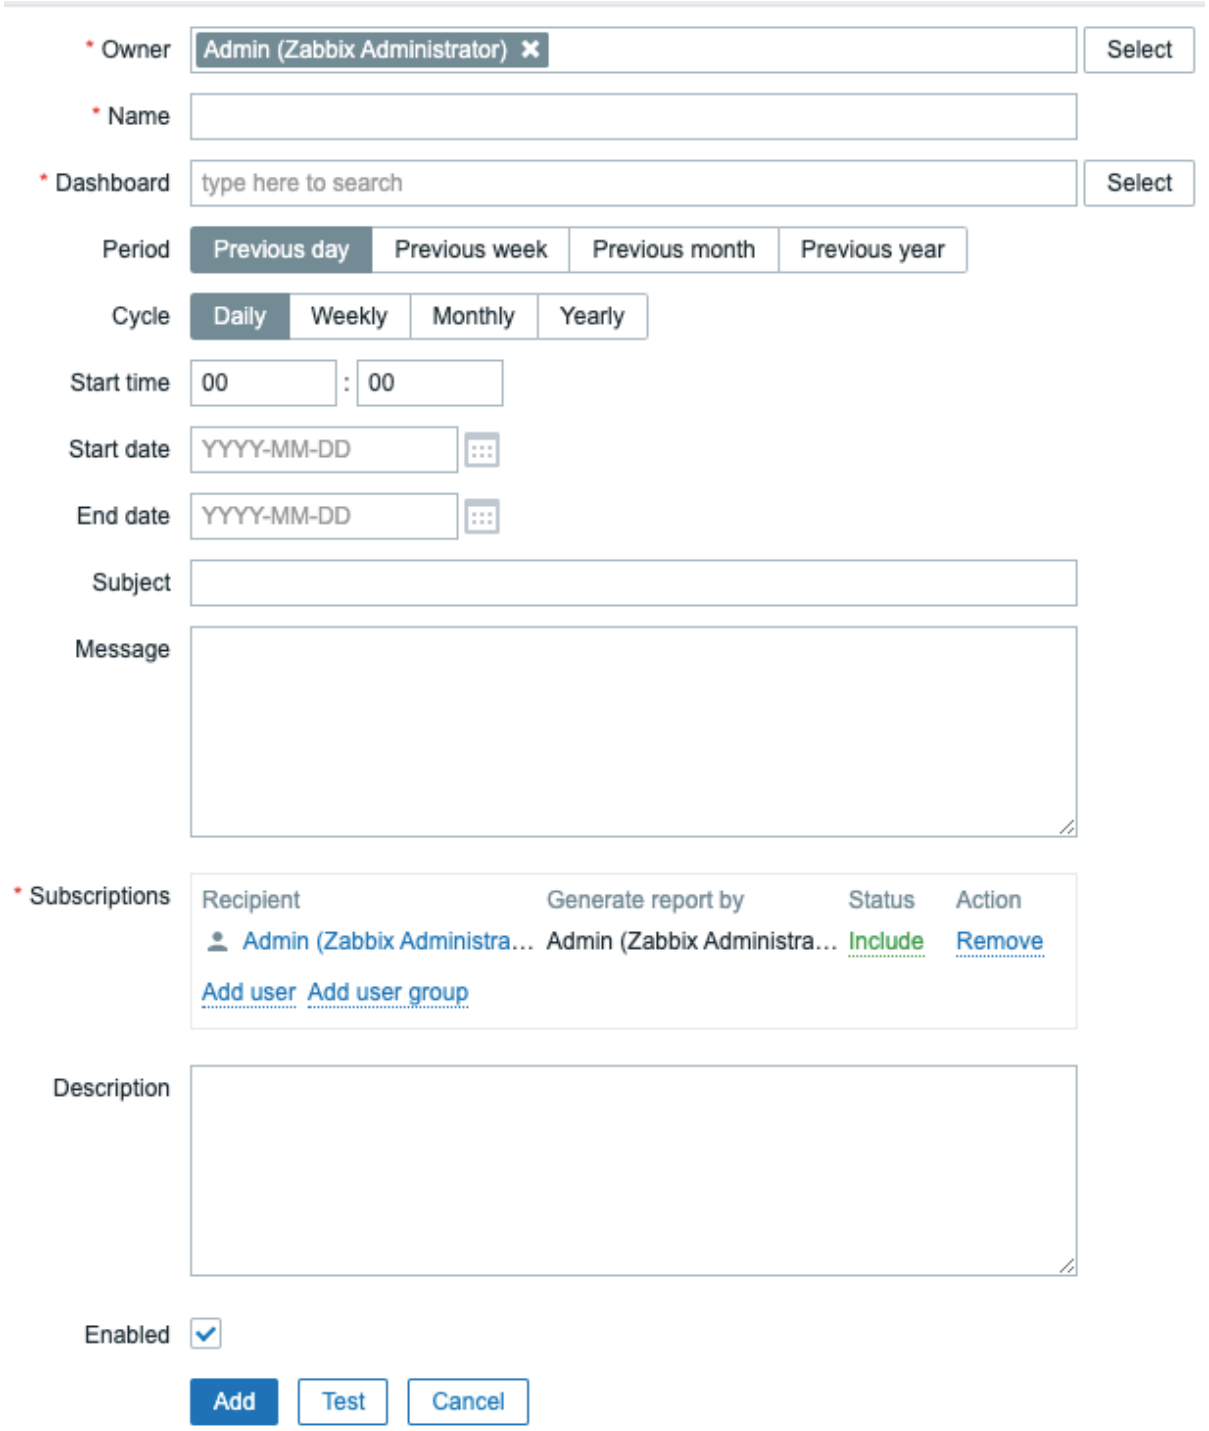
To create a scheduled report in Zabbix frontend, do the following:

1. Go to *Reports* > *Scheduled reports*.
2. Click *Create report* in the upper-right corner of the screen.
3. Enter parameters of the report in the form.

Alternatively, you can create a report by opening an existing one, clicking the *Clone* button, and saving it under a different name.

Configuration

The *Scheduled reports* tab contains general report attributes.



All mandatory input fields are marked with a red asterisk.

Parameter	Description
Owner	User that creates a report. *Super admin* level users are allowed to change the owner. For *Admin* level users, this field is read-only.
Name	Name of the report; must be unique.
Dashboard	Dashboard on which the report is based; only one dashboard can be selected at a time. To select a dashboard, start typing the name - a list of matching dashboards will appear; scroll down to select. Alternatively, you may click *Select* next to the field and select a dashboard from the displayed list.
Period	Time period for which the report will be prepared. Select the previous day, week, month, or year.
Cycle	Report generation frequency. The reports can be sent daily, weekly, monthly, or yearly. ”Weekly” mode allows you to select days of the week when the report will be sent.
Start time	Time of the day in the format hh:mm when the report will be prepared. Note that Zabbix server time zone will be used.
Repeat on	Days of the week when the report will be sent. This field is available only if *Cycle* is set to ”Weekly”.

Parameter	Description
Start date	Date when regular report generation should be started.
End date	Date when regular report generation should be stopped.
Subject	Subject of the report email. Supported macros: {TIME}, {TIMESTAMP}.
Message	Body of the report email. Supported macros: {TIME}, {TIMESTAMP}.
Subscriptions	List of report recipients. By default, includes only the report owner. Any Zabbix user with configured email media may be specified as a report recipient. Click *Add user* or *Add user group* to add more recipients. Click the username to edit settings: *Generate report by* - whether the report data should be generated based on the dashboard permissions of the current user or the recipient. *Status* - select ”Include” to send the report to the user or ”Exclude” to prevent sending the report to this user. At least one user must have the ”Include” status. The ”Exclude” status can be used to exclude specific users from a user group that is included. Note that users with insufficient permissions (that is, users with a role based on the *Admin* user type who are not members of the same user group as the recipient or report owner) will see ”Inaccessible user” or ”Inaccessible user group” instead of the actual names in the fields *Recipient* and *Generate report by*; the fields *Status* and *Action* will be displayed as read-only.
Enabled	Report status. Clearing this checkbox will disable the report.
Description	An optional description of the report. This description is for internal use and will not be sent to report recipients.

Form buttons

Buttons at the bottom of the form allow you to perform several operations.

Add	Add a report. This button is only available for new reports.
Update	Update the properties of a report.
Clone	Create another report based on the properties of the current report.
Test	Test if report configuration is correct by sending a report to the current user.
Delete	Delete the report.
Cancel	Cancel the editing of report properties.

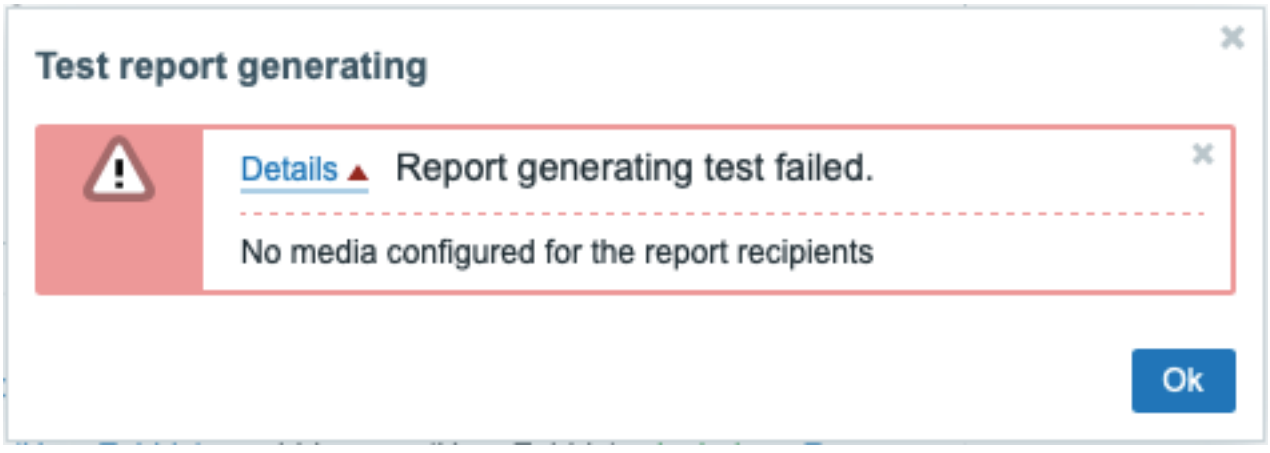
Testing

To test a report, click the *Test* button at the bottom of the report configuration form.

Note:
The *Test* button is not available if the report configuration form has been opened from the dashboard action menu.

If the configuration is correct, the test report is sent immediately to the current user. For test reports, subscribers and *Generate report by* user settings are ignored.

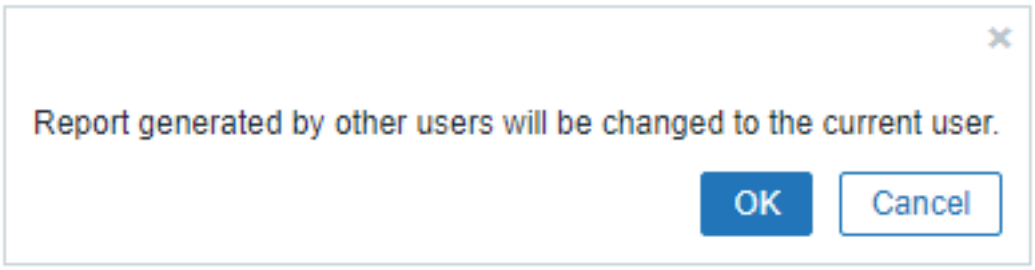
If the configuration is incorrect, an error message is displayed describing the possible cause.



Updating a report

To update an existing report, click the report name, make the required configuration changes, and then click the *Update* button.

If an existing report is updated by another user and this user changes the Dashboard, upon clicking the *Update* button, a warning message ”Report generated by other users will be changed to the current user” will be displayed.



Clicking *OK* at this step will lead to the following changes:

- *Generate report by* settings will be updated to display the user who edited the report last (unless *Generate report by* is set to the recipient).
- Users that have been displayed as ”Inaccessible user” or ”Inaccessible user group” will be deleted from the list of report subscribers.

Clicking *Cancel* will close the configuration form and cancel the report update.

Cloning a report

To quickly clone an existing report, click the *Clone* button at the bottom of an existing report configuration form. When cloning a report created by another user, the current user becomes the owner of the new report.

Report settings will be copied to the new report configuration form with respect to user permissions:

- If the user who clones a report has no permissions to a dashboard, the *Dashboard* field will be cleared.
- If the user who clones a report has no permissions to some users or user groups in the *Subscriptions* list, inaccessible recipients will not be cloned.
- *Generate report by* settings will be updated to display the current user (unless *Generate report by* is set to the recipient).

Change the required settings and the report name, then click *Add*.

Data export

Overview

Zabbix supports data export in real-time in two ways:

- export to files
- streaming to external systems

The following entities can be exported:

- trigger events
- item values
- trends (export to files only)

Export to files

Overview

It is possible to configure real-time exporting of trigger events, item values and trends in a newline-delimited JSON format.

Exporting is done into files, where each line of the export file is a JSON object. Value mappings are not applied.

In case of errors (data cannot be written to the export file or the export file cannot be renamed or a new one cannot be created after renaming it), the data item is dropped and never written to the export file. It is written only in the Zabbix database. Writing data to the export file is resumed when the writing problem is resolved.

The export file is created with read-write permissions for the file owner only. Additionally the file is readable by the owner group. All other permissions are denied.

Note that host/item can have no metadata (host groups, host name, item name) if the host/item was removed after the data was received, but before server exported data.

Configuration

Real-time export of trigger events, item values and trends is configured by specifying a directory for the export files - see the ExportDir parameter in server configuration.

Two other parameters are available:

- `ExportFileSize` may be used to set the maximum allowed size of an individual export file. When a process needs to write to a file it checks the size of the file first. If it exceeds the configured size limit, the file is renamed by appending .old to its name and a new file with the original name is created.

Attention:
A file will be created per each process that will write data (i.e. approximately 4-30 files). As the default size per export file is 1G, keeping large export files may drain the disk space fast.

- `ExportType` allows you to specify which entity types (events, history, trends) will be exported.

Newline-delimited JSON export protocol

This section presents details of the export protocol in a newline-delimited JSON format, used in:

- data export to files
- streaming to external systems

All files have an `.ndjson` extension. Each line of the export file is a JSON object.

Trigger event export

The following information is exported for a problem event:

Field		Type	Description
clock		number	Number of seconds since Epoch to the moment when problem was detected (integer part).
ns		number	Number of nanoseconds to be added to `clock` to get a precise problem detection time.
value		number	1 (always).
eventid		number	Problem event ID.
name		string	Problem event name.
severity		number	Problem event severity (0 - Not classified, 1 - Information, 2 - Warning, 3 - Average, 4 - High, 5 - Disaster).
hosts		array	List of hosts involved in the trigger expression; there should be at least one element in array.
-		object	
	host	string	Host name.
	name	string	Visible host name.
groups		array	List of host groups of all hosts involved in the trigger expression; there should be at least one element in array.
-		string	Host group name.
tags		array	List of problem tags (can be empty).
-		object	
	tag	string	Tag name.
	value	string	Tag value (can be empty).

The following information is exported for a recovery event:

Field	Type	Description
clock	number	Number of seconds since Epoch to the moment when problem was resolved (integer part).
ns	number	Number of nanoseconds to be added to `clock` to get a precise problem resolution time.
value	number	0 (always).
eventid	number	Recovery event ID.
p_eventid	number	Problem event ID.

Trigger event export examples

Problem:

```
{"clock":1519304285,"ns":123456789,"value":1,"name":"Either Zabbix agent is unreachable on Host B or polle
```

Recovery:

```
{"clock":1519304345,"ns":987654321,"value":0,"eventid":43,"p_eventid":42}
```

Problem (multiple problem event generation):

```
{"clock":1519304286,"ns":123456789,"value":1,"eventid":43,"name":"Either Zabbix agent is unreachable on Ho
```

```
{"clock":1519304286,"ns":123456789,"value":1,"eventid":43,"name":"Either Zabbix agent is unreachable on Ho
```

Recovery:

```
{"clock":1519304346,"ns":987654321,"value":0,"eventid":44,"p_eventid":43}
```

```
{"clock":1519304346,"ns":987654321,"value":0,"eventid":44,"p_eventid":42}
```

Item value export

The following information is exported for a collected item value:

Field	Type	Description
host	object	Host name of the item host.
host	string	Host name.
name	string	Visible host name.
groups	array	List of host groups of the item host; there should be at least one element in array.
-	string	Host group name.
item_tags	array	List of item tags (can be empty).
-	object	
tag	string	Tag name.
value	string	Tag value (can be empty).
itemid	number	Item ID.
name	string	Visible item name.
clock	number	Number of seconds since Epoch to the moment when value was collected (integer part).
ns	number	Number of nanoseconds to be added to `clock` to get a precise value collection time.
timestamp (*Log* only)	number	0 if not available.
source (*Log* only)	string	Empty string if not available.
severity (*Log* only)	number	0 if not available.
eventid (*Log* only)	number	0 if not available.
value	number (for numeric items) or string (for text items)	Collected item value.
type	number	Collected value type: 0 - numeric float, 1 - character, 2 - log, 3 - numeric unsigned, 4 - text, 5 - binary, 6 - JSON

Item value export example

Numeric (unsigned) value:

```
{"host":{"host":"Host B","name":"Host B visible"},"groups":["Group X","Group Y","Group Z"],"item_tags":[{"
```

Numeric (float) value:

{"host":{"host":"Host B","name":"Host B visible"},"groups":["Group X","Group Y","Group Z"],"item_tags":[{"

Character, text value:

{"host":{"host":"Host B","name":"Host B visible"},"groups":["Group X","Group Y","Group Z"],"item_tags":[{"

Log value:

{"host":{"host":"Host A","name":"Host A visible"},"groups":["Group X","Group Y","Group Z"],"item_tags":[{"

Trends export

The following information is exported for a calculated trend value:

Field	Type	Description
host	object	Host name of the item host.
host	string	Host name.
name	string	Visible host name.
groups	array	List of host groups of the item host; there should be at least one element in array.
-	string	Host group name.
item_tags	array	List of item tags (can be empty).
-	object	
tag	string	Tag name.
value	string	Tag value (can be empty).
itemid	number	Item ID.
name	string	Visible item name.
clock	number	Number of seconds since Epoch to the moment when value was collected (integer part).
count	number	Number of values collected for a given hour.
min	number	Minimum item value for a given hour.
avg	number	Average item value for a given hour.
max	number	Maximum item value for a given hour.
type	number	Value type: 0 - numeric float, 3 - numeric unsigned

Trends export examples

Numeric (unsigned) value:

{"host":{"host":"Host B","name":"Host B visible"},"groups":["Group X","Group Y","Group Z"],"item_tags":[{"

Numeric (float) value:

{"host":{"host":"Host B","name":"Host B visible"},"groups":["Group X","Group Y","Group Z"],"item_tags":[{"

Streaming to external systems

Overview

It is possible to stream item values and events from Zabbix to external systems over HTTP (see protocol details).

The tag filter can be used to stream subsets of item values or events.

Two Zabbix server process types are responsible for data streaming: `connector manager` and `connector worker`. A Zabbix internal item `zabbix[connector_queue]` allows you to monitor the count of values enqueued in the connector queue.

Configuration

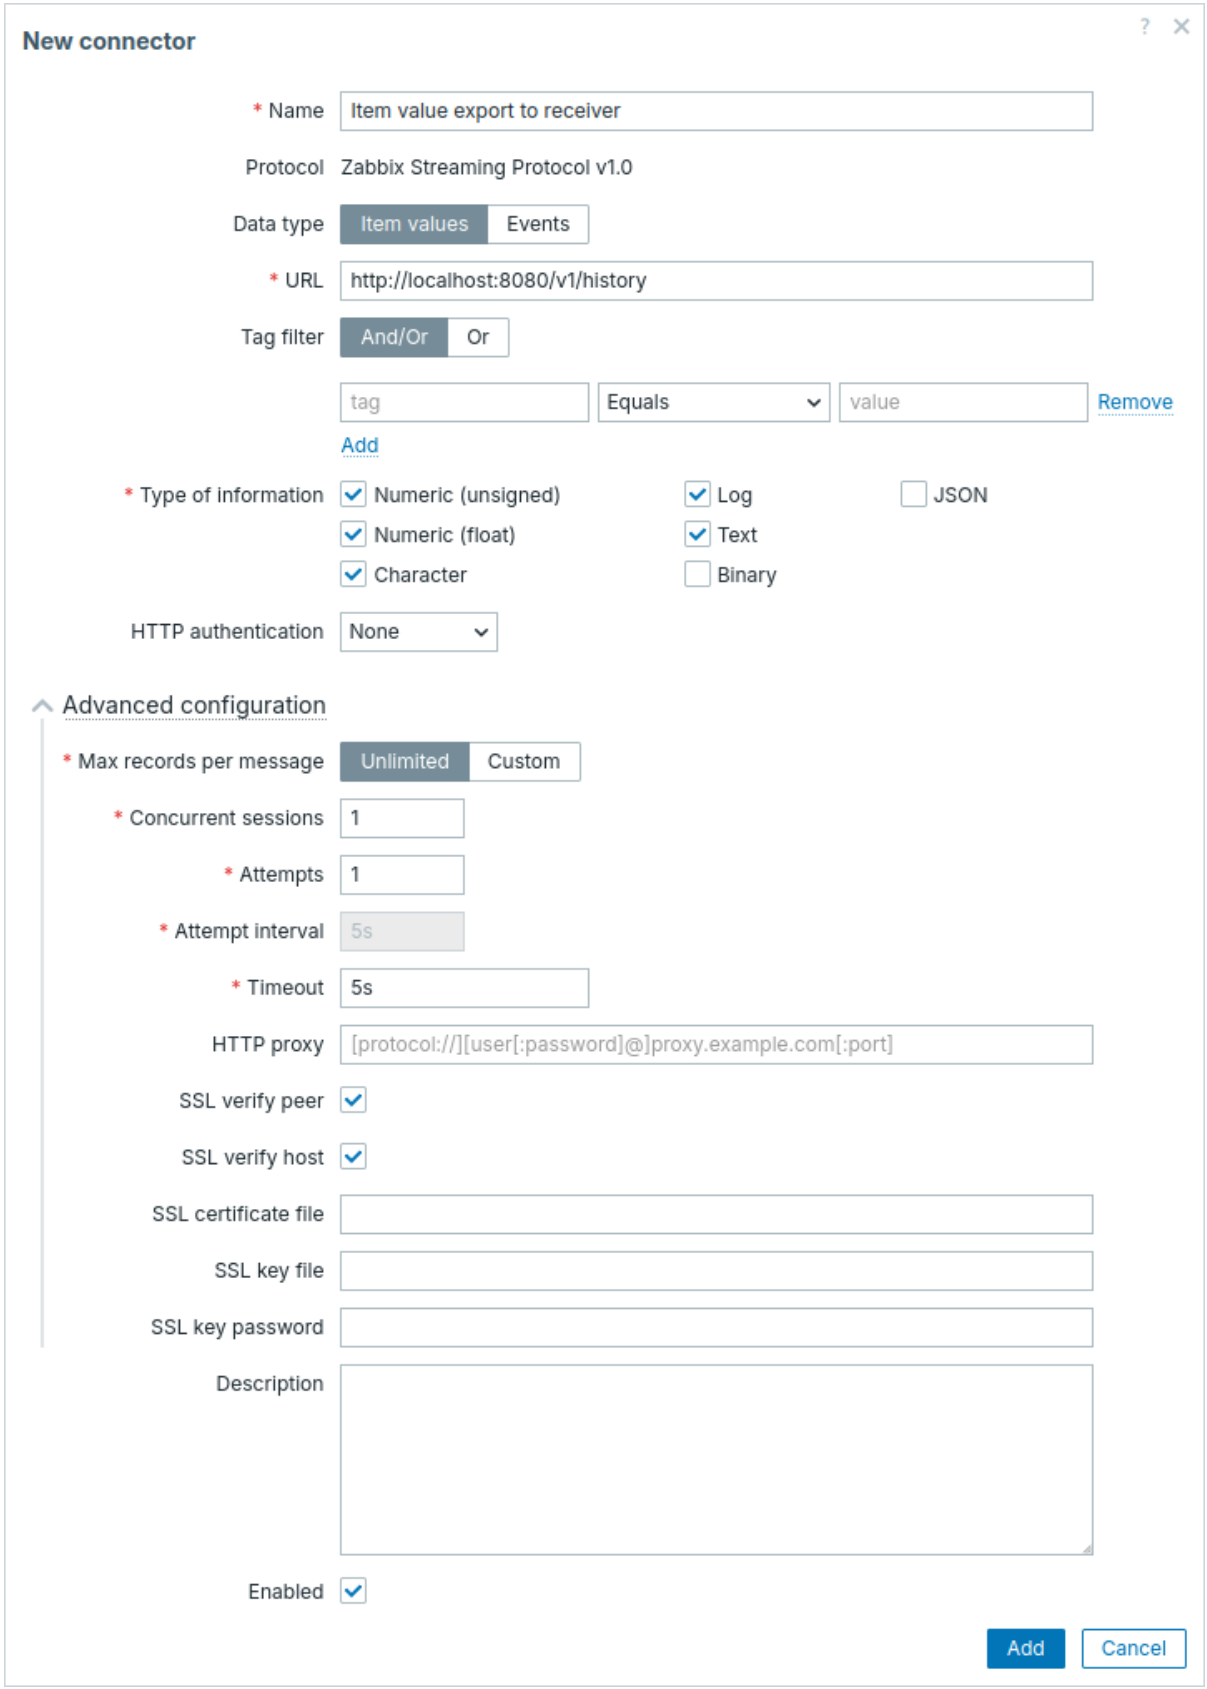
The following steps are required to configure data streaming to an external system:

1. Have a remote system set up for receiving data from Zabbix. For this purpose, the following tools are available:

- An example of a simple receiver that logs the received information in `events.ndjson` and `history.ndjson` files.
- Kafka connector for Zabbix server - a lightweight server written in Go, designed to forward item values and events from a Zabbix server to a Kafka broker.

2. Set the required number of connector workers in Zabbix by adjusting the `StartConnectors` parameter in `zabbix_server.conf`. The number of connector workers should match (or exceed if concurrent sessions are more than 1) the configured connector count in Zabbix frontend. Then, restart Zabbix server.

3. Configure a new connector in Zabbix frontend (*Administration* > *General* > *Connectors*) and reload the server cache with the `zabbix_server -R config_cache_reload` command.



Mandatory fields are marked by an asterisk.

Parameter	Description
Name	Enter the connector name.
Data type	Select the data type to stream: **Item values** - stream item values from Zabbix to external systems; **Events** - stream events from Zabbix to external systems.

Parameter	Description
URL	Enter the receiver URL. User macros are supported.
Tag filter	Export only item values or events matching the tag filter. If not set, then export everything. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met.
Type of information	Select the type of information (numeric (unsigned), numeric (float), character, etc.), by which to filter the item values that the connector should stream. This field is available if *Data type* is set to ”Item values”.
HTTP authentication	Select the authentication option: **None** - no authentication used; **Basic** - basic authentication is used; **NTLM** - NTLM (Windows NT LAN Manager) authentication is used; **Kerberos** - Kerberos authentication is used (see also: Configuring Kerberos with Zabbix); **Digest** - Digest authentication is used; **Bearer** - Bearer authentication is used.
Username	Enter the user name (up to 255 characters). User macros are supported. This field is available if *HTTP authentication* is set to ”Basic”, ”NTLM”, ”Kerberos”, or ”Digest”.
Password	Enter the user password (up to 255 characters). User macros are supported. This field is available if *HTTP authentication* is set to ”Basic”, ”NTLM”, ”Kerberos”, or ”Digest”.
Bearer token	Enter the Bearer token. User macros are supported. This field is available and required if *HTTP authentication* is set to ”Bearer”.
Advanced configuration	Click the *Advanced configuration* header to display advanced configuration options (see below).
Max records per message	Specify the maximum number of values or events that can be streamed within one message.
Concurrent sessions	Select the number of sender processes to run for this connector. Up to 100 sessions can be specified; the default value is ”1”.
Attempts	Number of attempts for streaming data. Up to 5 attempts can be specified; the default value is ”1”.
Attempt interval	Specify how long the connector should wait after an unsuccessful attempt to stream data. Up to 10s can be specified; the default value is ”5s”. This field is available if *Attempts* is set to ”2” or more. Unsuccessful attempts are those where establishing a connection has failed, or where the HTTP response code is not 200, 201, 202, 203, 204. Retries are **triggered** in case of communication errors or when the HTTP response code is not 200, 201, 202, 203, 204, 400, 401, 403, 404, 405, 415, 422. Redirects are followed, so 302 -> 200 is a positive response; whereas 302 -> 503 will trigger a retry.
Timeout	Specify the message timeout (1-60 seconds, default - 5 seconds). Time suffixes are supported (e.g., 30s, 1m). User macros are supported.

Parameter	Description
HTTP proxy	You can specify an HTTP proxy to use in the following format: `[protocol://][username[:password]@]proxy.example.com[:port]` User macros are supported. The optional `protocol://` prefix may be used to specify alternative proxy protocols (the protocol prefix support was added in cURL 7.21.7). With no protocol specified, the proxy will be treated as an HTTP proxy. By default, 1080 port will be used. If *HTTP proxy* is specified, the proxy will overwrite proxy related environment variables like `http_proxy`, `HTTPS_PROXY`. If not specified, the proxy will not overwrite proxy-related environment variables. The entered value is passed on as is, no sanity checking takes place. You may also enter a SOCKS proxy address. If you specify the wrong protocol, the connector will fail to stream item values or events from Zabbix. Note that only simple authentication is supported with HTTP proxy.
SSL verify peer	Mark the checkbox to verify the SSL certificate of the web server. The server certificate is verified against the trusted certificate authority (CA) certificates from the system-wide CA location. You can override the location of CA files using Zabbix server or proxy configuration parameter `SSLCALocation`.
SSL verify host	Mark the checkbox to verify that the *Common Name* field or the *Subject Alternate Name* field of the web server certificate matches. This sets the `CURLOPT_SSL_VERIFYHOST` cURL option.
SSL certificate file	Name of the SSL certificate file used for client authentication. The certificate file must be in PEM[1] format. User macros are supported. If the certificate file also contains the private key, leave the *SSL key file* field empty. If the key is encrypted, specify the password in the *SSL key password* field. The directory containing this file is specified by Zabbix server or proxy configuration parameter `SSLCertLocation`.
SSL key file	Name of the SSL private key file used for client authentication. The private key file must be in PEM[1] format. User macros are supported. The directory containing this file is specified by Zabbix server or proxy configuration parameter `SSLKeyLocation`.
SSL key password	SSL private key file password. User macros are supported.
Description	Enter the connector description.
Enabled	Mark the checkbox to enable the connector.

When the Kafka connector is configured with a comma-separated list of bootstrap broker addresses (for example, `Kafka.URL=kafka1.exampl` the Kafka client connects to the broker(s) that respond first and uses their cluster metadata. If the list contains addresses from different Kafka clusters, only the fastest-responding cluster will be used and other addresses will be logged as unavailable; as a result, startup warnings such as the following may appear even though the connector is connected:

```
kafka cluster connected, but broker(s) "kafka1.example.com:9093, kafka2.example.com:9093" unavailable; wil
```

In some environments (private networks, container networks or non-standard DNS/hosts setups) hostnames or IPs may resolve to loopback addresses (for example, `127.0.0.1/localhost`) or be normalized by the client, which can make such warnings misleading. To reduce confusion, ensure all `Kafka.URL` addresses belong to the same Kafka cluster, verify DNS resolution from the connector host and the brokers' `advertised.listeners`, and prefer addresses that resolve to the broker's advertised address.

Protocol

Communication between the server and the receiver is done over HTTP using REST API, NDJSON, "Content-Type: application/x-ndjson".

For more details, see Newline-delimited JSON export protocol.

Server request

Example of streaming item values:

```
POST /v1/history HTTP/1.1
Host: localhost:8080
Accept: */*
Accept-Encoding: deflate, gzip, br, zstd
Content-Length: 628
Content-Type: application/x-ndjson
```

```
{"host":{"host":"Zabbix server","name":"Zabbix server"},"groups":["Zabbix servers"],"item_tags":[{"tag":"f
{"host":{"host":"Zabbix server","name":"Zabbix server"},"groups":["Zabbix servers"],"item_tags":[{"tag":"f
{"host":{"host":"Zabbix server","name":"Zabbix server"},"groups":["Zabbix servers"],"item_tags":[{"tag":"b
```

Example of streaming events:

```
POST /v1/events HTTP/1.1
Host: localhost:8080
Accept: */*
Accept-Encoding: deflate, gzip, br, zstd
Content-Length: 333
Content-Type: application/x-ndjson

{"clock":1673454303,"ns":800155804,"value":1,"eventid":5,"name":"trigger for foo being 0","severity":0,"ho
{"clock":1673454303,"ns":832290669,"value":0,"eventid":6,"p_eventid":5}
```

Receiver response

The response consists of the HTTP response status code and the JSON payload. The HTTP response status code must be "200", "201", "202", "203", or "204" for requests that were handled successfully, other for failed requests.

Example of success:

```
HTTP/1.1 200 OK
Content-Type: application/json
X-Content-Type-Options: nosniff
Date: Tue, 21 Apr 2026 10:13:04 GMT
Content-Length: 23

{"response":"success"}
```

Example with errors:

```
HTTP/1.1 422 Unprocessable Entity
Content-Type: application/json
X-Content-Type-Options: nosniff
Date: Tue, 21 Apr 2026 12:15:01 GMT
Content-Length: 55

{"error":"invalid character '{' after top-level value"}
```

SNMP gateway

Overview

Zabbix SNMP gateway is an AgentX-extension for snmpd supporting both SNMP polling and trapping.

With Zabbix SNMP gateway it is possible to use the SNMP protocol to retrieve:

- trigger data;
- problem trigger data;
- host group status (count of triggers by trigger state per group)

The data is retrieved by the OID, which is a combination of a common base and a specific suffix. The common **base** is set in the configuration file of SNMP gateway, for example:

- `BaseOID=1.3.6.1.4.1.3043.7.55` - for any trigger data;
- `ProblemBaseOID=1.3.6.1.4.1.3047.7.55` - for problem trigger data;
- `BaseOID=1.3.6.1.4.1.3046.7.55` - for host group status.

The OID **suffix** is set in the configuration on host triggers as a **tag** (for example, `OIDSuffix:3`) in the frontend.

In this case all information for the trigger will be available under `OID=1.3.6.1.4.1.3043.7.55.X.3`. "X" here will be the number of trigger data fields (i.e. 1 - suffix, 2 - ID, 3 - expression, 4 - description, etc.).

For a more detailed description and the configuration file example, see the SNMP gateway readme file.

Installation and setup

See the readme file of SNMP gateway for instructions on:

- installing and configuring snmpd;
- enabling AgentX support;
- configuring Zabbix SNMP gateway;
- configuring SNMP traps for trigger state changes.

Retrieving data

With everything set up properly, you may use `snmpwalk` and `snmpget` commands to retrieve data:

```
[user@localhost ~]# snmpget -v2c -c public 127.0.0.1 1.3.6.1.4.1.3043.7.55.2.3
SNMPv2-SMI::enterprises.3043.7.55.2.3 = INTEGER: 15247
```

```
[user@localhost ~]# snmpwalk -v2c -c public 127.0.0.1 1.3.6.1.4.1.3043.7.55
SNMPv2-SMI::enterprises.3043.7.55.1.1 = INTEGER: 1
SNMPv2-SMI::enterprises.3043.7.55.1.3 = INTEGER: 3
SNMPv2-SMI::enterprises.3043.7.55.1.4 = INTEGER: 4
SNMPv2-SMI::enterprises.3043.7.55.1.5 = INTEGER: 5
SNMPv2-SMI::enterprises.3043.7.55.1.6 = INTEGER: 6
SNMPv2-SMI::enterprises.3043.7.55.1.10 = INTEGER: 10
SNMPv2-SMI::enterprises.3043.7.55.2.1 = INTEGER: 15367
SNMPv2-SMI::enterprises.3043.7.55.2.3 = INTEGER: 15247
SNMPv2-SMI::enterprises.3043.7.55.2.4 = INTEGER: 15365
SNMPv2-SMI::enterprises.3043.7.55.2.5 = INTEGER: 15366
SNMPv2-SMI::enterprises.3043.7.55.2.6 = INTEGER: 13493
SNMPv2-SMI::enterprises.3043.7.55.2.10 = INTEGER: 13503
...
```

Filtering options

You may limit the problem trigger information in SNMP gateway configuration:

- by severity (by default `ProblemMinSeverity=-1`)
- by hiding acknowledged problems (by default `ProblemHideAck=false`)

You may limit the problem count per host group in SNMP gateway configuration:

- by unknown state triggers (by default `CountUnknown=false`)
- by triggers with acknowledged/unacknowledged/all problems (by default `CountAcknowledgeStatus=all`)

Configuration export/import

Overview

Zabbix export/import functionality makes it possible to exchange various configuration entities between one Zabbix system and another.

Typical use cases for this functionality:

- share templates or network maps - Zabbix users may share their configuration parameters
- upload a template to Zabbix Community templates. Then others can download the template and import the file into Zabbix.
- integrate with third-party tools - universal YAML, XML and JSON formats make integration and data import/export possible with third-party tools and applications

What can be exported/imported

Objects that can be exported/imported are:

- Host groups (*through Zabbix API only*)
- Template groups (*through Zabbix API only*)
- Templates
- Hosts
- Network maps
- Media types
- Images
- Dashboards
- Regular expressions

Export format

Data can be exported using the Zabbix web frontend or Zabbix API. Supported export formats are YAML, XML and JSON.

Details about export

- All supported elements are exported in one file.
- Host and template entities (items, triggers, graphs, discovery rules) that are inherited from linked templates are not exported. Any changes made to those entities on a host level (such as changed item interval, modified regular expression or added prototypes to the low-level discovery rule) will be lost when exporting; when importing, all entities from linked templates are re-created as on the original linked template.
- Entities created by low-level discovery and any entities depending on them are not exported. For example, a trigger created for an LLD-rule generated item will not be exported.
- When the exported host/template contains entities supporting timeouts, the timeout values will be exported if these entities have their own timeouts configured.

Details about import

- Import stops at the first error.
- When updating existing images during image import, "imagetype" field is ignored, i.e., it is impossible to change image type via import.
- When importing hosts/templates using the "Delete missing" option, host/template macros not present in the import file will be deleted from
the host/template after the import.
- Empty tags for items, triggers, graphs, discoveryRules, itemPrototypes, triggerPrototypes, graphPrototypes are meaningless, i.e., it's the same as if it was missing.
- If entities of the imported host/template have their own timeouts configured, they will be applied; otherwise, proxy/global timeouts will be applied.
- Import supports YAML, XML and JSON, the import file must have a correct file extension: .yaml and .yml for YAML, .xml for XML and .json for JSON. See compatibility information about supported XML versions.
- Import supports configuration files only in UTF-8 encoding (with or without BOM); other encodings (UTF16LE, UTF16BE, UTF32LE, UTF32BE, etc.) will result in an import conversion error.

YAML base format

The YAML export format contains the following nodes:

- Root node for Zabbix YAML export
- Export version

```
zabbix_export:
  version: '8.2'
```

Other nodes are dependent on exported objects.

XML format

The XML export format contains the following tags:

- Default header for XML documents
- Root tag for Zabbix XML export
- Export version

```
<?xml version="1.0" encoding="UTF-8"?>
<zabbix_export>
    <version>8.2</version>
</zabbix_export>
```

Other tags are dependent on exported objects.

JSON format

The JSON export format contains the following objects:

- Root object for Zabbix JSON export
- Export version

```
{
    "zabbix_export": {
        "version": "8.2"
    }
}
```

Other objects are dependent on exported objects.

Template groups

Overview

In the frontend, template groups can be exported only with template export. When a template is exported, all groups it belongs to are exported with it automatically.

API allows exporting template groups independently of templates.

Export format

```
template_groups:
  - uuid: 36bff6c29af64692839d077febfc7079
    name: 'Network devices'
```

Exported elements

Element	Type	Description
uuid	string	Unique identifier for this template group.
name	string	Group name.

Host groups

Overview

In the frontend, host groups can be exported only with host export. When a host is exported, all groups it belongs to are exported with it automatically.

API allows exporting host groups independently of hosts.

Export format

```
host_groups:
  - uuid: 6f6799aa69e844b4b3918f779f2abf08
    name: 'Zabbix servers'
```

Exported elements

Element	Type	Description
uuid	string	Unique identifier for this host group.
name	string	Group name.

Templates

Overview

Templates are exported with many related objects and object relations.

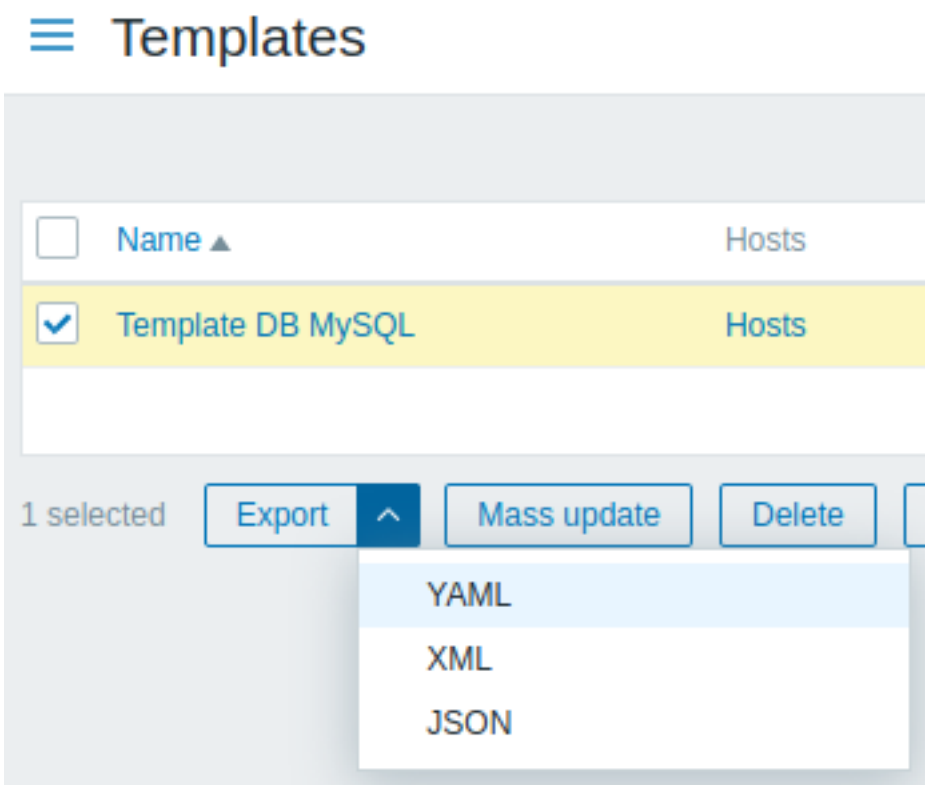
Template export contains:

- Linked template groups
- Linked host groups (if used in host prototype configuration)
- Template data
- Linkage to other templates
- Linkage to template groups
- Directly linked items
- Directly linked triggers
- Directly linked graphs
- Directly linked dashboards
- Directly linked discovery rules with all prototypes
- Directly linked web scenarios
- Value maps

Exporting

To export templates, do the following:

1. Go to *Data collection* > *Templates*.
2. Mark the checkboxes of the templates to export.
3. Click *Export* below the list to export in the default YAML format, or click the arrow to select a different export format.



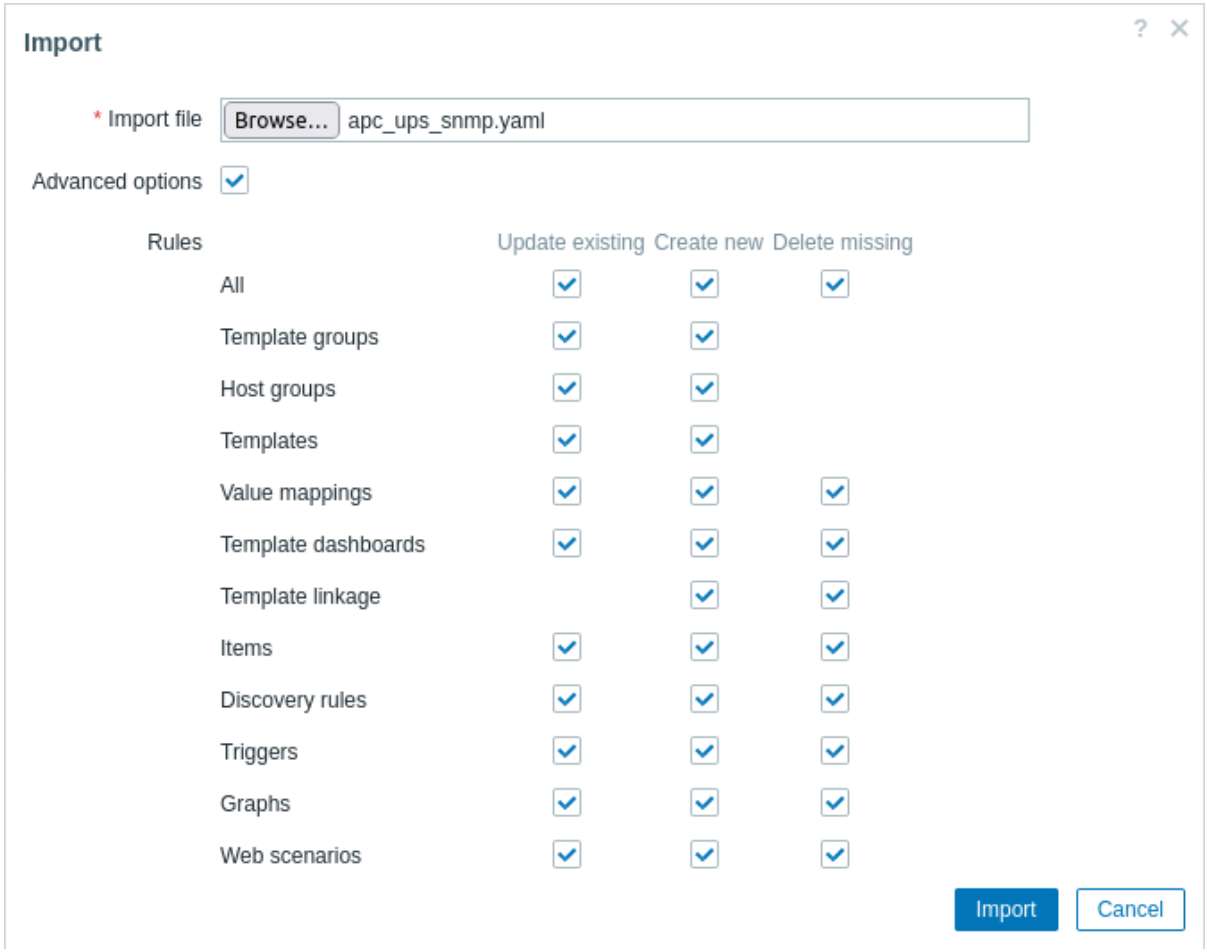
Depending on the selected format, templates are exported to a local file with a default name:

- `zabbix_export_templates.yaml` - in YAML export;
- `zabbix_export_templates.xml` - in XML export;
- `zabbix_export_templates.json` - in JSON export.

Importing

To import templates, do the following:

1. Go to *Data collection* > *Templates*.
2. Click *Import* in the upper-right corner.
3. Select the import file.
4. Mark the checkboxes of the import rules you want to apply (see below).
5. Click *Import* in the lower-right corner of the configuration form.



If you mark the *Advanced options* checkbox, a detailed list of all importable elements will be displayed - mark or unmark each import rule as required.

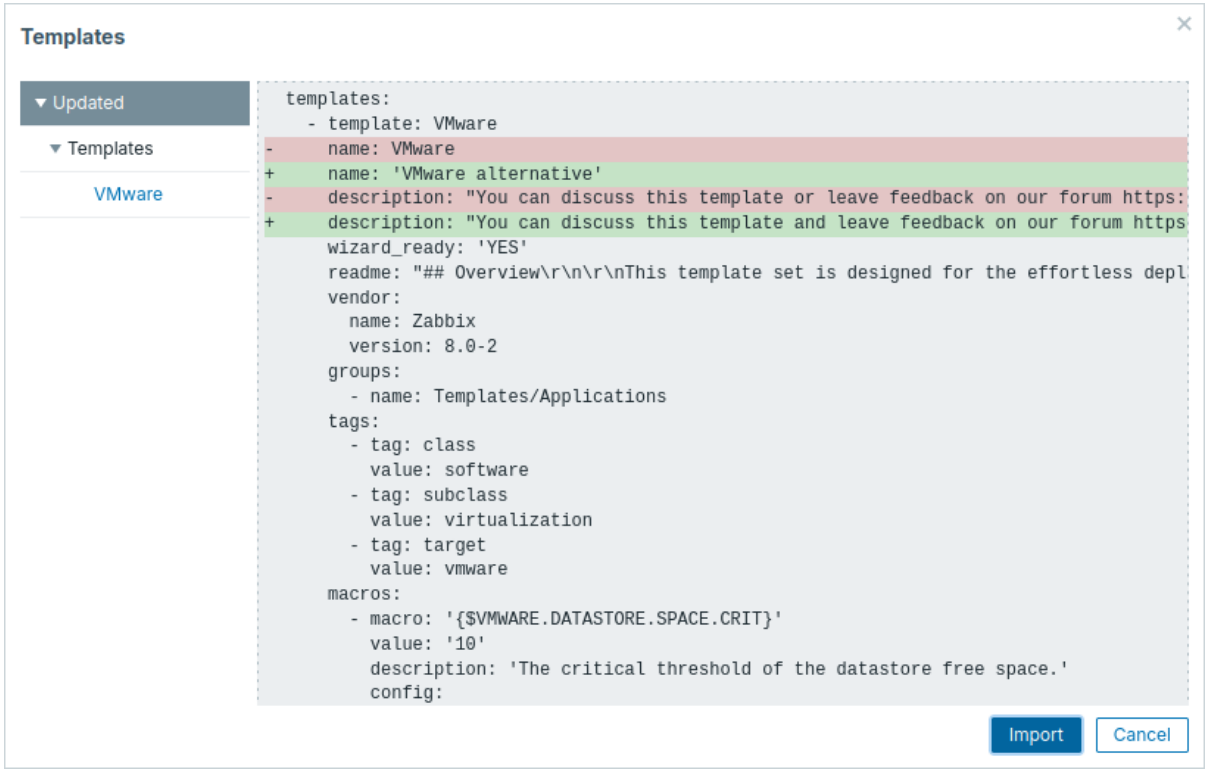
If you click the checkbox in the *All* row, all elements below it will be marked/unmarked.

Import rules:

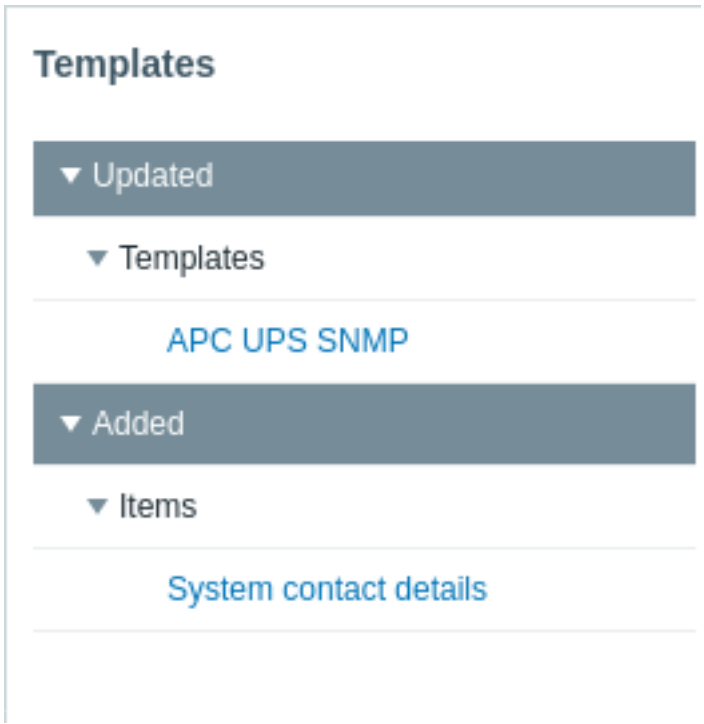
Rule	Description
Update existing	Existing elements will be updated using data from the import file. Otherwise, they will not be updated.
Create new	New elements will be created using data from the import file. Otherwise, they will not be created.
Delete missing	Existing elements not present in the import file will be removed. Otherwise, they will not be removed. If *Delete missing* is marked for *Template linkage*, current template linkage not present in the import file will be unlinked. Entities (items, triggers, graphs, etc.) inherited from the unlinked templates will not be removed (unless the *Delete missing* option is selected for each entity as well).

Attention:
If templates with the same names already exist, the *Delete missing* options should be checked when importing to achieve a clean import. This way, the old items that are no longer in the updated template will be removed (note that it will mean losing history of these old items).

On the next screen, you will be able to view the content of a template being imported. If this is a new template, all elements will be listed in green. If updating an existing template, new template elements will be highlighted in green; removed template elements will be highlighted in red; elements that have not changed will be listed on a gray background.



The menu on the left can be used to navigate through the list of changes. The *Updated* section highlights all changes made to existing template elements. The *Added* section lists new template elements. The elements in each section are grouped by element type; click the gray arrow to expand or collapse the group of elements.



Review template changes and then click *Import* to perform the template import. A success or failure message of the import will be displayed in the frontend.

Export format When a template is exported, Zabbix generates a structured YAML, JSON, or XML format. The export includes template elements, such as template metadata, items, macros, triggers, dashboards, and more.

Each element serves a specific purpose and may contain nested elements.

The following sections describe each element in the export format. The examples use the Linux by Zabbix agent template.

An ellipsis (`...`) indicates elements omitted for brevity. The note (`see table below`) is used when the element is explained in more detail in later sections.

```
zabbix_export:
  version: '8.2'
  template_groups:
    - uuid: 846977d1dfed4968bc5f8bdb363285bc
      name: 'Templates/Operating systems'
  templates: (see table below)
  triggers: (see table below)
  graphs: (see table below)
```

Template metadata

Element	Type	Description
version	string	(required) Export version.
template_groups		(required) Root element for template groups.
uuid	string	(required) Unique identifier for this template group.
name	string	(required) Template group name.
host_groups		Root element for host groups that are used by host prototypes.
uuid	string	(required) Unique identifier for this host group.
name	string	(required) Host group name.
templates		Root element for templates.
triggers		Root element for template trigger elements, which are the same as for template item triggers.
graphs		Root element for template graph elements, which are the same as for host graphs.

Templates

```
zabbix_export:
  (...)
  templates:
    - uuid: f8f7908280354f2abeed07dc788c3747
      template: 'Linux by Zabbix agent'
      name: 'Linux by Zabbix agent'
      description: |
        This is an official Linux template. It requires Zabbix agent 8.2 or newer. (...)
      wizard_ready: 'YES'
      readme: |
        ## Overview

        This is an official Linux template. It requires Zabbix agent 8.2 or newer. (...)
      vendor:
        name: Zabbix
        version: 8.2-2
      groups:
        - name: 'Templates/Operating systems'
      items: (see table below)
      discovery_rules: (see table below)
      tags:
        - tag: class
          value: os
        - tag: target
          value: linux
      macros:
        - macro: '{$AGENT.TIMEOUT}'
          value: 3m
          description: 'Timeout after which agent is considered unavailable. Works only for agents reachab
          config: (see table below)
        - macro: '{$CPU.UTIL.CRIT}'
          value: '90'
          description: 'Critical threshold of CPU utilization expressed in %.'
          config: (see table below)
        (...)
```

```
      dashboards: (see table below)
      valuemaps: (see table below)
    (...)
```

Element	Type	Description
uuid	string	(required) Unique identifier for this template.
template	string	(required) Unique template name.
name	string	Visible template name.
description	text	Template description.
wizard_ready	text	Whether the template is available for selection in the Host Wizard. Possible values:[1] NO (0, default), YES (1). See also: Template object (`wizard_ready`).
readme	text	Template-specific configuration instructions to display in the Host Wizard. Supports Markdown formatting.
vendor		Root element for template vendor (present if the exported template contains vendor data).
name	string	(required) Template vendor name.
version	string	(required) Template version. For out-of-the-box templates, version is displayed as follows: major version of Zabbix, delimiter ("-"), revision number (increased with each new version of the template, and reset with each major version of Zabbix). For example, 8.0-0, 8.0-3, 8.2-0, 8.2-3.
templates		Root element for linked templates.
name	string	(required) Template name.
groups		Root element for template groups.
name	string	(required) Template group name.
items		Root element for template items.
discovery_rules		Root element for template low-level discovery rules.
httptests		Root element for template web scenarios.
tags		Root element for template tags.
tag	string	(required) Tag name.
value	string	Tag value.
macros		Root element for template user macros.
macro	string	(required) User macro name.
type	string	User macro type. Possible values:[1] TEXT (0, default), SECRET_TEXT (1), VAULT (2). See also: User macro object (`type`).
value	string	User macro value.
description	string	User macro description.
config		Root element for template macro configuration responsible for how the macro is displayed in the Host Wizard.
dashboards		Root element for template dashboards.
valuemaps		Root element for template value maps.

Template items

```
zabbix_export:
  (...)
  templates:
    (...)
      items:
        - uuid: f94f9f4699e94c369e6c98b2a2f485ce
          name: 'Zabbix agent ping'
          key: agent.ping
          description: 'The agent always returns "1" for this item. May be used in combination with `nodat
          valuemap:
            name: 'Zabbix agent ping status'
          tags:
            - tag: component
              value: system
        (...)
        - uuid: 58818005e76d46dda14d6592f601ab00
          name: 'Number of installed packages'
```

```
      key: system.sw.packages.get
      delay: 1h
      preprocessing: (see table below)
      tags:
        - tag: component
          value: os
      triggers: (see table below)
    - uuid: 403cebed115441369e94d35d070ca7b8
      name: 'Memory utilization'
      type: DEPENDENT
      key: vm.memory.utilization
      value_type: FLOAT
      units: '%'
      description: 'The percentage of used memory is calculated as `100-pavailable`.'
      preprocessing: (see table below)
      master_item:
        key: 'vm.memory.size[pavailable]'
      tags:
        - tag: component
          value: memory
      triggers: (see table below)
    (...)
```

Element	Type	Description
uuid	string	(required) Unique identifier for this item.
name	string	(required) Item name.
type	string	Item type. Possible values:[1] ZABBIX_PASSIVE (0, default), TRAP (2), SIMPLE (3), INTERNAL (5), ZABBIX_ACTIVE (7), EXTERNAL (10), ODBC (11), IPMI (12), SSH (13), TELNET (14), CALCULATED (15), JMX (16), SNMP_TRAP (17), DEPENDENT (18), HTTP_AGENT (19), SNMP_AGENT (20), ITEM_TYPE_SCRIPT (21), ITEM_TYPE_BROWSER (22), TELEMETRY_QUERY (24).
snmp_oid	string	(required for SNMP_AGENT items) SNMP object ID.
key	string	(required) Item key.
delay	string	Update interval of the item. Default: 1m. The value will always be 0 for TRAP items.
history	string	Time period (using time suffix, user macro or LLD macro) of how long the history data should be stored. Default: 31d.
trends	string	Time period (using time suffix, user macro or LLD macro) of how long the trends data should be stored. Default: 365d.
status	string	Item status. Possible values:[1] ENABLED (0, default), DISABLED (1).
value_type	string	Received value type. Possible values:[1] FLOAT (0), CHAR (1), LOG (2), UNSIGNED (3, default), TEXT (4), BINARY (5), JSON (6).
allowed_hosts	string	List of comma-delimited IP addresses of hosts allowed to send data for the item. Supported for TRAP and HTTP_AGENT items.
units	string	Received value units (bps, B, etc.).
params	text	Additional parameters depending on the type of the item (executed script for SSH and TELNET items; SQL query for ODBC items; formula for CALCULATED items; the script for ITEM_TYPE_SCRIPT and ITEM_TYPE_BROWSER items).
ipmi_sensor	string	IPMI sensor. Supported for IPMI items.
authtype	string	Authentication type. Supported for SSH and HTTP_AGENT items. Possible values for SSH items:[1] PASSWORD (0, default), PUBLIC_KEY (1). Possible values for HTTP_AGENT items:[1] NONE (0, default), BASIC (1), NTLM (2), Kerberos (3), or Digest (4).

Element	Type	Description
username	string	(required for SSH and TELNET items) Username for authentication. Supported for SIMPLE, ODBC, JMX and HTTP_AGENT items. When used for JMX items, `password` (see below) should also be specified or both elements should be left blank.
password	string	(required for SSH and TELNET items) Password for authentication. Supported for SIMPLE, ODBC, JMX and HTTP_AGENT items. When used for JMX items, `username` (see above) should also be specified or both elements should be left blank.
publickey	string	(required for SSH items) Name of the public key file.
privatekey	string	(required for SSH items) Name of the private key file.
description	text	Item description.
inventory_link	string	Host inventory field that is populated by the item. Possible values:[1] NONE (0), ALIAS (4), etc. (see Host inventory for supported fields).
valuemap		Root element for item value maps.
name	string	(required) Name of the value map to use for the item.
logtimefmt	string	Format of the time in log entries. Supported for items of LOG value type.
preprocessing		Root element for item value preprocessing.
step		Root element for template item value preprocessing steps.
jmx_endpoint	string	JMX endpoint. Supported for JMX items.
master_item		(required for DEPENDENT items) Root element for dependent item's master item.
key	string	(required) Dependent item's master item key.
timeout	string	Item data polling request timeout. Supported for Timeouts list of item types.
url	string	(required for HTTP_AGENT items) URL string.
query_fields		Root element for query parameters. Supported for HTTP_AGENT items.
name	string	(required for HTTP_AGENT items) Query parameter name.
value	string	Query parameter value. Supported for HTTP_AGENT items.
parameters		Root element for user-defined parameters. Supported for ITEM_TYPE_SCRIPT and ITEM_TYPE_BROWSER items.
name	string	(required for ITEM_TYPE_SCRIPT and ITEM_TYPE_BROWSER items) User-defined parameter name.
value	string	User-defined parameter value. Supported for ITEM_TYPE_SCRIPT and ITEM_TYPE_BROWSER items.
posts	string	HTTP(S) request body data. Supported for HTTP_AGENT items.
status_codes	string	Ranges of required HTTP status codes, separated by commas. Supported for HTTP_AGENT items.
follow_redirects	string	Follow response redirects while polling data. Supported for HTTP_AGENT items. Possible values:[1] NO (0), YES (1, default).
post_type	string	Type of post data body. Supported for HTTP_AGENT items. Possible values:[1] RAW (0, default), JSON (2), XML (3).
http_proxy	string	HTTP(S) proxy connection string. Supported for HTTP_AGENT items.
headers		Root element for HTTP(S) request headers. Supported for HTTP_AGENT items.
name	string	(required for HTTP_AGENT items) Header name.
value	string	(required for HTTP_AGENT items) Header value.
retrieve_mode	string	What part of response should be stored. Supported for HTTP_AGENT items. Possible values:[1] BODY (0, default), HEADERS (1), BOTH (2).
request_method	string	Request method type. Supported for HTTP_AGENT items. Possible values:[1] GET (0, default), POST (1), PUT (2), HEAD (3).
output_format	string	How to process response. Supported for HTTP_AGENT items. Possible values:[1] RAW (0, default), JSON (1).

Element	Type	Description
allow_traps	string	Allow the value to be populated similarly to the trapper item. Supported for HTTP_AGENT items. Possible values:[1] NO (0, default), YES (1).
ssl_cert_file	string	Public SSL Key file path. Supported for HTTP_AGENT items.
ssl_key_file	string	Private SSL Key file path. Supported for HTTP_AGENT items.
ssl_key_password	string	Password for SSL Key file. Supported for HTTP_AGENT items.
verify_peer	string	Whether to validate that the host's certificate is authentic. Supported for HTTP_AGENT items. Possible values:[1] NO (0, default), YES (1).
verify_host	string	Whether to validate that the host name for the connection matches the one in the host's certificate. Supported for HTTP_AGENT items. Possible values:[1] NO (0, default), YES (1).
time_shift	string	How far back (before the current time) ends the time range that the item queries. Supported for TELEMETRY_QUERY items.
lookback_limit	string	How far back (before the current time) the item queries for data when there is no previous query to continue from (e.g., on the item's first run), or when the previous query is older than `delay` (e.g., because Zabbix server was down). Supported for TELEMETRY_QUERY items.
granularity	string	The length of a piece that the query's full time range is split into. Supported for TELEMETRY_QUERY items.
query		(required for TELEMETRY_QUERY items) Root element for the telemetry query. Defines what the item queries from the data source, how it filters the data, and how it aggregates the result.
signal_type	string	(required) Signal type that the item queries. Supported for TELEMETRY_QUERY items. Possible values:[1] TRACES (0), METRICS (1), LOGS (2).
metric_point_type	string	(required for METRICS signal type) Metric point type that the item queries. Supported for TELEMETRY_QUERY items. Possible values:[1] SUM (0), GAUGE (1), HISTOGRAM (2), EXPONENTIAL_HISTOGRAM (3).
columns		(required) Root element for telemetry data columns that the item queries. This element has the following elements: `column (string)` - (required) Name of the telemetry data column that the item queries. `attribute_key (string)` - (required for columns that hold key-value pairs) Name of the key that the item queries, within the column that holds key-value pairs.
aggregated_columns		(required) Root element for telemetry data aggregated columns that the item queries. This element has the following elements: `column (string)` - (required for MIN, MAX, AVG, SUM, and PERCENTILE functions) Name of the telemetry data column that the item queries. `function (string)` - (required) Function that the item applies to the column records. Possible values:[1] MIN (1), MAX (2), AVG (3), COUNT (4), SUM (5), PERCENTILE (8). `parameters (array)` - (required for PERCENTILE functions) Percentile that the item calculates from the column records. `alias (string)` - (required) Unique name for the calculated value, used to identify it in the item value.

Element	Type	Description
filter		(required) Root element for telemetry query filter for column records in the query. This element has the following elements: `evaltype` (`string`) - (required) Filter condition evaluation method. Possible values:[1] AND_OR (0), AND (1), OR (2), FORMULA (3). `formula` (`string`) - (required for FORMULA evaluation methods) User-defined expression to be used for evaluating conditions of filters with a custom expression. `conditions` (`string`) - (required) Root element for telemetry data filter conditions a record in a telemetry data column must meet. The `conditions` element has the following elements: `column` (`string`) - (required) Name of the telemetry data column that the filter conditions apply to. `attribute_key` (`string`) - (required for columns that hold key-value pairs) Name of the key that the condition applies to, within the column that holds key-value pairs. `operator` (`string`) - (required) Filter condition operator. Possible values for single-value columns:[1] EQUAL (0), NOT_EQUAL (1), LIKE (2), NOT_LIKE (3). Possible values for columns that hold key-value pairs:[1] EQUAL (0), NOT_EQUAL (1), EXISTS (12). `formulaid` (`string`) - (required for FORMULA evaluation methods) Arbitrary unique ID that is used to reference the condition from a custom expression. `value` (`string`) - (required for EQUAL, NOT_EQUAL, LIKE, and NOT LIKE operations) Value to compare with.
tags		Root element for item tags.
tag	string	(required) Tag name.
value	string	Tag value.
triggers		Root element for template item triggers.

Note:
See also: Item object (refer to the relevant property with a matching name).

Template item value preprocessing steps

```
zabbix_export:
  (...)
  templates:
    (...)
      items:
        (...)
        - uuid: 58818005e76d46dda14d6592f601ab00
          (...)
          preprocessing:
            - type: JSONPATH
              parameters:
                - $.length()
            - type: DISCARD_UNCHANGED_HEARTBEAT
              parameters:
                - 12h
          (...)
```

Element	Type	Description
type	string	(required) The item value preprocessing step type. Possible values:[1] MULTIPLIER (1), RTRIM (2), LTRIM (3), TRIM (4), REGEX (5), BOOL_TO_DECIMAL (6), OCTAL_TO_DECIMAL (7), HEX_TO_DECIMAL (8), SIMPLE_CHANGE (9, calculated: received value - previous value), CHANGE_PER_SECOND (10, calculated: (received value - previous value)/(time now - time of last check)), XMLPATH (11), JSONPATH (12), IN_RANGE (13), MATCHES_REGEX (14), NOT_MATCHES_REGEX (15), CHECK_JSON_ERROR (16), CHECK_XML_ERROR (17), CHECK_REGEX_ERROR (18), DISCARD_UNCHANGED (19), DISCARD_UNCHANGED_HEARTBEAT (20), JAVASCRIPT (21), PROMETHEUS_PATTERN (22), PROMETHEUS_TO_JSON (23), CSV_TO_JSON (24), STR_REPLACE (25), CHECK_NOT_SUPPORTED (26), XML_TO_JSON (27), SNMP_WALK_VALUE (28), SNMP_WALK_TO_JSON (29), SNMP_GET_VALUE (30).
parameters		(required) Root element for parameters of the item value preprocessing step.
parameter	string	Individual parameter of the item value preprocessing step.
error_handler	string	Action type used in case of preprocessing step failure. Possible values:[1] ORIGINAL_ERROR (0, default), DISCARD_VALUE (1), CUSTOM_VALUE (2), CUSTOM_ERROR (3).
error_handler_params	string	Error handler parameters.

Note:
See also: Item preprocessing object (refer to the relevant property with a matching name).

Template item triggers

```
zabbix_export:
  (...)
  templates:
    (...)
      items:
        (...)
        - uuid: 58818005e76d46dda14d6592f601ab00
          (...)
          triggers:
            - uuid: b950c306394f4b3c902060a8273cbcde
              expression: 'change(/Linux by Zabbix agent/system.sw.packages.get)<>0'
              name: 'Linux: Number of installed packages has been changed'
              priority: WARNING
              manual_close: 'YES'
              tags:
                - tag: scope
                  value: notice
        (...)
        - uuid: 403cebed115441369e94d35d070ca7b8
          (...)
          triggers:
            - uuid: cfd395b1cde74ef18a5e5f840bd5142a
              expression: 'min(/Linux by Zabbix agent/vm.memory.utilization,5m)>{$MEMORY.UTIL.MAX}'
              name: 'Linux: High memory utilization'
              event_name: 'Linux: High memory utilization (>{$MEMORY.UTIL.MAX}% for 5m)'
              priority: AVERAGE
              description: 'The system is running out of free memory.'
              dependencies:
                - name: 'Linux: Lack of available memory'
                  expression: 'max(/Linux by Zabbix agent/vm.memory.size[available],5m)<{$MEMORY.AVAILABLE
              tags:
                - tag: scope
                  value: capacity
                - tag: scope
                  value: performance
        (...)
```

Element	Type	Description
uuid	string	(required) Unique identifier for this trigger.
expression	string	(required) Trigger expression.
recovery_mode	string	Basis for generating OK events. Possible values:[1] EXPRESSION (0, default), RECOVERY_EXPRESSION (1), NONE (2).
recovery_expression	string	Trigger recovery expression.
correlation_mode	string	Correlation mode (no event correlation or event correlation by tag). Possible values:[1] DISABLED (0, default), TAG_VALUE (1).
correlation_tag	string	The tag name to be used for event correlation.
name	string	(required) Trigger name.
event_name	string	Event name.
opdata	string	Operational data.
url_name	string	Label for the URL associated with the trigger.
url	string	URL associated with the trigger.
status	string	Trigger status. Possible values:[1] ENABLED (0, default), DISABLED (1).
priority	string	Trigger severity. Possible values:[1] NOT_CLASSIFIED (0, default), INFO (1), WARNING (2), AVERAGE (3), HIGH (4), DISASTER (5).
description	text	Trigger description.
type	string	Event generation type (single problem event or multiple problem events). Possible values:[1] SINGLE (0, default), MULTIPLE (1).
manual_close	string	Manual closing of problem events. Possible values:[1] NO (0, default), YES (1).
dependencies		Root element for dependencies.
name	string	(required) Dependency trigger name.
expression	string	(required) Dependency trigger expression.
recovery_expression	string	Dependency trigger recovery expression.
tags		Root element for trigger tags.
tag	string	(required) Tag name.
value	string	Tag value.

Note:
See also: Trigger object (refer to the relevant property with a matching name).

Template low-level discovery rules

```
zabbix_export:
  (...)
  templates:
    (...)
      discovery_rules:
        - uuid: acfdea9c46ef48c68e6636f43b8f96a2
          name: 'Network interface discovery'
          key: net.if.discovery
          delay: 1h
          filter: (see table below)
          description: 'The discovery of network interfaces.'
          item_prototypes: (see table below)
          trigger_prototypes: (see table below)
          graph_prototypes: (see table below)
        (...)
```

Attention:
Most template low-level discovery rule elements are the same as for template items. The table below describes those elements that differ from template items.

Element	Type	Description
type	string	Item type. Possible values:[1] ZABBIX_PASSIVE (0, default), TRAP (2), SIMPLE (3), INTERNAL (5), ZABBIX_ACTIVE (7), EXTERNAL (10), ODBC (11), IPMI (12), SSH (13), TELNET (14), JMX (16), DEPENDENT (18), HTTP_AGENT (19), SNMP_AGENT (20), ITEM_TYPE_SCRIPT (21), ITEM_TYPE_BROWSER (22).
key	string	(required) The low-level discovery rule key.
filter		Root element for template low-level discovery rule filters.
lifetime	string	Time period (using seconds, time suffix, or user macro), after which no longer discovered resources will be deleted. Default: 7d.
lifetime_type	string	Scenario for deleting lost LLD resources. Possible values: DELETE_NEVER, DELETE_IMMEDIATELY, DELETE_AFTER.
enabled_lifetime	string	Time period (using seconds, time suffix, or user macro), after which no longer discovered resources will be disabled.
enabled_lifetime_type	string	Scenario for disabling lost LLD resources. Possible values: DISABLE_NEVER, DISABLE_IMMEDIATELY, DISABLE_AFTER.
item_prototypes		Root element for template item prototype elements, which are the same as for template items.
trigger_prototypes		Root element for template trigger prototype elements, which are the same as for template item triggers.
graph_prototypes		Root element for template graph prototype elements, which are the same as for host graphs.
host_prototypes		Root element for template host prototype elements, which are the same as for hosts.
parent_discovery_rule		Root element for the parent low-level discovery rule (or rule prototype) of the low-level discovery rule prototype. This property denotes that it is an LLD rule prototype, direct child of the referenced rule/rule prototype.
key	string	(required) The parent low-level discovery rule (or rule prototype) key.
master_item	string	(required for DEPENDENT rules) Root element for the dependent rule's master item.
lld_macro_paths		Root element for low-level discovery rule macro paths.
lld_macro	string	(required) Low-level discovery rule macro name.
path	string	(required) Selector for value, which will be assigned to the corresponding macro.
preprocessing		Root element for low-level discovery rule value preprocessing.
step		Root element for low-level discovery rule value preprocessing step elements, which are the same as for template item value preprocessing steps, except with fewer possible values. See also: LLD rule preprocessing object.
overrides		Root element for low-level discovery rule override rules.
name	string	(required) Unique override name.
step	string	(required) Unique order number of the override.
stop	string	Stop processing next overrides if matches.
filter		Root element for template low-level discovery rule override rule filter elements, which are the same as for template low-level discovery rule filters.
operations		Root element for template low-level discovery rule override operations.

Note:
See also: LLD rule object (refer to the relevant property with a matching name).

Template low-level discovery rule filters

```
zabbix_export:
  (...)
  templates:
    (...)
      discovery_rules:
        - uuid: acfdea9c46ef48c68e6636f43b8f96a2
          (...)
          filter:
            evaltype: AND
            conditions:
              - macro: '{#IFNAME}'
                value: '{$NET.IF.IFNAME.MATCHES}'
              - macro: '{#IFNAME}'
```

```
            value: '{$NET.IF.IFNAME.NOT_MATCHES}'
            operator: NOT_MATCHES_REGEX
      (...)
    - uuid: 2bbdc79f082d4c618e01bec625e9c90a
      (...)
      filter:
        evaltype: AND
        conditions:
          - macro: '{#DEVNAME}'
            value: '{$VFS.DEV.DEVNAME.MATCHES}'
          - macro: '{#DEVNAME}'
            value: '{$VFS.DEV.DEVNAME.NOT_MATCHES}'
            operator: NOT_MATCHES_REGEX
          - macro: '{#DEVTYPE}'
            value: disk
      (...)
```

Element	Type	Description
evaltype	string	Override filter condition evaluation method. Possible values:[1] AND_OR (0, default), AND (1), OR (2), FORMULA (3).
formula	string	Custom calculation formula for filter conditions.
conditions		Root element for filter conditions.
macro	string	(required) The name of the low-level discovery macro to perform the check on.
value	string	Value to compare with.
operator	string	Condition operator. Possible values:[1] MATCHES_REGEX (8, default), NOT_MATCHES_REGEX (9).
formulaid	string	(required) Arbitrary unique ID that is used to reference a condition from the custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward.

Note:
See also: LLD rule filter object (refer to the relevant property with a matching name).

Template low-level discovery rule override operations

Element	Type	Description
operationobject	string	Object to which the operation is applied. Possible values:[1] ITEM_PROTOTYPE (0), TRIGGER_PROTOTYPE (1), GRAPH_PROTOTYPE (2), HOST_PROTOTYPE (3).
operator	string	Override condition operator. Possible values:[1] EQUAL (1), NOT_EQUAL (2), LIKE (3), NOT_LIKE (4), REGEXP (5), NOT_REGEXP (6).
value	string	A regular expression or a string for the override condition operator.
status	string	Status of the object upon the override operation.
discover	string	Whether the object is added as a result of the discovery.
delay	string	Update interval set for the item prototype upon the override operation.
history	string	History storage period set for the item prototype upon the override operation.
trends	string	Trend storage period set for the item prototype upon the override operation.
severity	string	Trigger prototype severity set upon the override operation.
tags		Root element for the tags set for the object upon the override operation.
tag	string	(required) Tag name.
value	string	Tag value.
templates		Root element for the templates linked to the host prototype upon the override operation.
name	string	(required) Template name.
inventory_mode	string	Host prototype inventory mode set upon the override operation.

Note:
See also: LLD rule override operation object (refer to the relevant property with a matching name).

Template web scenarios

Element	Type	Description
uuid	string	(required) Unique identifier for this web scenario.
name	string	(required) Web scenario name.
delay	string	Frequency (using seconds, time suffix, or user macro) of executing web scenario. Default: 1m.
attempts	integer	The number of attempts for executing web scenario steps. Possible values:[1] 1-10 (default: 1).
agent	string	Client agent. Zabbix will pretend to be the selected browser. This is useful when a website returns different content for different browsers. Default: Zabbix.
http_proxy	string	Proxy that will be used by the web scenario given as: `http://[username[:password]@]proxy.example.com[:port]`
variables		Root element for web scenario variables that may be used in scenario steps.
name	string	(required) Variable name.
value	text	(required) Variable value.
headers		Root element for HTTP headers to be sent when performing a request. Headers should be listed using the same syntax as they would appear in the HTTP protocol.
name	string	(required) Header name.
value	text	(required) Header value.
status	string	Web scenario status. Possible values:[1] ENABLED (0, default), DISABLED (1).
authentication	string	Authentication method. Possible values:[1] NONE (0, default), BASIC (1), NTLM (2), Kerberos (3), or Digest (4).
http_user	string	User name used for BASIC (HTTP), NTLM, Kerberos, or Digest authentication.
http_password	string	Password used for BASIC (HTTP), NTLM, Kerberos, or Digest authentication.
verify_peer	string	Verify the SSL certificate of the web server. Possible values:[1] NO (0, default), YES (1).
verify_host	string	Verify that the *Common Name* field or the *Subject Alternate Name* field of the web server certificate matches. Possible values:[1] NO (0, default), YES (1).
ssl_cert_file	string	Name of the SSL certificate file used for client authentication (must be in PEM format).
ssl_key_file	string	Name of the SSL private key file used for client authentication (must be in PEM format).
ssl_key_password	string	SSL private key file password.
steps		(required) Root element for template web scenario steps.
tags		Root element for web scenario tags.
tag	string	(required) Tag name.
value	string	Tag value.

Note:
See also: Web scenario object (refer to the relevant property with a matching name).

Template web scenario steps

Element	Type	Description
name	string	(required) Web scenario step name.
url	string	(required) URL for monitoring.
query_fields		Root element for query parameters (an array of HTTP fields to be added to the URL when performing a request).
name	string	(required) Query parameter name.
value	string	Query parameter value.
posts		Root element for HTTP POST variables (a string (raw post data) or an array of HTTP fields (form field data)).
name	string	(required) Post field name.
value	string	(required) Post field value.

Element	Type	Description
variables		Root element of step-level variables (macros) that should be applied after this step. If the variable value has a 'regex:' prefix, then its value is extracted from the data returned by this step according to the regular expression pattern following the 'regex:' prefix
name	string	(required) Variable name.
value	text	(required) Variable value.
headers		Root element for HTTP headers to be sent when performing a request.
name	string	(required) Header name.
value	text	(required) Header value.
follow_redirects	string	Follow HTTP redirects. Possible values:[1] NO (0), YES (1, default).
retrieve_mode	string	HTTP response retrieve mode. Possible values:[1] BODY (0, default), HEADERS (1), BOTH (2).
timeout	string	Timeout (using seconds, time suffix, or user macro) of step execution. Default: 15s.
required	string	Text that must be present in the response (ignored if empty).
status_codes	string	A comma-delimited list of accepted HTTP status codes (e.g., 200-201,210-299; ignored if empty).

Note:
See also: Web scenario step object (refer to the relevant property with a matching name).

Template macro configuration

```
zabbix_export:
  (...)
  templates:
    - uuid: f8f7908280354f2abeed07dc788c3747
      (...)
      macros:
        - macro: '{$AGENT.TIMEOUT}'
          (...)
          config:
            type: TEXT
            priority: '1'
            label: 'Seconds since the last Zabbix agent seen'
            description: 'Timeout after which agent is considered unavailable.'
        - macro: '{$CPU.UTIL.CRIT}'
          (...)
          config:
            type: TEXT
            priority: '2'
            section_name: Thresholds
            label: 'Threshold of CPU utilization expressed'
            description: 'Critical threshold of CPU utilization expressed in %. In the range from 0 to 100
            regex: '^-?([0-9]+|(([0-9]+)\.([0-9]+)))$'
        (...)
        - macro: '{$IFCONTROL}'
          (...)
          config:
            type: CHECKBOX
            priority: '19'
            label: 'Interfaces control'
            description: 'Fire a trigger when the interface operational status changes to "Link down".'
            options:
              - checked: '1'
                unchecked: '0'
        (...)
```

Element	Type	Description
type	string	(required) Type of macro input field. Possible values:[1] NOCONF (0), TEXT (1), LIST (2), CHECKBOX (3).
priority	string	Macro position in the macro list.
section_name	string	Label of the collapsible section where the macro is grouped.
label	string	(required for TEXT, LIST, CHECKBOX macros) Macro label.
description	text	Macro help text. Supports Markdown formatting.
required	string	Marks the macro as mandatory. Possible values:[1] NO (0), YES (1). Supported for TEXT and LIST macros.
regex	string	Regular expression to validate user input in a textbox field. Supported for TEXT macros.
options		Root element for LIST items or CHECKBOX values.
value	string	(required for LIST macros) LIST item value.
text	string	(required for LIST macros) LIST item text.
checked	string	(required for CHECKBOX macros) Value representing a checked state. Possible values:[1] 0 (false), 1 (true). Supported for CHECKBOX macros.
unchecked	string	(required for CHECKBOX macros) Value representing an unchecked state. Possible values:[1] 0 (false), 1 (true). Supported for CHECKBOX macros.

Note:
See also: Macro configuration object (refer to the relevant property with a matching name).

Template dashboards

```
zabbix_export:
  (...)
  templates:
    (...)
      dashboards:
        - uuid: c689ad3115fd46a4b927d1f70ee2e5a4
          name: Filesystems
          pages:
            - name: Overview
              widgets: (see table below)
        (...)
```

Element	Type	Description
uuid	string	(required) Unique identifier for this dashboard.
name	string	(required) Template dashboard name.
display pe-riod	integer	Display period of dashboard pages.
auto_start	string	Slideshow auto start. Possible values:[1] NO (0), YES (1, default).
pages		Root element for template dashboard pages.
name	string	Page name.
display period	integer	Slideshow interval.
sortorder	integer	Page sorting order.
widgets		Root element for template dashboard widgets.

Note:
See also: Template dashboard object (refer to the relevant property with a matching name).

Template dashboard widgets

```yaml
zabbix_export:
  (...)
  templates:
    (...)
      dashboards:
        - uuid: c689ad3115fd46a4b927d1f70ee2e5a4
          name: Filesystems
          pages:
            - name: Overview
              widgets:
                - type: graphprototype
                  width: '72'
                  height: '5'
                  fields:
                    - type: INTEGER
                      name: columns
                      value: '3'
                    - type: GRAPH_PROTOTYPE
                      name: graphid.0
                      value:
                        host: 'Linux by Zabbix agent'
                        name: 'FS [{#FSTYPE}({#FSNAME})]: Space utilization chart (relative to total)'
                    - type: STRING
                      name: reference
                      value: AAAEB
                - type: graphprototype
                  'y': '5'
                  width: '72'
                  height: '5'
                  fields:
                    - type: INTEGER
                      name: columns
                      value: '1'
                    - type: GRAPH_PROTOTYPE
                      name: graphid.0
                      value:
                        host: 'Linux by Zabbix agent'
                        name: 'FS [{#FSTYPE}({#FSNAME})]: Space usage graph, in % (relative to max availab
                    - type: STRING
                      name: reference
                      value: AAAEC
        (...)
```

Element	Type	Description
type	string	(required) Widget type.
name	string	Widget name.
x	integer	Horizontal position from the left side of the template dashboard. Possible values:[1] 0-71.
y	integer	Vertical position from the top of the template dashboard. Possible values:[1] 0-63.
width	integer	Widget width. Possible values:[1] 1-72.
height	integer	Widget height. Possible values:[1] 1-64.
hide_header	string	Hide widget header. Possible values:[1] NO (0, default), YES (1).
fields		Root element for the template dashboard widget fields.
type	string	(required) Widget field type. Possible values:[1] INTEGER (0), STRING (1), ITEM (4), ITEM_PROTOTYPE (5), GRAPH (6), GRAPH_PROTOTYPE (7), MAP (8), SERVICE (9), SLA (10), USER (11), ACTION (12), MEDIA_TYPE (13).
name	string	(required) Widget field name.

Element	Type	Description
value	mixed	(required) Widget field value, depending on the field type.

Note:
See also: Template dashboard widget object (refer to the relevant property with a matching name).

Template value maps

```
zabbix_export:
  (...)
  templates:
    (...)
      valuemaps:
        - uuid: 8c048c6cca8248f2860c208e8db0f59e
          name: ifOperStatus
          mappings:
            - value: '0'
              newvalue: unknown
            - value: '1'
              newvalue: notpresent
            - value: '2'
              newvalue: down
            - value: '3'
              newvalue: lowerlayerdown
            - value: '4'
              newvalue: testing
            - value: '5'
              newvalue: dormant
            - value: '6'
              newvalue: up
        (...)
```

Element	Type	Description
uuid	string	(required) Unique identifier for this value map.
name	string	(required) Value map name.
mapping		Root element for mappings.
type	string	Mapping match type. Possible values:[1] EQUAL (0, default), GREATER_OR_EQUAL (2), LESS_OR_EQUAL (3), IN_RANGE (4), REGEXP (5), DEFAULT (6).
value	string	Original value.
newvalue	string	(required) Value to which the original value is mapped to.

Note:
See also: Value map object (refer to the relevant property with a matching name).

Footnotes [1] API integer values in brackets, for example, ENABLED (0), are mentioned only for reference. For more information, see the linked API object page in the table entry or at the end of each section.

Hosts

Overview

Hosts are exported with many related objects and object relations.

Host export contains:

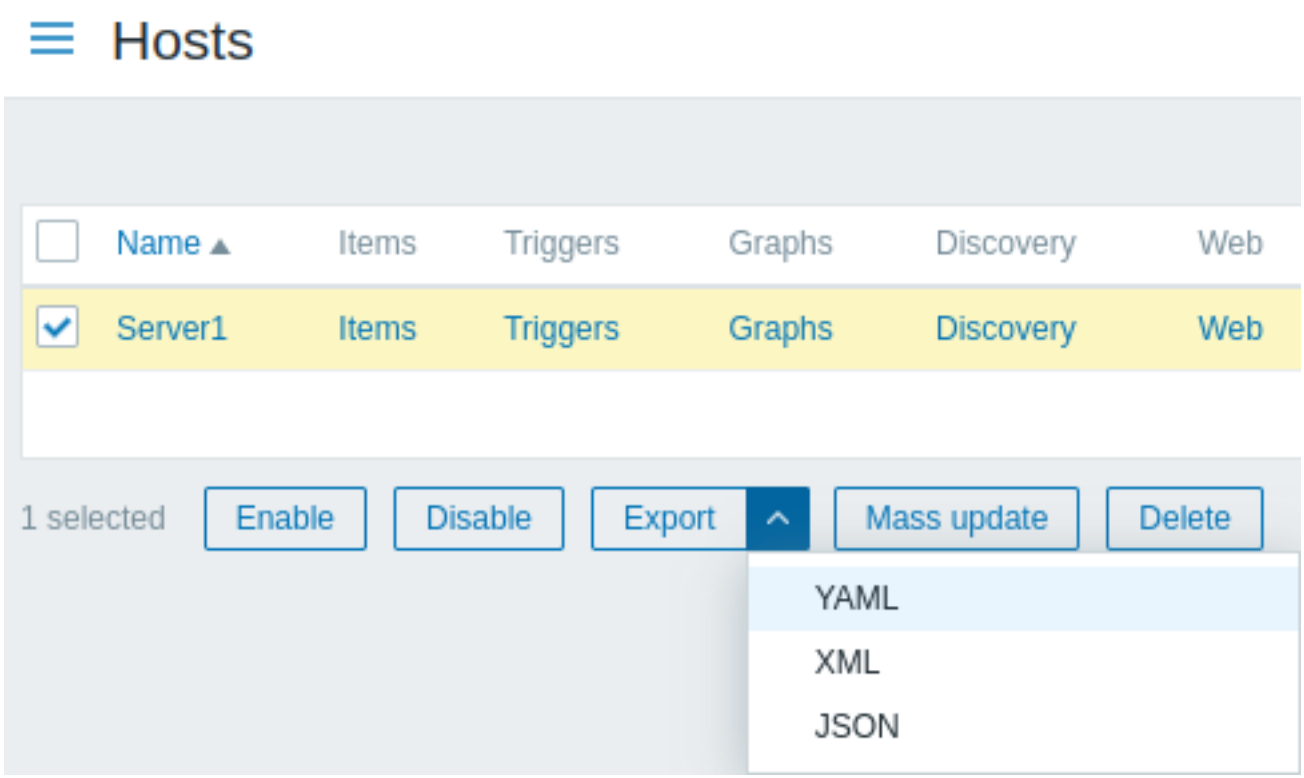
- Linked host groups
- Host data
- Template linkage
- Host group linkage

- Host interfaces
- Directly linked items
- Directly linked triggers
- Directly linked discovery rules with all prototypes
- Directly linked web scenarios
- Host macros
- Host inventory data
- Value maps
- Linked graphs

Exporting

To export hosts, do the following:

1. Go to *Data collection* > *Hosts*.
2. Mark the checkboxes of the hosts to export.
3. Click *Export* below the list to export in the default YAML format, or click the arrow to select a different export format.



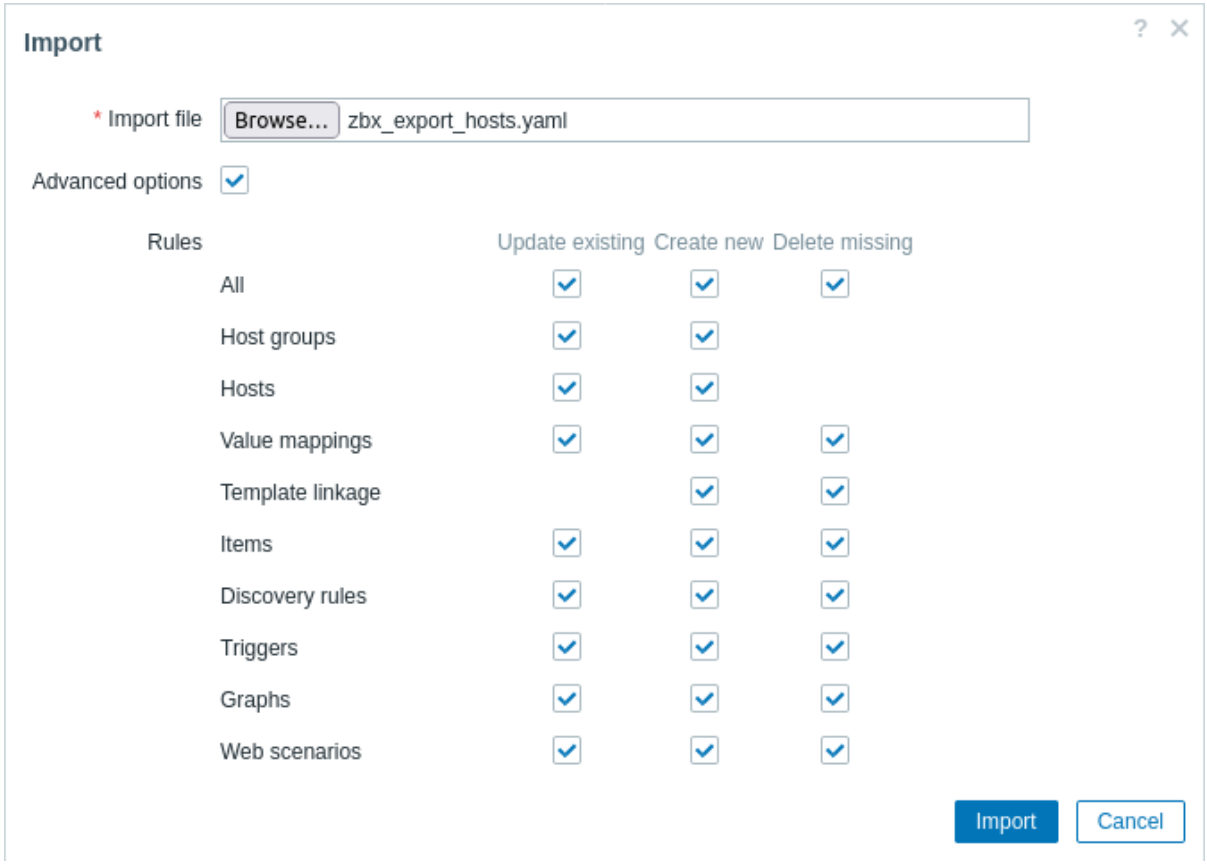
Depending on the selected format, hosts are exported to a local file with a default name:

- `zabbix_export_hosts.yaml` - in YAML export;
- `zabbix_export_hosts.xml` - in XML export;
- `zabbix_export_hosts.json` - in JSON export.

Importing

To import hosts, do the following:

1. Go to *Data collection* > *Hosts*.
2. Click *Import* in the upper-right corner.
3. Select the import file.
4. Mark the checkboxes of the import rules you want to apply (see below).
5. Click *Import* in the lower-right corner of the configuration form.



If you mark the *Advanced options* checkbox, a detailed list of all importable elements will be displayed - mark or unmark each import rule as required.

If you click the checkbox in the *All* row, all elements below it will be marked/unmarked.

Import rules:

Rule	Description
Update existing	Existing elements will be updated using data from the import file. Otherwise, they will not be updated.
Create new	New elements will be created using data from the import file. Otherwise, they will not be created.
Delete missing	Existing elements not present in the import file will be removed. Otherwise, they will not be removed. If *Delete missing* is marked for *Template linkage*, current template linkage not present in the import file will be unlinked. Entities (items, triggers, graphs, etc.) inherited from the unlinked templates will not be removed (unless the *Delete missing* option is selected for each entity as well).

You cannot review changes to hosts during the import. A success or failure message of the import will be displayed in the frontend.

Export format

When a host is exported, Zabbix generates a structured YAML, JSON, or XML format. The export includes host elements, such as host metadata, items, macros, triggers, and more.

Each element serves a specific purpose and may contain nested elements.

The following sections describe each element in the export format. The examples use a host with the MySQL by Zabbix agent 2 template. For element demonstration purposes, the template is unlinked after the host is created. Unlinking keeps all elements in the configuration (unlike unlinking and clearing, which removes them).

An ellipsis (`...`) indicates elements omitted for brevity. The note (`see table below`) is used when the element is explained in more detail in later sections.

```
zabbix_export:
  version: '8.2'
```

```
  host_groups:
    - uuid: 748ad4d098d447d492bb935c907f652f
      name: Databases
  hosts: (see table below)
  graphs: (see table below)
```

Host metadata

Element		Type	Description
version		string	(required) Zabbix version.
host_groups			(required) Root element for host groups.
	uuid	string	(required) Unique identifier for this host group.
	name	string	(required) Host group name.
hosts			Root element for hosts.
graphs			Root element for host graphs.

Hosts

```
zabbix_export:
  (...)
  hosts:
    - host: 'MySQL server'
      name: 'MySQL server'
      groups:
        - name: Databases
      interfaces: (see table below)
      items: (see table below)
      discovery_rules: (see table below)
      tags:
        - tag: location
          value: Riga
      macros:
        (...)
        - macro: '{$MYSQL.DSN}'
          value: 192.0.2.1
        (...)
        - macro: '{$MYSQL.PASSWORD}'
          type: SECRET_TEXT
          description: 'MySQL user password.'
        (...)
      valuemaps: (see table below)
  (...)
```

Element	Type	Description
host	string	(required) Unique host name.
name	string	Visible host name.
description	text	Host description.
monitored_by	string	How the host is monitored. Possible values:[1] SERVER (0, default), PROXY (1), or PROXY_GROUP (2).
proxy		Root element for proxy.
name	string	(required) Name of the proxy (if any) that monitors the host.
proxy_group		Root element for proxy group.
name	string	(required) Name of the proxy group (if any) that is used to monitor the host.
status	string	Host status. Possible values:[1] ENABLED (0, default), DISABLED (1).
ipmi_authtype	string	IPMI session authentication type. Possible values:[1] DEFAULT (-1, default), NONE (0), MD2 (1), MD5 (2), STRAIGHT (4), OEM (5), RMCP_PLUS (6).
ipmi_privilege	string	IPMI session privilege level. Possible values:[1] CALLBACK (1), USER (2, default), OPERATOR (3), ADMIN (4), OEM (5).
ipmi_username	string	Username for IPMI checks.

Element	Type	Description
ipmi_password	string	Password for IPMI checks.
templates		Root element for linked templates.
name	string	(required) Template name.
groups		Root element for host groups to which the host belongs.
name	string	(required) Host group name.
interfaces		Root element for host interfaces.
items		Root element for host items.
discovery_rules		Root element for host low-level discovery rules.
http_tests		Root element for host web scenarios.
tags		Root element for host tags.
tag	string	(required) Tag name.
value	string	Tag value.
macros		Root element for host macros.
macro	string	(required) User macro name.
type	string	User macro type. Possible values:[1] TEXT (0, default), SECRET_TEXT (1), VAULT (2).
value	string	User macro value.
description	string	User macro description.
inventory		Root element for host inventory.
<inventory_property>	string	Inventory property. All properties have their respective element (`type`, `name`, `os`, etc.; for example, see Export format).
inventory_mode	string	Inventory mode. Possible values:[1] DISABLED (-1), MANUAL (0, default), AUTOMATIC (1).
valuemaps		Root element for host value maps.

Note:
See also: Host object (refer to the relevant property with a matching name).

Host interfaces

```
zabbix_export:
  (...)
  hosts:
    - host: 'MySQL server'
      (...)
      interfaces:
        - ip: 192.0.2.1
          interface_ref: if1
      (...)
```

Element	Type	Description
default	string	Whether this is the primary host interface. Note that there can be only one primary interface of one type on a host. Possible values:[1] NO (0), YES (1, default).
type	string	Interface type. Possible values:[1] ZABBIX (1, default), SNMP (2), IPMI (3), JMX (4).
useip	string	Whether to use IP as the interface for connecting to the host (otherwise, DNS will be used). Possible values:[1] NO (0), YES (1, default).
ip	string	(required for IP connections) IP address (IPv4 or IPv6).
dns	string	(required for DNS connections) DNS name.
port	string	Port number.
details		Root element for interface details.
version	string	Use this SNMP version. Possible values:[1] SNMPV1 (1), SNMP_V2C (2, default), SNMP_V3 (3).
community	string	(required for SNMPv1 and SNMPv2 items) SNMP community.
max_repetitions	string	
retries	string	
contextname	string	SNMPv3 context name. Supported for SNMPv3 items.

Element	Type	Description
securityname	string	SNMPv3 security name. Supported for SNMPv3 items.
securitylevel	string	SNMPv3 security level. Supported for SNMPv3 items. Possible values:[1] NOAUTHNOPRIV (0, default), AUTHNOPRIV (1), AUTHPRIV (2).
authprotocol	string	SNMPv3 authentication protocol. Supported for SNMPv3 items. Possible values:[1] MD5 (0, default), SHA1 (1), SHA224 (2), SHA256 (3), SHA384 (4), SHA512 (5).
authpassphrase	string	SNMPv3 authentication passphrase. Supported for SNMPv3 items.
privprotocol	string	SNMPv3 privacy protocol. Supported for SNMPv3 items. Possible values:[1] DES (0, default), AES128 (1), AES192 (2), AES256 (3), AES192C (4), AES256C (5).
privpassphrase	string	SNMPv3 privacy passphrase. Supported for SNMPv3 items.
bulk	string	Use bulk requests for SNMP. Possible values:[1] NO (0), YES (1, default).
interface_ref	string	Interface reference name to be used in items (format: `if<N>`).

Note:
See also: Host interface object (refer to the relevant property with a matching name).

Host items

```
zabbix_export:
  (...)
  hosts:
    - host: 'MySQL server'
      (...)
      items:
        (...)
        - name: 'Binlog cache disk use'
          type: DEPENDENT
          key: mysql.binlog_cache_disk_use
          value_type: FLOAT
          description: 'Number of transactions that used a temporary disk cache because they could not fit
          preprocessing: (see table below)
          master_item:
            key: 'mysql.get_status_variables["{$MYSQL.DSN}","{$MYSQL.USER}","{$MYSQL.PASSWORD}"]'
          tags:
            - tag: component
              value: cache
        (...)
        - name: 'Buffer pool utilization'
          type: CALCULATED
          key: mysql.buffer_pool_utilization
          value_type: FLOAT
          units: '%'
          params: |
            ( last(//mysql.innodb_buffer_pool_pages_total) -
            last(//mysql.innodb_buffer_pool_pages_free) ) /
            ( last(//mysql.innodb_buffer_pool_pages_total) +
            ( last(//mysql.innodb_buffer_pool_pages_total) = 0 ) ) * 100 *
            ( last(//mysql.innodb_buffer_pool_pages_total) > 0 )
          description: 'Ratio of used to total pages in the buffer pool.'
          tags:
            - tag: component
              value: memory
          triggers: (see table below)
```

```
      (...)
      - name: Uptime
        type: DEPENDENT
        key: mysql.uptime
        units: uptime
        description: 'Number of seconds that the server has been up.'
        preprocessing: (see table below)
        master_item:
          key: 'mysql.get_status_variables["{$MYSQL.DSN}","{$MYSQL.USER}","{$MYSQL.PASSWORD}"]'
        tags:
          - tag: component
            value: application
        triggers: (see table below)
      (...)
```

Element	Type	Description
name	string	(required) Item name.
type	string	Item type. Possible values:[1] ZABBIX_PASSIVE (0, default), TRAP (2), SIMPLE (3), INTERNAL (5), ZABBIX_ACTIVE (7), EXTERNAL (10), ODBC (11), IPMI (12), SSH (13), TELNET (14), CALCULATED (15), JMX (16), SNMP_TRAP (17), DEPENDENT (18), HTTP_AGENT (19), SNMP_AGENT (20), ITEM_TYPE_SCRIPT (21), ITEM_TYPE_BROWSER (22), ITEM_TYPE_NESTED (23), TELEMETRY_QUERY (24).
snmp_oid	string	(required for SNMP_AGENT items) SNMP object ID.
key	string	(required) Item key.
delay	string	Update interval of the item. Default: 1m. The value will always be 0 for TRAP items.
history	string	Time period (using time suffix, user macro or LLD macro) of how long the history data should be stored. Default: 31d.
trends	string	Time period (using time suffix, user macro or LLD macro) of how long the trends data should be stored. Default: 365d.
status	string	Item status. Possible values:[1] ENABLED (0, default), DISABLED (1).
value_type	string	Received value type. Possible values:[1] FLOAT (0), CHAR (1), LOG (2), UNSIGNED (3, default), TEXT (4), BINARY (5), JSON (6).
allowed_hosts	string	List of comma-delimited IP addresses of hosts allowed to send data for the item. Supported for TRAP and HTTP_AGENT items.
units	string	Received value units (bps, B, etc.).
params	text	Additional parameters depending on the type of the item (executed script for SSH and TELNET items; SQL query for ODBC items; formula for CALCULATED items; the script for ITEM_TYPE_SCRIPT and ITEM_TYPE_BROWSER items).
ipmi_sensor	string	IPMI sensor. Supported for IPMI items.
authtype	string	Authentication type. Supported for SSH and HTTP_AGENT items. Possible values for SSH items:[1] PASSWORD (0, default), PUBLIC_KEY (1). Possible values for HTTP_AGENT items:[1] NONE (0, default), BASIC (1), NTLM (2), Kerberos (3), or Digest (4).
username	string	(required for SSH and TELNET items) Username for authentication. Supported for SIMPLE, ODBC, JMX and HTTP_AGENT items. When used for JMX items, `password` (see below) should also be specified or both elements should be left blank.
password	string	(required for SSH and TELNET items) Password for authentication. Supported for SIMPLE, ODBC, JMX and HTTP_AGENT items. When used for JMX items, `username` (see above) should also be specified or both elements should be left blank.
publickey	string	(required for SSH items) Name of the public key file.
privatekey	string	(required for SSH items) Name of the private key file.
description	text	Item description.

Element	Type	Description
inventory_link	string	Host inventory field that is populated by the item. Possible values:[1] NONE (0), ALIAS (4), etc. (see Host inventory for supported fields).
valuemap		Root element for item value maps.
name	string	(required) Name of the value map to use for the item.
logtimefmt	string	Format of the time in log entries. Supported for items of LOG value type.
preprocessing		Root element for item value preprocessing.
step		Root element for host item value preprocessing steps.
interface_ref	string	Reference to the host interface (format: `if<N>`).
jmx_endpoint	string	JMX endpoint. Supported for JMX items.
master_item		(required for DEPENDENT items) Root element for dependent item's master item.
key	string	(required) Dependent item's master item key.
timeout	string	Item data polling request timeout. Supported for the Timeouts list of item types.
url	string	(required for HTTP_AGENT items) URL string.
query_fields		Root element for query parameters. Supported for HTTP_AGENT items.
name	string	(required for HTTP_AGENT items) Query parameter name.
value	string	Query parameter value. Supported for HTTP_AGENT items.
parameters		Root element for user-defined parameters. Supported for ITEM_TYPE_SCRIPT and ITEM_TYPE_BROWSER items.
name	string	(required for ITEM_TYPE_SCRIPT and ITEM_TYPE_BROWSER items) User-defined parameter name.
value	string	User-defined parameter value. Supported for ITEM_TYPE_SCRIPT and ITEM_TYPE_BROWSER items.
posts	string	HTTP(S) request body data. Supported for HTTP_AGENT items.
status_codes	string	Ranges of required HTTP status codes, separated by commas. Supported for HTTP_AGENT items.
follow_redirects	string	Follow response redirects while polling data. Supported for HTTP_AGENT items. Possible values:[1] NO (0), YES (1, default).
post_type	string	Type of post data body. Supported for HTTP_AGENT items. Possible values:[1] RAW (0, default), JSON (2), XML (3).
http_proxy	string	HTTP(S) proxy connection string. Supported for HTTP_AGENT items.
headers		Root element for HTTP(S) request headers. Supported for HTTP_AGENT items.
name	string	(required for HTTP_AGENT items) Header name.
value	string	(required for HTTP_AGENT items) Header value.
retrieve_mode	string	What part of response should be stored. Supported for HTTP_AGENT items. Possible values:[1] BODY (0, default), HEADERS (1), BOTH (2).
request_method	string	Request method type. Supported for HTTP_AGENT items. Possible values:[1] GET (0, default), POST (1), PUT (2), HEAD (3).
output_format	string	How to process response. Supported for HTTP_AGENT items. Possible values:[1] RAW (0, default), JSON (1).
allow_traps	string	Allow the value to be populated in the same way as for a trapper item. Supported for HTTP_AGENT items. Possible values:[1] NO (0, default), YES (1).
ssl_cert_file	string	Public SSL Key file path. Supported for HTTP_AGENT items.
ssl_key_file	string	Private SSL Key file path. Supported for HTTP_AGENT items.
ssl_key_password	string	Password for SSL Key file. Supported for HTTP_AGENT items.

Element	Type	Description
verify_peer	string	Whether to validate that the host's certificate is authentic. Supported for HTTP_AGENT items. Possible values:[1] NO (0, default), YES (1).
verify_host	string	Whether to validate that the host name for the connection matches the one in the host's certificate. Supported for HTTP_AGENT items. Possible values:[1] NO (0, default), YES (1).
time_shift	string	How far back (before the current time) ends the time range that the item queries. Supported for TELEMETRY_QUERY items.
lookback_limit	string	How far back (before the current time) the item queries for data when there is no previous query to continue from (e.g., on the item's first run), or when the previous query is older than `delay` (e.g., because Zabbix server was down). Supported for TELEMETRY_QUERY items.
granularity	string	The length of a piece that the query's full time range is split into. Supported for TELEMETRY_QUERY items.
query		(required for TELEMETRY_QUERY items) Root element for the telemetry query. Defines what the item queries from the data source, how it filters the data, and how it aggregates the result.
signal_type	string	(required) Signal type that the item queries. Supported for TELEMETRY_QUERY items. Possible values:[1] TRACES (0), METRICS (1), LOGS (2).
metric_point_type	string	(required for METRICS signal type) Metric point type that the item queries. Supported for TELEMETRY_QUERY items. Possible values:[1] SUM (0), GAUGE (1), HISTOGRAM (2), EXPONENTIAL_HISTOGRAM (3).
columns		(required) Root element for telemetry data columns that the item queries. This element has the following elements: `column (string)` - (required) Name of the telemetry data column that the item queries. `attribute_key (string)` - (required for columns that hold key-value pairs) Name of the key that the item queries, within the column that holds key-value pairs.
aggregated_columns		(required) Root element for telemetry data aggregated columns that the item queries. This element has the following elements: `column (string)` - (required for MIN, MAX, AVG, SUM, and PERCENTILE functions) Name of the telemetry data column that the item queries. `function (string)` - (required) Function that the item applies to the column records. Possible values:[1] MIN (1), MAX (2), AVG (3), COUNT (4), SUM (5), PERCENTILE (8). `parameters (array)` - (required for PERCENTILE functions) Percentile that the item calculates from the column records. `alias (string)` - (required) Unique name for the calculated value, used to identify it in the item value.

Element	Type	Description
filter		(required) Root element for telemetry query filter for column records in the query. This element has the following elements: `evaltype` (`string`) - (required) Filter condition evaluation method. Possible values:[1] AND_OR (0), AND (1), OR (2), FORMULA (3). `formula` (`string`) - (required for FORMULA evaluation methods) User-defined expression to be used for evaluating conditions of filters with a custom expression. `conditions` (`string`) - (required) Root element for telemetry data filter conditions a record in a telemetry data column must meet. The `conditions` element has the following elements: `column` (`string`) - (required) Name of the telemetry data column that the filter conditions apply to. `attribute_key` (`string`) - (required for columns that hold key-value pairs) Name of the key that the condition applies to, within the column that holds key-value pairs. `operator` (`string`) - (required) Filter condition operator. Possible values for single-value columns:[1] EQUAL (0), NOT_EQUAL (1), LIKE (2), NOT_LIKE (3). Possible values for columns that hold key-value pairs:[1] EQUAL (0), NOT_EQUAL (1), EXISTS (12). `formulaid` (`string`) - (required for FORMULA evaluation methods) Arbitrary unique ID that is used to reference the condition from a custom expression. `value` (`string`) - (required for EQUAL, NOT_EQUAL, LIKE, and NOT LIKE operations) Value to compare with.
tags		Root element for item tags.
tag	string	(required) Tag name.
value	string	Tag value.
triggers		Root element for host item triggers.

Note:
See also: Item object (refer to the relevant property with a matching name).

Host item value preprocessing steps

```
zabbix_export:
  (...)
  hosts:
    - host: 'MySQL server'
      (...)
      items:
        (...)
        - name: 'Binlog cache disk use'
          (...)
          preprocessing:
            - type: JSONPATH
              parameters:
                - $.Binlog_cache_disk_use
            - type: DISCARD_UNCHANGED_HEARTBEAT
              parameters:
                - 6h
          (...)
```

Element	Type	Description
type	string	(required) The item value preprocessing step type. Possible values:[1] MULTIPLIER (1), RTRIM (2), LTRIM (3), TRIM (4), REGEX (5), BOOL_TO_DECIMAL (6), OCTAL_TO_DECIMAL (7), HEX_TO_DECIMAL (8), SIMPLE_CHANGE (9, calculated: received value - previous value), CHANGE_PER_SECOND (10, calculated: (received value - previous value)/(time now - time of last check)), XMLPATH (11), JSONPATH (12), IN_RANGE (13), MATCHES_REGEX (14), NOT_MATCHES_REGEX (15), CHECK_JSON_ERROR (16), CHECK_XML_ERROR (17), CHECK_REGEX_ERROR (18), DISCARD_UNCHANGED (19), DISCARD_UNCHANGED_HEARTBEAT (20), JAVASCRIPT (21), PROMETHEUS_PATTERN (22), PROMETHEUS_TO_JSON (23), CSV_TO_JSON (24), STR_REPLACE (25), CHECK_NOT_SUPPORTED (26), XML_TO_JSON (27), SNMP_WALK_VALUE (28), SNMP_WALK_TO_JSON (29), SNMP_GET_VALUE (30).
parameters		(required) Root element for parameters of the item value preprocessing step.
parameter	string	Individual parameter of the item value preprocessing step.
error_handler	string	Action type used in case of preprocessing step failure. Possible values:[1] ORIGINAL_ERROR (0, default), DISCARD_VALUE (1), CUSTOM_VALUE (2), CUSTOM_ERROR (3).
error_handler_params	string	Error handler parameters.

Note:
See also: Item preprocessing object (refer to the relevant property with a matching name).

Host item triggers

```
zabbix_export:
  (...)
  hosts:
    - host: 'MySQL server'
      (...)
      items:
        (...)
        - name: Uptime
          (...)
          triggers:
            - expression: 'nodata(/MySQL server/mysql.uptime,30m)=1'
              name: 'MySQL: Failed to fetch info data'
              event_name: 'MySQL: Failed to fetch info data (or no data for 30m)'
              priority: INFO
              description: 'Zabbix has not received any data for items for the last 30 minutes.'
              dependencies:
                - name: 'MySQL: Service is down'
                  expression: 'last(/MySQL server/mysql.ping["{$MYSQL.DSN}","{$MYSQL.USER}","{$MYSQL.PASSW
              tags:
                - tag: scope
                  value: availability
            - expression: 'last(/MySQL server/mysql.uptime)<10m'
              name: 'MySQL: Service has been restarted'
              event_name: 'MySQL: Service has been restarted (uptime < 10m)'
              priority: INFO
              description: 'MySQL uptime is less than 10 minutes.'
              tags:
                - tag: scope
                  value: notice
        (...)
```

Element	Type	Description
uuid	string	(required) Unique identifier for this trigger.
expression	string	(required) Trigger expression.
recovery_mode	string	Basis for generating OK events. Possible values:[1] EXPRESSION (0, default), RECOVERY_EXPRESSION (1), NONE (2).
recovery_expression	string	Trigger recovery expression.

Element	Type	Description
correlation_mode	string	Correlation mode (no event correlation or event correlation by tag). Possible values:[1] DISABLED (0, default), TAG_VALUE (1).
correlation_tag	string	The tag name to be used for event correlation.
name	string	(required) Trigger name.
event_name	string	Event name.
opdata	string	Operational data.
url_name	string	Label for the URL associated with the trigger.
url	string	URL associated with the trigger.
status	string	Trigger status. Possible values:[1] ENABLED (0, default), DISABLED (1).
priority	string	Trigger severity. Possible values:[1] NOT_CLASSIFIED (0, default), INFO (1), WARNING (2), AVERAGE (3), HIGH (4), DISASTER (5).
description	text	Trigger description.
type	string	Event generation type (single problem event or multiple problem events). Possible values:[1] SINGLE (0, default), MULTIPLE (1).
manual_close	string	Manual closing of problem events. Possible values:[1] NO (0, default), YES (1).
dependencies		Root element for dependencies.
name	string	(required) Dependency trigger name.
expression	string	(required) Dependency trigger expression.
recovery_expression	string	Dependency trigger recovery expression.
tags		Root element for trigger tags.
tag	string	(required) Tag name.
value	string	Tag value.

Note:
See also: Trigger object (refer to the relevant property with a matching name).

Host low-level discovery rules

```
zabbix_export:
  (...)
  hosts:
    - host: 'MySQL server'
      (...)
      discovery_rules:
        - name: 'Database discovery'
          key: 'mysql.db.discovery["{$MYSQL.DSN}","{$MYSQL.USER}","{$MYSQL.PASSWORD}"]'
          delay: 1h
          filter: (see table below)
          description: 'Scanning databases in DBMS.'
          interface_ref: if1
          item_prototypes:
            - name: 'Size of database {#DATABASE}'
              key: 'mysql.db.size["{$MYSQL.DSN}","{$MYSQL.USER}","{$MYSQL.PASSWORD}","{#DATABASE}"]'
              delay: 5m
              units: B
              description: 'Database size.'
              preprocessing:
                - type: DISCARD_UNCHANGED_HEARTBEAT
                  parameters:
                    - 1h
              interface_ref: if1
              tags:
                - tag: component
                  value: storage
                - tag: database
                  value: '{#DATABASE}'
          lld_macro_paths:
            - lld_macro: '{#DATABASE}'
```

```
            path: $.Database
          preprocessing:
            - type: DISCARD_UNCHANGED_HEARTBEAT
              parameters:
                - 1d
        (...)
```

Attention:
Most host low-level discovery rule elements are the same as for host items. The table below describes those elements that differ from host items.

Element	Type	Description
type	string	Item type. Possible values:[1] ZABBIX_PASSIVE (0, default), TRAP (2), SIMPLE (3), INTERNAL (5), ZABBIX_ACTIVE (7), EXTERNAL (10), ODBC (11), IPMI (12), SSH (13), TELNET (14), JMX (16), DEPENDENT (18), HTTP_AGENT (19), SNMP_AGENT (20), ITEM_TYPE_SCRIPT (21), ITEM_TYPE_BROWSER (22).
key	string	(required) The low-level discovery rule key.
filter		Root element for host low-level discovery rule filters.
lifetime	string	Time period (using seconds, time suffix, or user macro), after which no longer discovered resources will be deleted. Default: 7d.
lifetime_type	string	Scenario for deleting lost LLD resources. Possible values: DELETE_NEVER, DELETE_IMMEDIATELY, DELETE_AFTER.
enabled_lifetime	string	Time period (using seconds, time suffix, or user macro), after which no longer discovered resources will be disabled.
enabled_lifetime_type	string	Scenario for disabling lost LLD resources. Possible values: DISABLE_NEVER, DISABLE_IMMEDIATELY, DISABLE_AFTER.
item_prototypes		Root element for host item prototype elements, which are the same as for host items.
trigger_prototypes		Root element for host trigger prototype elements, which are the same as for host item triggers.
graph_prototypes		Root element for host graph prototypes, which are the same as for host graphs.
host_prototypes		Root element for host prototypes, which are the same as for hosts.
parent_discovery_rule		Root element for the parent low-level discovery rule (or rule prototype) of the low-level discovery rule prototype. This property denotes that it is an LLD rule prototype, direct child of the referenced rule/rule prototype.
key	string	(required) The parent low-level discovery rule (or rule prototype) key.
master_item	string	(required for DEPENDENT rules) Root element for the dependent rule's master item.
lld_macro_paths		Root element for low-level discovery rule macro paths.
lld_macro	string	(required) Low-level discovery rule macro name.
path	string	(required) Selector for value, which will be assigned to the corresponding macro.
preprocessing		Root element for low-level discovery rule value preprocessing.
step		Root element for low-level discovery rule value preprocessing step elements, which are the same as for host item value preprocessing steps, except with fewer possible values. See also: LLD rule preprocessing object.
overrides		Root element for low-level discovery rule override rules.
name	string	(required) Unique override name.
step	string	(required) Unique order number of the override.
stop	string	Stop processing next overrides if matches.
filter		Root element for low-level discovery rule override rule filter elements, which are the same as for host low-level discovery rule filters.
operations		Root element for host low-level discovery rule override operations.

Note:
See also: LLD rule object (refer to the relevant property with a matching name).

Host low-level discovery rule filters

```
zabbix_export:
  (...)
  hosts:
```

```
  - host: 'MySQL server'
    (...)
    discovery_rules:
      - name: 'Database discovery'
        (...)
        filter:
          evaltype: AND
          conditions:
            - macro: '{#DATABASE}'
              value: '{$MYSQL.DBNAME.MATCHES}'
            - macro: '{#DATABASE}'
              value: '{$MYSQL.DBNAME.NOT_MATCHES}'
              operator: NOT_MATCHES_REGEX
        (...)
```

Element	Type	Description
evaltype	string	Override filter condition evaluation method. Possible values:[1] AND_OR (0, default), AND (1), OR (2), FORMULA (3).
formula	string	Custom calculation formula for filter conditions.
conditions		Root element for filter conditions.
macro	string	(required) The name of the low-level discovery macro to perform the check on.
value	string	Value to compare with.
operator	string	Condition operator. Possible values:[1] MATCHES_REGEX (8, default), NOT_MATCHES_REGEX (9).
formulaid	string	(required) Arbitrary unique ID that is used to reference a condition from the custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward.

Note:
See also: LLD rule filter object (refer to the relevant property with a matching name).

Host low-level discovery rule override operations

Element	Type	Description
operationobject	string	Object to which the operation is applied. Possible values:[1] ITEM_PROTOTYPE (0), TRIGGER_PROTOTYPE (1), GRAPH_PROTOTYPE (2), HOST_PROTOTYPE (3), DISCOVERY_PROTOTYPE (4).
operator	string	Override condition operator. Possible values:[1] EQUAL (1), NOT_EQUAL (2), LIKE (3), NOT_LIKE (4), REGEXP (5), NOT_REGEXP (6).
value	string	A regular expression or a string for the override condition operator.
status	string	Status of the object upon the override operation.
discover	string	Whether the object is added as a result of the discovery.
delay	string	Update interval set for the item prototype upon the override operation.
history	string	History storage period set for the item prototype upon the override operation.
trends	string	Trend storage period set for the item prototype upon the override operation.
severity	string	Trigger prototype severity set upon the override operation.
tags		Root element for the tags set for the object upon the override operation.
tag	string	(required) Tag name.
value	string	Tag value.
templates		Root element for the templates linked to the host prototype upon the override operation.
name	string	(required) Template name.
inventory_mode	string	Host prototype inventory mode set upon the override operation.

Note:
See also: LLD rule override operation object (refer to the relevant property with a matching name).

Host web scenarios

Element	Type	Description
uuid	string	(required) Unique identifier for this web scenario.
name	string	(required) Web scenario name.
delay	string	Frequency (using seconds, time suffix, or user macro) of executing web scenario. Default: 1m.
attempts	integer	The number of attempts for executing web scenario steps. Possible values: 1-10 (default: 1).
agent	string	Client agent. Zabbix will pretend to be the selected browser. This is useful when a website returns different content for different browsers. Default: Zabbix.
http_proxy	string	Proxy that will be used by the web scenario given as: `http://[username[:password]@]proxy.example.com[:port]`
variables		Root element for web scenario variables that may be used in scenario steps.
name	string	(required) Variable name.
value	text	(required) Variable value.
headers		Root element for HTTP headers to be sent when performing a request. Headers should be listed using the same syntax as they would appear in the HTTP protocol.
name	string	(required) Header name.
value	text	(required) Header value.
status	string	Web scenario status. Possible values:[1] ENABLED (0, default), DISABLED (1).
authentication	string	Authentication method. Possible values:[1] NONE (0, default), BASIC (1), NTLM (2), Kerberos (3), or Digest (4).
http_user	string	User name used for BASIC (HTTP), NTLM, Kerberos, or Digest authentication.
http_password	string	Password used for BASIC (HTTP), NTLM, Kerberos, or Digest authentication.
verify_peer	string	Verify the SSL certificate of the web server. Possible values:[1] NO (0, default), YES (1).
verify_host	string	Verify that the *Common Name* field or the *Subject Alternate Name* field of the web server certificate matches. Possible values:[1] NO (0, default), YES (1).
ssl_cert_file	string	Name of the SSL certificate file used for client authentication (must be in PEM format).
ssl_key_file	string	Name of the SSL private key file used for client authentication (must be in PEM format).
ssl_key_password	string	SSL private key file password.
steps		(required) Root element for host web scenario steps.
tags		Root element for web scenario tags.
tag	string	(required) Tag name.
value	string	Tag value.

Note:
See also: Web scenario object (refer to the relevant property with a matching name).

Host web scenario steps

Element	Type	Description
name	string	(required) Web scenario step name.
url	string	(required) URL for monitoring.
query_fields		Root element for query parameters (an array of HTTP fields to be added to the URL when performing a request).
name	string	(required) Query parameter name.
value	string	Query parameter value.
posts		Root element for HTTP POST variables (a string (raw post data) or an array of HTTP fields (form field data)).
name	string	(required) Post field name.
value	string	(required) Post field value.
variables		Root element of step-level variables (macros) that should be applied after this step. If the variable value has a 'regex:' prefix, then its value is extracted from the data returned by this step according to the regular expression pattern following the 'regex:' prefix
name	string	(required) Variable name.

Element	Type	Description
value	text	(required) Variable value.
headers		Root element for HTTP headers to be sent when performing a request.
name	string	(required) Header name.
value	text	(required) Header value.
follow_redirects	string	Follow HTTP redirects. Possible values:[1] NO (0), YES (1, default).
retrieve_mode	string	HTTP response retrieve mode. Possible values:[1] BODY (0, default), HEADERS (1), BOTH (2).
timeout	string	Timeout (using seconds, time suffix, or user macro) of step execution. Default: 15s.
required	string	Text that must be present in the response (ignored if empty).
status_codes	string	A comma-delimited list of accepted HTTP status codes (e.g., 200-201,210-299; ignored if empty).

Note:
See also: Web scenario step object (refer to the relevant property with a matching name).

Host graphs

```
zabbix_export:
  (...)
  graphs:
    - name: 'MySQL: Bandwidth'
      graph_items: (see table below)
    (...)
```

Element	Type	Description
uuid	string	Unique identifier of the graph.
name	string	(required) Graph name.
width	integer	Graph width, in pixels. Used for preview and for pie/exploded graphs. Possible values:[1] 20-65535 (default: 900).
height	integer	Graph height, in pixels. Used for preview and for pie/exploded graphs. Possible values:[1] 20-65535 (default: 900).
yaxismin	double	Value of Y axis minimum. Supported for FIXED minimum value of Y axis. Default: 0.
yaxismax	double	Value of Y axis maximum. Supported for FIXED minimum value of X axis. Default: 0.
show_work_period	string	Highlight non-working hours. Supported for NORMAL and STACKED graphs. Possible values:[1] NO (0), YES (1, default).
show_triggers	string	Display simple trigger values as a line. Supported for NORMAL and STACKED graphs. Possible values:[1] NO (0), YES (1, default).
type	string	Graph type. Possible values:[1] NORMAL (0, default), STACKED (1), PIE (2), EXPLODED (3).
show_legend	string	Display graph legend. Possible values:[1] NO (0), YES (1, default).
show_3d	string	Enable 3D style. Supported for NORMAL and STACKED graphs. Possible values:[1] NO (0, default), YES (1).
percent_left	double	Show the percentile line for left axis. Supported for NORMAL graphs. Default: 0.
percent_right	double	Show the percentile line for right axis. Supported for NORMAL graphs. Default: 0.

Element	Type	Description
ymin_type_1	string	Minimum value of Y axis. Supported for NORMAL and STACKED graphs. Possible values:[1] CALCULATED (0, default), FIXED (1), ITEM (2).
ymin_item_1		(required if `ymin_type_1` is set to `ITEM`) Root element for individual item details.
host	string	(required) Item host.
key	string	(required) Item key.
ymax_type_1	string	Maximum value of Y axis. Supported for NORMAL and STACKED graphs. Possible values:[1] CALCULATED (0, default), FIXED (1), ITEM (2).
ymax_item_1		(required if `ymax_type_1` is set to `ITEM`) Root element for individual item details.
host	string	(required) Item host.
key	string	(required) Item key.
graph_items		(required) Root element for host graph items.

Note:
See also: Graph object (refer to the relevant property with a matching name).

Host graph items

```
zabbix_export:
  (...)
  graphs:
    - name: 'MySQL: Bandwidth'
      graph_items:
        - drawtype: GRADIENT_LINE
          color: 199C0D
          item:
            host: 'MySQL server'
            key: mysql.bytes_received.rate
        - sortorder: '1'
          drawtype: GRADIENT_LINE
          color: F63100
          item:
            host: 'MySQL server'
            key: mysql.bytes_sent.rate
    (...)
```

Element	Type	Description
sortorder	integer	Draw order. The smaller value is drawn first. Can be used to draw lines or regions behind (or in front of) another.
drawtype	string	Draw style of the graph item. Supported for NORMAL graphs. Possible values:[1] SINGLE_LINE (0, default), FILLED_REGION (1), BOLD_LINE (2), DOTTED_LINE (3), DASHED_LINE (4), GRADIENT_LINE (5).
color	string	Element color (6 symbols, hex).
yaxisside	string	Side of the graph where the graph item's Y scale will be drawn. Supported for NORMAL and STACKED graphs. Possible values:[1] LEFT (0, default), RIGHT (1).
calc_fnc	string	Data to draw if more than one value exists for an item. Possible values:[1] MIN (1), AVG (2, default), MAX (4), ALL (7; minimum, average, and maximum; supported for simple graphs), LAST (9, supported for pie/exploded graphs).
type	string	Graph item type. Possible values:[1] SIMPLE (0, default), GRAPH_SUM (2; value of the item represents the whole pie; supported for pie/exploded graphs).
item		(required) Individual item.
host	string	(required) Item host.
key	string	(required) Item key.

Note:
See also: Graph item object (refer to the relevant property with a matching name).

Host value maps

```
zabbix_export:
  (...)
  hosts:
    - host: 'MySQL server'
      (...)
      valuemaps:
        - name: Example value map
          mappings:
            - value: '1'
              newvalue: Example value
  (...)
```

Element	Type	Description
uuid	string	(required) Unique identifier for this value map.
name	string	(required) Value map name.
mapping		Root element for mappings.
type	string	Mapping match type. Possible values:[1] EQUAL (0, default), GREATER_OR_EQUAL (2), LESS_OR_EQUAL (3), IN_RANGE (4), REGEXP (5), DEFAULT (6).
value	string	Original value.
newvalue	string	(required) Value to which the original value is mapped to.

Note:
See also: Value map object (refer to the relevant property with a matching name).

Footnotes [1] API integer values in brackets, for example, ENABLED (0), are mentioned only for reference. For more information, see the linked API object page in the table entry or at the end of each section.

Network maps

Overview

Network map export contains:

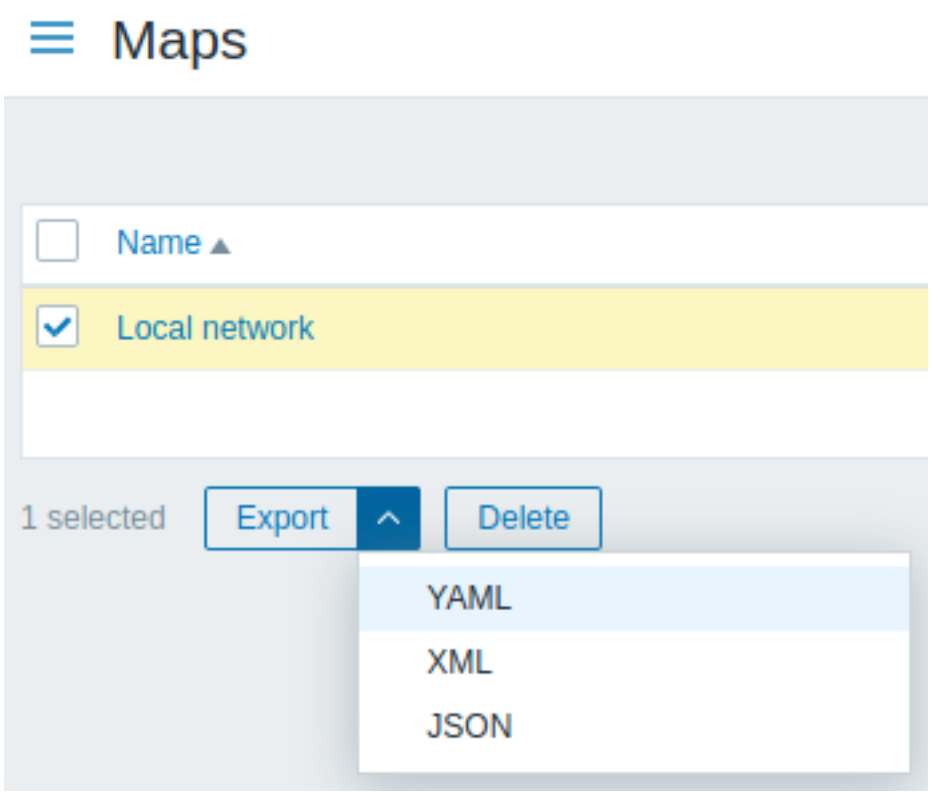
- All related images
- Map structure (all map settings, all contained elements with their settings, map links and map link status indicators)

Warning:
Any host groups, hosts, triggers, other maps or other elements that may be related to the exported map are not exported. Thus, if at least one of the elements the map refers to is missing, importing it will fail.

Exporting

To export network maps, do the following:

1. Go to *Monitoring* > *Maps*.
2. Mark the checkboxes of the network maps to export.
3. Click *Export* below the list to export in the default YAML format, or click the arrow to select a different export format.



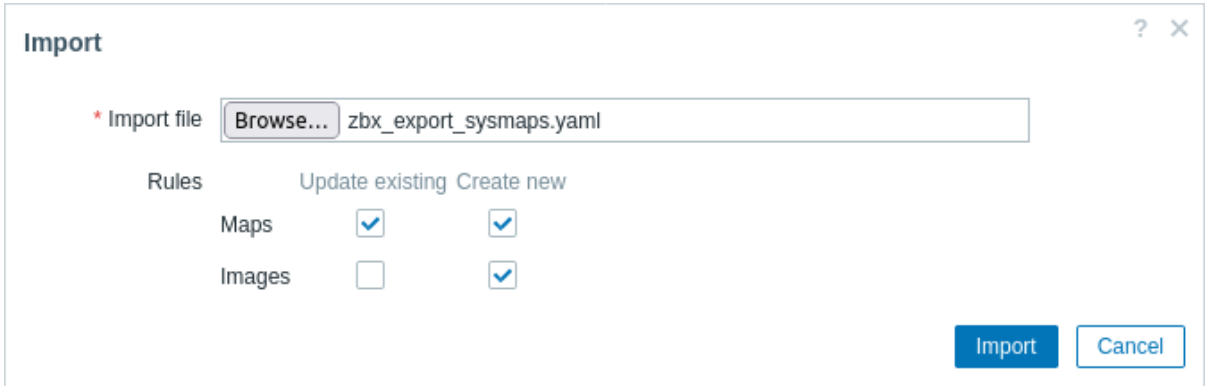
Depending on the selected format, maps are exported to a local file with a default name:

- `zabbix_export_maps.yaml` - in YAML export;
- `zabbix_export_maps.xml` - in XML export;
- `zabbix_export_maps.json` - in JSON export.

Importing

To import network maps, do the following:

1. Go to *Monitoring* > *Maps*.
2. Click *Import* in the upper-right corner.
3. Select the import file.
4. Mark the checkboxes of the import rules you want to apply (see below).
5. Click *Import* in the lower-right corner of the configuration form.



Import rules:

Rule	Description
Update existing	Existing maps will be updated using data from the import file. Otherwise, they will not be updated.
Create new	New maps will be created using data from the import file. Otherwise, they will not be created.

If you uncheck both map options and check the respective options for images, images only will be imported. Image importing is only available to *Super admin* users.

You cannot review changes to network maps during the import. A success or failure message of the import will be displayed in the frontend.

Warning:
If replacing an existing image, it will affect all maps that are using this image.

Export format

Export to YAML:

```yaml
zabbix_export:
  version: '8.2'
  images:
    - name: Zabbix_server_3D_(128)
      imagetype: '1'
      encodedImage: iVBOR...5CYII=
  maps:
    - name: 'Local network'
      width: '680'
      height: '200'
      label_type: '0'
      label_location: '0'
      highlight: '1'
      expandproblem: '1'
      markelements: '1'
      show_unack: '0'
      severity_min: '0'
      show_suppressed: '0'
      grid_size: '50'
      grid_show: '1'
      grid_align: '1'
      label_format: '0'
      label_type_host: '2'
      label_type_hostgroup: '2'
      label_type_trigger: '2'
      label_type_map: '2'
      label_type_image: '2'
      label_string_host: ''
      label_string_hostgroup: ''
      label_string_trigger: ''
      label_string_map: ''
      label_string_image: ''
      expand_macros: '1'
      background: {  }
      iconmap: {  }
      urls: {  }
      selements:
        - elementtype: '0'
          elements:
            - host: 'Zabbix server'
          label: |
            {HOST.NAME}
            {HOST.CONN}
          label_location: '0'
          x: '111'
          'y': '61'
          elementsubtype: '0'
          areatype: '0'
          width: '200'
          height: '200'
          viewtype: '0'
          use_iconmap: '0'
          selementid: '1'
          icon_off:
```

```
            name: Zabbix_server_3D_(128)
          icon_on: {  }
          icon_disabled: {  }
          icon_maintenance: {  }
          urls: {  }
          evaltype: '0'
      shapes:
        - type: '0'
          x: '0'
          'y': '0'
          width: '680'
          height: '15'
          text: '{MAP.NAME}'
          font: '9'
          font_size: '11'
          font_color: '000000'
          text_halign: '0'
          text_valign: '0'
          border_type: '0'
          border_width: '0'
          border_color: '000000'
          background_color: ''
          zindex: '0'
      lines: {  }
      links: {  }
```

Exported elements

Exported elements are explained in the table below.

Element	Type	Description
images		Root element for images.
name	string	Unique image name.
imagetype	integer	Image type. Possible values: 1 - Image; 2 - Background.
encodedImage	string	Base64 encoded image.
maps		Root element for maps.

Maps

Element	Type	Description
name	string	Unique map name.
width	integer	Map width, in pixels.
height	integer	Map height, in pixels.
label_type	integer	Map element label type. Possible values: 0 - Label; 1 - Host IP address; 2 - Element name; 3 - Status only; 4 - Nothing.
label_location	integer	Map element label location by default. Possible values: 0 - Bottom; 1 - Left; 2 - Right; 3 - Top.

Element	Type	Description
highlight	integer	Enable icon highlighting for active triggers and host statuses. Possible values: 0 - No; 1 - Yes.
expandproblem	integer	Display problem trigger for elements with a single problem. Possible values: 0 - No; 1 - Yes.
markelements	integer	Highlight map elements that have recently changed their status. Possible values: 0 - No; 1 - Yes.
show_unack	integer	Problem display. Possible values: 0 - Count of all problems; 1 - Count of unacknowledged problems; 2 - Count of acknowledged and unacknowledged problems separately.
severity_min	integer	Minimum trigger severity to show on the map by default. Possible values: 0 - Not classified; 1 - Information; 2 - Warning; 3 - Average; 4 - High; 5 - Disaster.
show_suppressed	integer	Display problems, which would otherwise be suppressed (not shown) because of host maintenance. Possible values: 0 - No; 1 - Yes.
grid_size	integer	Cell size of a map grid in pixels. Supported if `grid_show` is set to 0. Possible values: 20, 40, 50, 75 or 100.
grid_show	integer	Display a grid in map configuration. Possible values: 0 - Yes; 1 - No.
grid_align	integer	Automatically align icons in map configuration. Possible values: 0 - Yes; 1 - No.
label_format	integer	Use advanced label configuration. Possible values: 0 - No; 1 - Yes.
label_type_host	integer	Display label as host label. Supported if `label_format` is set to 1. Possible values: 0 - Label; 1 - Host IP address; 2 - Element name; 3 - Status only; 4 - Nothing; 5 - Custom label.
label_type_hostgroup	integer	Display label as host group label. Supported if `label_format` is set to 1. Possible values: 0 - Label; 2 - Element name; 3 - Status only; 4 - Nothing; 5 - Custom label.

Element	Type	Description
label_type_trigger	integer	Display label as trigger label. Supported if `label_format` is set to 1. Possible values: 0 - Label; 2 - Element name; 3 - Status only; 4 - Nothing; 5 - Custom label.
label_type_map	integer	Display label as map label. Supported if `label_format` is set to 1. Possible values: 0 - Label; 2 - Element name; 3 - Status only; 4 - Nothing; 5 - Custom label.
label_type_image	integer	Display label as image label. Supported if `label_format` is set to 1. Possible values: 0 - Label; 2 - Element name; 4 - Nothing; 5 - Custom label.
label_string_host	string	Custom label for host elements. Supported if `label_type_host` is set to 5.
label_string_hostgroup	string	Custom label for host group elements. Supported if `label_type_hostgroup` is set to 5.
label_string_trigger	string	Custom label for trigger elements. Supported if `label_type_trigger` is set to 5.
label_string_map	string	Custom label for map elements. Supported if `label_type_map` is set to 5.
label_string_image	string	Custom label for image elements. Supported if `label_type_image` is set to 5.
expand_macros	integer	Expand macros in labels in map configuration. Possible values: 0 - No; 1 - Yes.
background		Root element for background image (if any). Supported if `imagetype` is set to 2.
name	string	Background image name.
iconmap		Root element for icon mapping (if any).
name	string	Icon mapping name.
urls		Root element for URLs used by maps or each map element.
name	string	Link name.
url	string	Link URL.
elementtype	integer	Map item type the link belongs to. Possible values: 0 - Host; 1 - Map; 2 - Trigger; 3 - Host group; 4 - Image.
selements		Root element for map selements.
shapes		Root element for map shapes.
type	integer	Shape type. Possible values: 0 - Rectangle; 1 - Ellipse.
x	integer	X coordinates of the shape in pixels.
y	integer	Y coordinates of the shape in pixels.
width	integer	Shape width.

Element		Type	Description
	height	integer	Shape height.
	text	string	Text inside of shape.
	font	integer	Text font style. Possible values: 0 - Georgia, serif; 1 - ”Palatino Linotype”, ”Book Antiqua”, Palatino, serif; 2 - ”Times New Roman”, Times, serif; 3 - Arial, Helvetica, sans-serif; 4 - ”Arial Black”, Gadget, sans-serif; 5 - ”Comic Sans MS”, cursive, sans-serif; 6 - Impact, Charcoal, sans-serif; 7 - ”Lucida Sans Unicode”, ”Lucida Grande”, sans-serif; 8 - Tahoma, Geneva, sans-serif; 9 - ”Trebuchet MS”, Helvetica, sans-serif; 10 - Verdana, Geneva, sans-serif; 11 - ”Courier New”, Courier, monospace; 12 - ”Lucida Console”, Monaco, monospace.
	font_size	integer	Font size in pixels.
	font_color	string	Font color represented in hexadecimal code.
	text_halign	integer	Horizontal alignment of text. Possible values: 0 - Center; 1 - Left; 2 - Right.
	text_valign	integer	Vertical alignment of text. Possible values: 0 - Middle; 1 - Top; 2 - Bottom.
	border_type	integer	Type of the border for the shape. Possible values: 0 - None; 1 - Bold line; 2 - Dotted line; 3 - Dashed line.
	border_width	integer	Width of the border in pixels.
	border_color	string	Border color represented in hexadecimal code.
	background_color	string	Background (fill) color represented in hexadecimal code.
	zindex	integer	Value for ordering all shapes and lines (z-index).
lines			Root element for map lines.
	x1	integer	X coordinates of the line point 1 in pixels.
	y1	integer	Y coordinates of the line point 1 in pixels.
	x2	integer	X coordinates of the line point 2 in pixels.
	y2	integer	Y coordinates of the line point 2 in pixels.
	line_type	integer	Line type. Possible values: 0 - None; 1 - Bold line; 2 - Dotted line; 3 - Dashed line.
	line_width	integer	Line width in pixels.
	line_color	string	Line color represented in hexadecimal code.
	zindex	integer	Value for ordering all shapes and lines (z-index).
links			Root element for links between map elements.
	drawtype	integer	Link style. Possible values: 0 - Line; 2 - Bold line; 3 - Dotted line; 4 - Dashed line.
	color	string	Link color (6 symbols, hex).

Element	Type	Description
label	string	Link label.
selementid1	id	ID of one element to connect.
selementid2	id	ID of the other element to connect.
linktriggers		Root element for link status indicators.

Note:
See also: Map object (refer to the relevant property with a matching name).

Maps elements

Element	Type	Description
elementtype	integer	Map element type. Possible values: 0 - Host; 1 - Map; 2 - Trigger; 3 - Host group; 4 - Image.
elements		Root element for Zabbix entities (host, host group, map, etc.) that are represented on the map. All entities have their respective element (`host`, etc.; see, for example, Export format).
label	string	Icon label.
label_location	integer	Label location. Possible values: -1 - Use map default; 0 - Bottom; 1 - Left; 2 - Right; 3 - Top.
x	integer	Location on the X axis.
y	integer	Location on the Y axis.
elementsubtype	integer	Element subtype. Supported if `elementtype` is set to 3. Possible values: 0 - Single host group; 1 - All host groups.
areatype	integer	Area size. Supported if `elementtype` is set to 1. Possible values: 0 - Same as whole map; 1 - Custom size.
width	integer	Width of area. Supported if `areatype` is set to 1.
height	integer	Height of area. Supported if `areatype` is set to 1.
viewtype	integer	Area placement algorithm. Supported if `elementsubtype` is set to 1. Possible values: 0 - Place evenly in the area.
use_iconmap	integer	Use icon mapping for this element. Relevant only if icon mapping is activated on map level. Possible values: 0 - No; 1 - Yes.
selementid	id	Unique element record ID.
icon_off		Root element for image to use when element is in `OK` state.
name	string	Unique image name.
icon_on		Root element for image to use when element is in `Problem` state.
name	string	Unique image name.

Element	Type	Description
icon_disabled		Root element for image to use when element is disabled.
name	string	Unique image name.
icon_maintenance		Root element for image to use when element is in maintenance.
name	string	Unique image name.
urls		Root element for URLs used by maps or each map element.
name	string	Link name.
url	string	Link URL.
evaltype	integer	Evaluation type for tags.
tags		Root element for problem tags (for host and host group elements). If tags are given, only problems with these tags will be displayed on the map.
tag	string	Tag name.
value	string	Tag value.
operator	integer	Operator.

Note:
See also: Map element object (refer to the relevant property with a matching name).

Map link status indicators

Element	Type	Description
drawtype	integer	Link style when trigger is in the 'Problem' state. Possible values: 0 - Line; 2 - Bold line; 3 - Dotted line; 4 - Dashed line.
color	string	Link color (6 symbols, hex) when trigger is in the 'Problem' state.
trigger		Root element for trigger used to indicate link status.
description	string	Trigger name.
expression	string	Trigger expression.
recovery_expression	string	Trigger recovery expression.

Note:
See also: Map link trigger object (refer to the relevant property with a matching name).

Media types

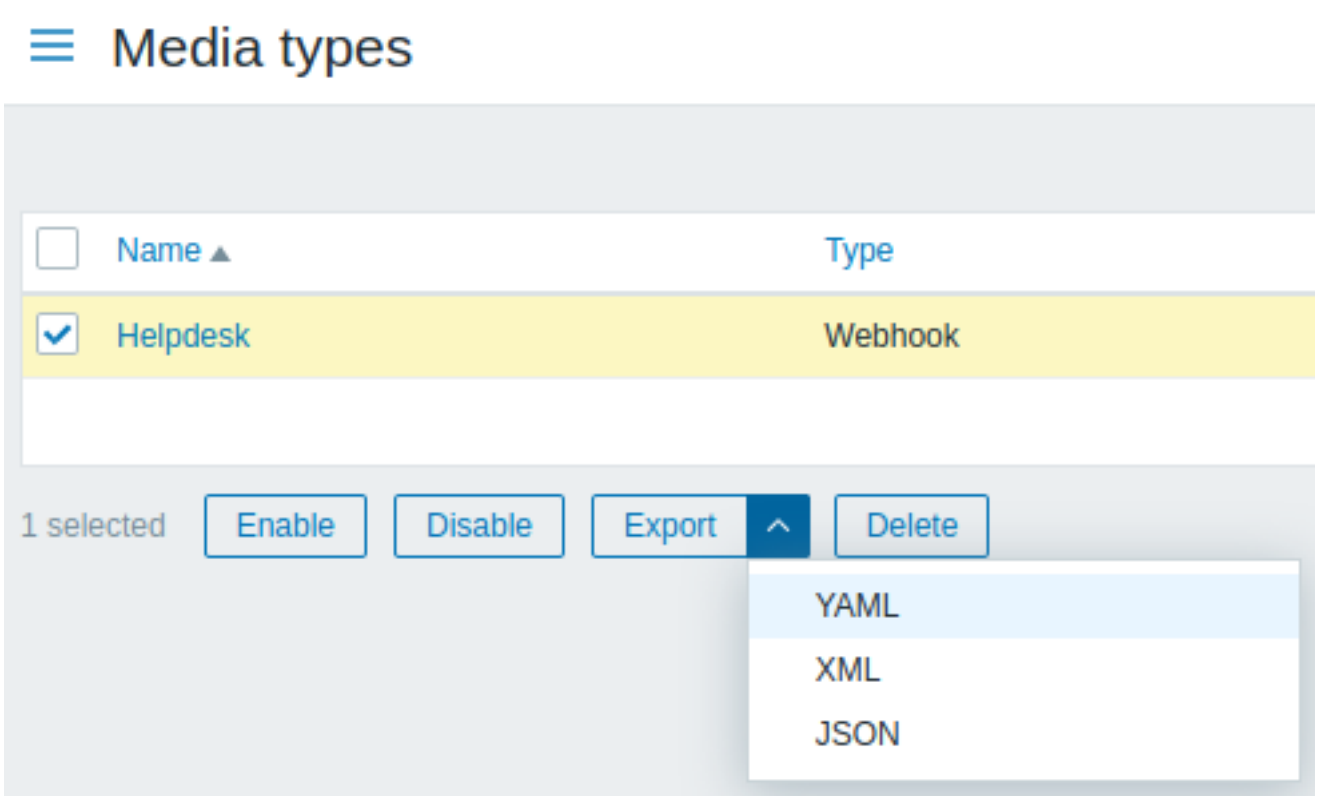
Overview

Media types are exported with all related objects and object relations.

Exporting

To export media types, do the following:

1. Go to *Alerts* > *Media types*.
2. Mark the checkboxes of the media types to export.
3. Click *Export* below the list to export in the default YAML format, or click the arrow to select a different export format.



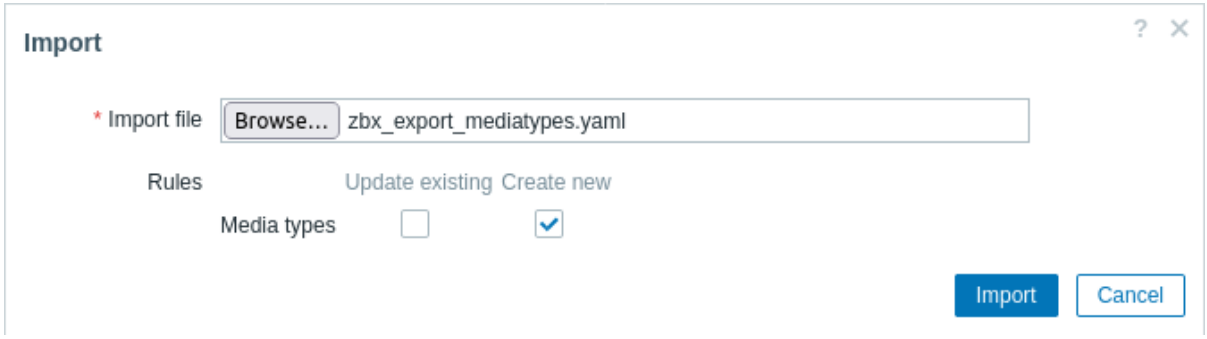
Depending on the selected format, media types are exported to a local file with a default name:

- `zabbix_export_mediatypes.yaml` - in YAML export;
- `zabbix_export_mediatypes.xml` - in XML export;
- `zabbix_export_mediatypes.json` - in JSON export.

Importing

To import media types, do the following:

1. Go to *Alerts* > *Media types*.
2. Click *Import* in the upper-right corner.
3. Select the import file.
4. Mark the checkboxes of the import rules you want to apply (see below).
5. Click *Import* in the lower-right corner of the configuration form.



Import rules:

Rule	Description
Update existing	Mark this checkbox to update existing media types using data from the import file.
Create new	Mark this checkbox to create new media types using data from the import file. If the import file contains a media type that already exists (e.g., *Email*), a new one will not be created, and the existing media type will also not be updated.

You cannot review changes to media types during the import. A success or failure message of the import will be displayed in the frontend.

Export format

Export to YAML:

```yaml
zabbix_export:
  version: '8.2'
  media_types:
    - name: Pushover
      type: WEBHOOK
      parameters:
        - name: endpoint
          value: 'https://api.pushover.net/1/messages.json'
        - name: eventid
          value: '{EVENT.ID}'
        - name: event_nseverity
          value: '{EVENT.NSEVERITY}'
        - name: event_source
          value: '{EVENT.SOURCE}'
        - name: event_value
          value: '{EVENT.VALUE}'
        - name: expire
          value: '1200'
        - name: message
          value: '{ALERT.MESSAGE}'
        - name: priority_average
          value: '0'
        - name: priority_default
          value: '0'
        - name: priority_disaster
          value: '0'
        - name: priority_high
          value: '0'
        - name: priority_information
          value: '0'
        - name: priority_not_classified
          value: '0'
        - name: priority_warning
          value: '0'
        - name: retry
          value: '60'
        - name: title
          value: '{ALERT.SUBJECT}'
        - name: token
          value: '<PUSHOVER TOKEN HERE>'
        - name: triggerid
          value: '{TRIGGER.ID}'
        - name: url
          value: '{$ZABBIX.URL}'
        - name: url_title
          value: Zabbix
        - name: user
          value: '{ALERT.SENDTO}'
      status: DISABLED
      max_sessions: '0'
      script: |
        try {
            var params = JSON.parse(value),
                request = new HttpRequest(),
                data,
                response,
                severities = [
                    {name: 'not_classified', color: '#97AAB3'},
                    {name: 'information', color: '#7499FF'},
                    {name: 'warning', color: '#FFC859'},
                    {name: 'average', color: '#FFA059'},
```

```
            {name: 'high', color: '#E97659'},
            {name: 'disaster', color: '#E45959'},
            {name: 'resolved', color: '#009900'},
            {name: 'default', color: '#000000'}
        ],
        priority;

    if (typeof params.HTTPProxy === 'string' && params.HTTPProxy.trim() !== '') {
        request.setProxy(params.HTTPProxy);
    }

    if ([0, 1, 2, 3].indexOf(parseInt(params.event_source)) === -1) {
        throw 'Incorrect "event_source" parameter given: "' + params.event_source + '".\nMust be 0
    }

    if (params.event_value !== '0' && params.event_value !== '1'
        && (params.event_source === '0' || params.event_source === '3')) {
        throw 'Incorrect "event_value" parameter given: ' + params.event_value + '\nMust be 0 or 1
    }

    if ([0, 1, 2, 3, 4, 5].indexOf(parseInt(params.event_nseverity)) === -1) {
        params.event_nseverity = '7';
    }

    if (params.event_value === '0') {
        params.event_nseverity = '6';
    }

    priority = params['priority_' + severities[params.event_nseverity].name] || params.priority_de

    if (isNaN(priority) || priority < -2 || priority > 2) {
        throw '"priority" should be -2..2';
    }

    if (params.event_source === '0' && isNaN(params.triggerid)) {
        throw 'field "triggerid" is not a number';
    }

    if (isNaN(params.eventid)) {
        throw 'field "eventid" is not a number';
    }

    if (typeof params.message !== 'string' || params.message.trim() === '') {
        throw 'field "message" cannot be empty';
    }

    data = {
        token: params.token,
        user: params.user,
        title: params.title,
        message: params.message,
        url: (params.event_source === '0')
            ? params.url + '/tr_events.php?triggerid=' + params.triggerid + '&eventid=' + params.e
            : params.url,
        url_title: params.url_title,
        priority: priority
    };

    if (priority == 2) {
        if (isNaN(params.retry) || params.retry < 30) {
            throw 'field "retry" should be a number with value of at least 30 if "priority" is set
        }
```

```
            if (isNaN(params.expire) || params.expire > 10800) {
                throw 'field "expire" should be a number with value of at most 10800 if "priority" is
            }

            data.retry = params.retry;
            data.expire = params.expire;
        }

        data = JSON.stringify(data);
        Zabbix.log(4, '[ Pushover Webhook ] Sending request: ' + params.endpoint + '\n' + data);

        request.addHeader('Content-Type: application/json');
        response = request.post(params.endpoint, data);

        Zabbix.log(4, '[ Pushover Webhook ] Received response with status code ' + request.getStatus()

        if (response !== null) {
            try {
                response = JSON.parse(response);
            }
            catch (error) {
                Zabbix.log(4, '[ Pushover Webhook ] Failed to parse response received from Pushover');
                response = null;
            }
        }

        if (request.getStatus() != 200 || response === null || typeof response !== 'object' || respons
            if (response !== null && typeof response === 'object' && typeof response.errors === 'objec
                    && typeof response.errors[0] === 'string') {
                throw response.errors[0];
            }
            else {
                throw 'Unknown error. Check debug log for more information.';
            }
        }

        return 'OK';
    }
    catch (error) {
        Zabbix.log(4, '[ Pushover Webhook ] Pushover notification failed: ' + error);
        throw 'Pushover notification failed: ' + error;
    }
description: |
    Please refer to setup guide here: https://git.zabbix.com/projects/ZBX/repos/zabbix/browse/template

    Set token parameter with to your Pushover application key.
    When assigning Pushover media to the Zabbix user - add user key into send to field.
message_templates:
    - event_source: TRIGGERS
      operation_mode: PROBLEM
      subject: 'Problem: {EVENT.NAME}'
      message: |
        Problem started at {EVENT.TIME} on {EVENT.DATE}
        Problem name: {EVENT.NAME}
        Host: {HOST.NAME}
        Severity: {EVENT.SEVERITY}
        Operational data: {EVENT.OPDATA}
        Original problem ID: {EVENT.ID}
        {TRIGGER.URL}
    - event_source: TRIGGERS
      operation_mode: RECOVERY
      subject: 'Resolved in {EVENT.DURATION}: {EVENT.NAME}'
```

```
      message: |
        Problem has been resolved at {EVENT.RECOVERY.TIME} on {EVENT.RECOVERY.DATE}
        Problem name: {EVENT.NAME}
        Problem duration: {EVENT.DURATION}
        Host: {HOST.NAME}
        Severity: {EVENT.SEVERITY}
        Original problem ID: {EVENT.ID}
        {TRIGGER.URL}
    - event_source: TRIGGERS
      operation_mode: UPDATE
      subject: 'Updated problem in {EVENT.AGE}: {EVENT.NAME}'
      message: |
        {USER.FULLNAME} {EVENT.UPDATE.ACTION} problem at {EVENT.UPDATE.DATE} {EVENT.UPDATE.TIME}.
        {EVENT.UPDATE.MESSAGE}

        Current problem state is {EVENT.STATUS}, age is {EVENT.AGE}, acknowledged: {EVENT.ACK.STATUS}.
    - event_source: DISCOVERY
      operation_mode: PROBLEM
      subject: 'Discovery: {DISCOVERY.DEVICE.STATUS} {DISCOVERY.DEVICE.IPADDRESS}'
      message: |
        Discovery rule: {DISCOVERY.RULE.NAME}

        Device IP: {DISCOVERY.DEVICE.IPADDRESS}
        Device DNS: {DISCOVERY.DEVICE.DNS}
        Device status: {DISCOVERY.DEVICE.STATUS}
        Device uptime: {DISCOVERY.DEVICE.UPTIME}

        Device service name: {DISCOVERY.SERVICE.NAME}
        Device service port: {DISCOVERY.SERVICE.PORT}
        Device service status: {DISCOVERY.SERVICE.STATUS}
        Device service uptime: {DISCOVERY.SERVICE.UPTIME}
    - event_source: AUTOREGISTRATION
      operation_mode: PROBLEM
      subject: 'Autoregistration: {HOST.HOST}'
      message: |
        Host name: {HOST.HOST}
        Host IP: {HOST.IP}
        Agent port: {HOST.PORT}
```

Exported elements

Exported elements are explained in the table below.

Element	Type	Description
name	string	(required) Media type name.
type	string	(required) Transport used by the media type. Possible values:[1] EMAIL (0), SMS (1), SCRIPT (2), WEBHOOK (4).
status	string	Whether the media type is enabled. Possible values:[1] ENABLED (0, default), DISABLED (1)
max_sessions	integer	The maximum number of alerts that can be processed in parallel. Possible values for SMS:[1] 1 (default). Possible values for other media types:[1] 0-100 (where 0 - unlimited).
attempts	integer	The maximum number of attempts to send an alert. Possible values:[1] 1-10 (default: 3).
attempt_interval	string	The interval (using seconds or time suffix) between retry attempts. Possible values:[1] 0-60s (default: 10s).
description	string	Media type description.
message_templates		Root element for media type message templates.
event_source	string	(required) Event source. Possible values:[1] TRIGGERS (0), DISCOVERY (1), AUTOREGISTRATION (2), INTERNAL (3), SERVICE (4).
operation_mode	string	Operation mode. Possible values:[1] PROBLEM (0), RECOVERY (1), UPDATE (2).

Element	Type	Description
subject	string	Message subject.
message	string	Message body.

Note:
See also: Media type object (refer to the relevant property with a matching name).

Email

The following additional elements are exported only for the *Email* media type.

Element	Type	Description
provider	string	Email provider.
smtp_server	string	SMTP server.
smtp_port	integer	SMTP server port to connect to. Default: 25.
smtp_helo	string	SMTP helo.
smtp_email	string	Email address from which notifications will be sent.
smtp_security	string	SMTP connection security level to use. Possible values:[1] NONE (0, default), STARTTLS (1), SSL_OR_TLS (2).
smtp_verify_host	string	SSL verify host for SMTP. Possible values:[1] NO (0, default), YES (1).
smtp_verify_peer	string	SSL verify peer for SMTP. Possible values:[1] NO (0, default), YES (1).
smtp_authentication	string	SMTP authentication method to use. Possible values:[1] NONE (0, default), PASSWORD (1), OAUTH (2).
username	string	Username.
password	string	Authentication password.
redirection_url	string	Zabbix frontend URL to redirect back OAuth authorization.
client_id	string	The client identifier registered within the OAuth authorization server.
authorization_url	string	OAuth URL, with parameters, to get access and refresh tokens.
token_url	string	OAuth URL to exchange authorization token to access and refresh tokens.
message_format	string	Message format. Possible values:[1] TEXT (0), HTML (1, default).

Note:
See also: Media type object (refer to the relevant property with a matching name).

SMS

The following additional elements are exported only for the *SMS* media type.

Element	Type	Description
gsm_modem	string	(required) Serial device name of the GSM modem.

Note:
See also: Media type object (refer to the relevant property with a matching name).

Script

The following additional elements are exported only for the *Script* media type.

Element	Type	Description
script name	string	(required) Script name.
parameters		Root element for script parameters.

Element	Type	Description
sortorder	string	(required) Order of the script parameters passed to the script as command-line arguments.
value	string	Script parameter value.

Note:
See also: Media type object (refer to the relevant property with a matching name).

Webhook

The following additional elements are exported only for the *Webhook* media type.

Element	Type	Description
script	string	Script.
timeout	string	Javascript script HTTP request timeout interval. Possible values:[1] 1-60s (default: 30s).
process_tags	string	Whether to process returned tags. Possible values:[1] NO (0, default), YES (1).
show_event_menu	string	Indicates the presence of an entry in the event menu if the {EVENT.TAGS.*} macro was successfully resolved in `event_menu_url` and `event_menu_name` fields. Possible values:[1] NO (0, default), YES (1).
event_menu_url	string	URL of the event menu entry. Supports {EVENT.TAGS.*} macro.
event_menu_name	string	Name of the event menu entry. Supports {EVENT.TAGS.*} macro.
parameters		Root element for webhook media type parameters.
name	string	(required) Webhook parameter name.
value	string	Webhook parameter value.

Note:
See also: Media type object (refer to the relevant property with a matching name).

Footnotes

[1] API integer values in brackets, for example, ENABLED (0), are mentioned only for reference. For more information, see the linked API object page in the table entry or at the end of each section.

Dashboards

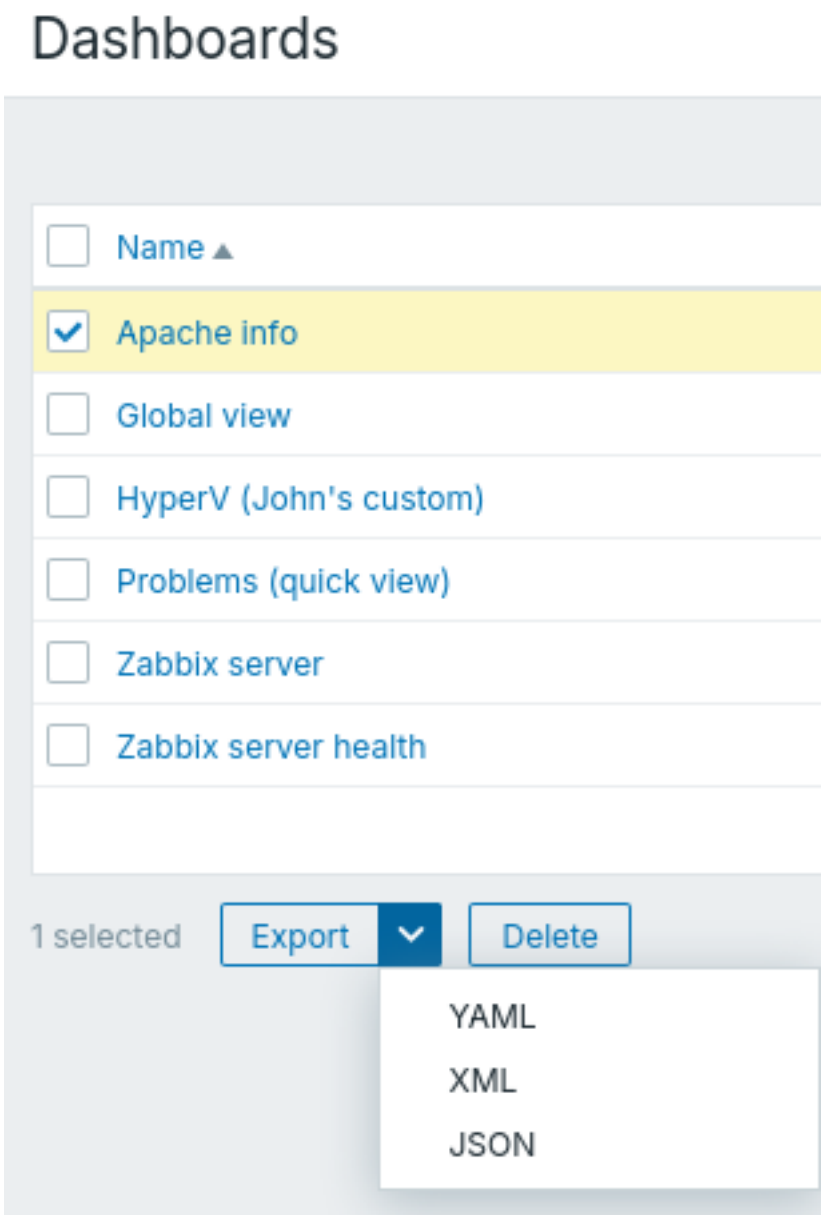
Overview

Dashboards are exported with all related objects and object relations.

Exporting

To export dashboards, do the following:

1. Go to *Dashboards* > *All dashboards*.
2. Mark the checkboxes of the dashboards to export.
3. Click *Export* below the list to export in the default YAML format, or click the arrow to select a different export format.



Depending on the selected format, dashboards are exported to a local file with a default name:

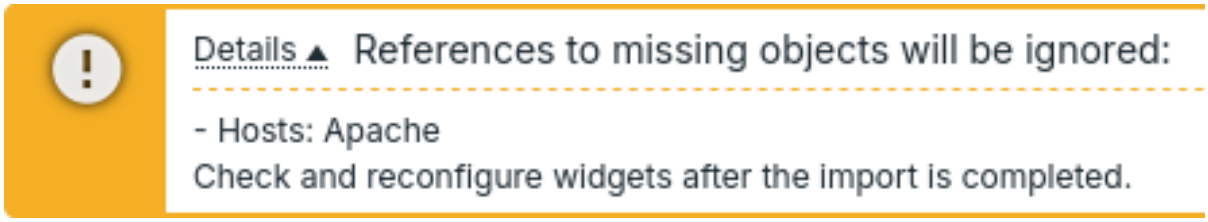
- `zbx_export_dashboards.yaml` - in YAML export;
- `zbx_export_dashboards.xml` - in XML export;
- `zbx_export_dashboards.json` - in JSON export.

Importing

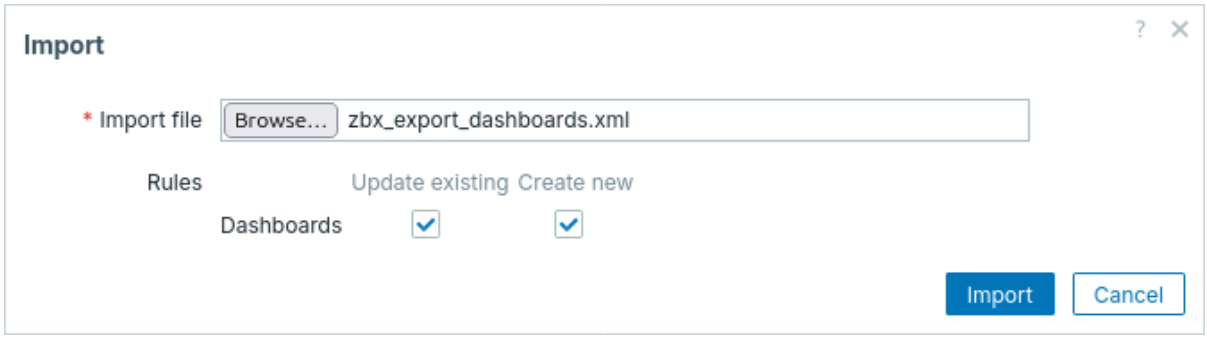
To import dashboards, do the following:

1. Go to *Dashboards* > *All dashboards*.
2. Click *Import* in the upper-right corner. Disabled if the user does not have permission to *Create and edit dashboards*.
3. Select the import file.
4. Mark the checkboxes of the import rules you want to apply (see below).
5. Click *Import* in the lower-right corner of the configuration form.

If the import file references objects that do not exist on the target system, an alert message appears. This is expected for widgets that refer to specific hosts, items, or other objects. In such cases, Zabbix allows the dashboard to be imported, clears the missing references, and either marks the widget as not fully configured when the data is mandatory or drops the filtering criteria when the fields are used only for filtering.



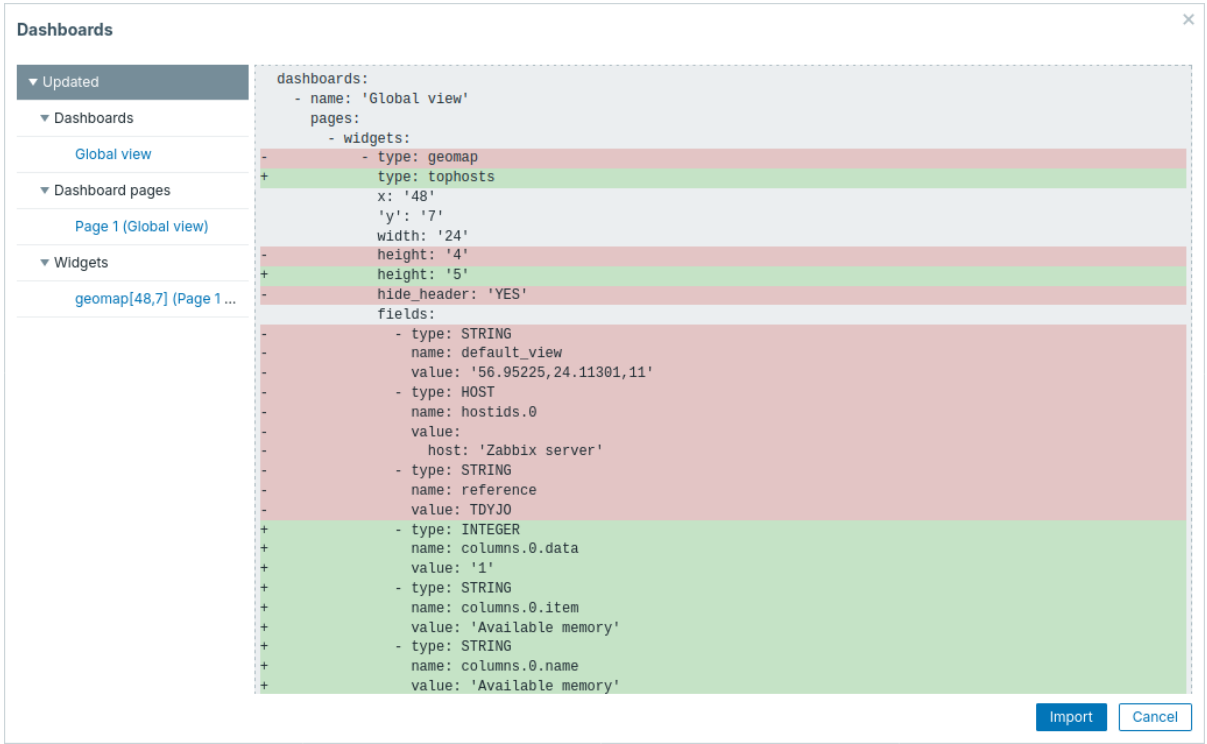
Clicking the *Import* button again will add and/or update the dashboard.



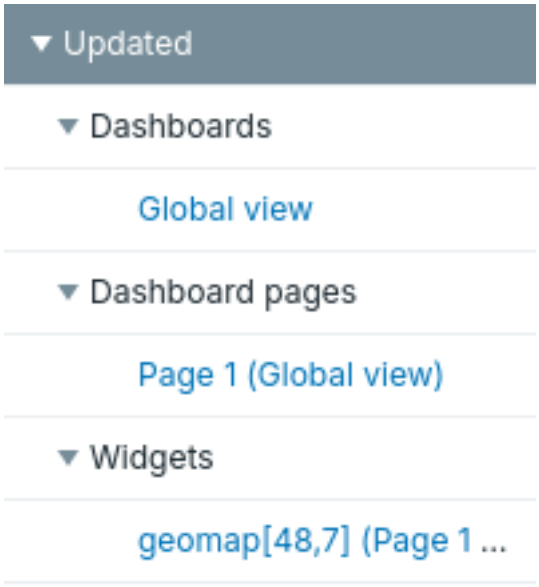
Import rules:

Rule	Description
Update existing	Existing elements will be updated using data from the import file. Otherwise, they will not be updated.
Create new	New elements will be created using data from the import file. Otherwise, they will not be created.

On the next screen, you will be able to view the content of dashboards being imported. If this is a new dashboard, all elements will be listed in green. If updating an existing dashboard, new dashboard elements will be highlighted in green; removed dashboard elements will be highlighted in red; elements that have not changed will be listed on a gray background.



The menu on the left can be used to navigate through the list of changes. The *Updated* section highlights all changes made to existing dashboard elements. The *Added* section lists new dashboard elements. The elements in each section are grouped by element type; click the gray arrow to expand or collapse the group of elements.



Review dashboard changes and then click *Import* to perform the dashboard import. A success or failure message of the import will be displayed in the frontend.

Export format

Export to YAML:

```yaml
zabbix_export:
  version: '8.2'
  dashboards:
    - name: 'Global view'
      pages:
        - widgets:
            - type: tophosts
              name: 'Top hosts by CPU utilization'
              width: '25'
              height: '4'
              fields:
                - type: INTEGER
                  name: column
                  value: '1'
                - type: INTEGER
                  name: columns.0.data
                  value: '2'
                - type: STRING
                  name: columns.0.item
                  value: 'System name'
                - type: STRING
                  name: columns.0.name
                  value: 'Host name'
                - type: STRING
                  name: columns.1.base_color
                  value: 4CAF50
                - type: INTEGER
                  name: columns.1.data
                  value: '1'
                - type: INTEGER
                  name: columns.1.display
                  value: '3'
                - type: STRING
                  name: columns.1.item
                  value: 'CPU utilization'
```

```
      - type: STRING
        name: columns.1.max
        value: '100'
      - type: STRING
        name: columns.1.min
        value: '0'
      - type: STRING
        name: columns.1.name
        value: Utilization
      - type: INTEGER
        name: columns.2.data
        value: '1'
      - type: STRING
        name: columns.2.item
        value: 'Load average (1m avg)'
      - type: STRING
        name: columns.2.name
        value: '1m avg'
      - type: INTEGER
        name: columns.3.data
        value: '1'
      - type: STRING
        name: columns.3.item
        value: 'Load average (5m avg)'
      - type: STRING
        name: columns.3.name
        value: '5m avg'
      - type: INTEGER
        name: columns.4.data
        value: '1'
      - type: STRING
        name: columns.4.item
        value: 'Load average (15m avg)'
      - type: STRING
        name: columns.4.name
        value: '15m avg'
      - type: INTEGER
        name: columns.5.data
        value: '1'
      - type: INTEGER
        name: columns.5.decimal_places
        value: '0'
      - type: STRING
        name: columns.5.item
        value: 'Number of processes'
      - type: STRING
        name: columns.5.name
        value: Processes
      - type: STRING
        name: columnsthresholds.1.color.0
        value: FFFF00
      - type: STRING
        name: columnsthresholds.1.color.1
        value: FF8000
      - type: STRING
        name: columnsthresholds.1.color.2
        value: FF4000
      - type: STRING
        name: columnsthresholds.1.threshold.0
        value: '50'
      - type: STRING
        name: columnsthresholds.1.threshold.1
```

```
            value: '80'
          - type: STRING
            name: columnsthresholds.1.threshold.2
            value: '90'
          - type: STRING
            name: reference
            value: NLVMO
        - type: hostavail
          'y': '4'
          width: '25'
          fields:
            - type: INTEGER
              name: interface_type.0
              value: '1'
        - type: problems
          name: 'Current problems'
          'y': '6'
          width: '48'
          height: '5'
          fields:
            - type: STRING
              name: reference
              value: BRUQE
            - type: INTEGER
              name: show
              value: '3'
            - type: INTEGER
              name: show_tags
              value: '3'
        - type: systeminfo
          x: '25'
          width: '23'
          height: '4'
        - type: problemsbysv
          name: 'Problems by severity'
          x: '25'
          'y': '4'
          width: '23'
          fields:
            - type: STRING
              name: reference
              value: LBSVB
            - type: INTEGER
              name: show_type
              value: '1'
        - type: gauge
          name: 'Zabbix server'
          x: '48'
          width: '12'
          height: '4'
          hide_header: 'YES'
          fields:
            - type: INTEGER
              name: angle
              value: '270'
            - type: INTEGER
              name: desc_size
              value: '8'
            - type: ITEM
              name: itemid.0
              value:
                host: 'Zabbix server'
```

```
      key: vm.memory.utilization
  - type: STRING
    name: max
    value: '100'
  - type: STRING
    name: min
    value: '0'
  - type: INTEGER
    name: scale_size
    value: '11'
  - type: INTEGER
    name: show.0
    value: '1'
  - type: INTEGER
    name: show.1
    value: '2'
  - type: INTEGER
    name: show.2
    value: '3'
  - type: INTEGER
    name: show.3
    value: '4'
  - type: INTEGER
    name: show.4
    value: '5'
  - type: STRING
    name: thresholds.0.color
    value: 42A5F5
  - type: STRING
    name: thresholds.0.threshold
    value: '0'
  - type: STRING
    name: thresholds.1.color
    value: 9CCC65
  - type: STRING
    name: thresholds.1.threshold
    value: '20'
  - type: STRING
    name: thresholds.2.color
    value: FFEE58
  - type: STRING
    name: thresholds.2.threshold
    value: '40'
  - type: STRING
    name: thresholds.3.color
    value: FFA726
  - type: STRING
    name: thresholds.3.threshold
    value: '60'
  - type: STRING
    name: thresholds.4.color
    value: FF7043
  - type: STRING
    name: thresholds.4.threshold
    value: '80'
  - type: INTEGER
    name: th_arc_size
    value: '9'
  - type: INTEGER
    name: th_show_arc
    value: '1'
  - type: INTEGER
```

```
            name: th_show_labels
            value: '1'
          - type: INTEGER
            name: units_bold
            value: '1'
          - type: INTEGER
            name: units_size
            value: '24'
          - type: INTEGER
            name: value_arc_size
            value: '30'
          - type: INTEGER
            name: value_bold
            value: '1'
          - type: INTEGER
            name: value_size
            value: '24'
      - type: svggraph
        x: '48'
        'y': '4'
        width: '24'
        height: '3'
        hide_header: 'YES'
        fields:
          - type: INTEGER
            name: axisx
            value: '0'
          - type: STRING
            name: ds.0.color
            value: AB47BC
          - type: INTEGER
            name: ds.0.fill
            value: '0'
          - type: STRING
            name: ds.0.hosts.0
            value: 'Zabbix server'
          - type: STRING
            name: ds.0.items.0
            value: 'Available memory'
          - type: INTEGER
            name: ds.0.transparency
            value: '2'
          - type: INTEGER
            name: ds.0.type
            value: '2'
          - type: INTEGER
            name: ds.0.width
            value: '2'
          - type: INTEGER
            name: ds.1.axisy
            value: '1'
          - type: STRING
            name: ds.1.color
            value: 42A5F5
          - type: INTEGER
            name: ds.1.fill
            value: '0'
          - type: STRING
            name: ds.1.hosts.0
            value: 'Zabbix server'
          - type: STRING
            name: ds.1.items.0
```

```
            value: 'CPU utilization'
          - type: INTEGER
            name: ds.1.transparency
            value: '0'
          - type: INTEGER
            name: ds.1.type
            value: '2'
          - type: INTEGER
            name: ds.1.width
            value: '2'
          - type: STRING
            name: reference
            value: TOYGS
        - type: geomap
          x: '48'
          'y': '7'
          width: '24'
          height: '4'
          hide_header: 'YES'
          fields:
            - type: STRING
              name: default_view
              value: '56.95225,24.11301,11'
            - type: HOST
              name: hostids.0
              value:
                host: 'Zabbix server'
            - type: STRING
              name: reference
              value: TDYJO
        - type: clock
          name: 'Local time'
          x: '60'
          width: '12'
          hide_header: 'YES'
          fields:
            - type: INTEGER
              name: clock_type
              value: '1'
            - type: INTEGER
              name: show.0
              value: '1'
            - type: INTEGER
              name: show.1
              value: '2'
            - type: INTEGER
              name: show.2
              value: '3'
            - type: INTEGER
              name: time_sec
              value: '0'
            - type: INTEGER
              name: tzone_format
              value: '1'
        - type: item
          name: 'Server performance'
          x: '60'
          'y': '2'
          width: '12'
          hide_header: 'YES'
          fields:
            - type: INTEGER
```

```
      name: aggregate_function
      value: '2'
    - type: INTEGER
      name: decimal_size
      value: '32'
    - type: STRING
      name: description
      value: 'Values per second'
    - type: INTEGER
      name: desc_h_pos
      value: '0'
    - type: INTEGER
      name: desc_v_pos
      value: '0'
    - type: ITEM
      name: itemid.0
      value:
        host: 'Zabbix server'
        key: 'zabbix[wcache,values]'
    - type: INTEGER
      name: show.0
      value: '1'
    - type: INTEGER
      name: show.1
      value: '2'
    - type: INTEGER
      name: show.2
      value: '5'
    - type: INTEGER
      name: sparkline.fill
      value: '2'
    - type: INTEGER
      name: sparkline.width
      value: '2'
    - type: STRING
      name: thresholds.0.color
      value: DCEDC8
    - type: STRING
      name: thresholds.0.threshold
      value: '20'
    - type: STRING
      name: thresholds.1.color
      value: FFF9C4
    - type: STRING
      name: thresholds.1.threshold
      value: '40'
    - type: STRING
      name: thresholds.2.color
      value: FFCCBC
    - type: STRING
      name: thresholds.2.threshold
      value: '60'
    - type: INTEGER
      name: units_bold
      value: '0'
    - type: INTEGER
      name: units_show
      value: '0'
    - type: INTEGER
      name: value_h_pos
      value: '0'
    - type: INTEGER
```

```
                name: value_size
                value: '32'
```

Exported elements

Exported elements are explained in the table below.

Element	Type	Description
name	string	(required) Dashboard name.
pages		Root element for dashboard pages
widgets		Root element for dashboard pages

Note:
See also: Dashboard object (refer to the relevant property with a matching name).

Dashboard widget

Element	Type	Description
type	string	(required) Type of the dashboard widget. Possible values: actionlog - Action log; clock - Clock; discovery - Discovery status; favgraphs - Favorite graphs; favmaps - Favorite maps; gauge - Gauge; geomap - Geomap; graph - Graph (classic); graphprototype - Graph prototype; honeycomb - Honeycomb; hostavail - Host availability; hostcard - Host card; hostnavigator - Host navigator; itemcard - Item card; itemhistory - Item history; itemnavigator - Item navigator; item - Item value; map - Map; navtree - Map Navigation Tree; piechart - Pie chart; problemhosts - Problem hosts; problems - Problems; problemsbysv - Problems by severity; scatterplot - Scatter plot; slareport - SLA report; svggraph - Graph; systeminfo - System information; tophosts - Top hosts; topitems - Top items; toptriggers - Top triggers; trigover - Trigger overview; url - URL; web - Web monitoring.
x	integer	A horizontal position from the left side of the dashboard. Possible values:[1] 0-71.
y	integer	A vertical position from the top of the dashboard. Possible values:[1] 0-63.
name	string	Custom widget name.
width	integer	The widget width. Possible values:[1] 1-72.

Element	Type	Description
height	integer	The widget height. Possible values:[1] 1-64.
hide_header	string	Hide widget header. Possible values:[1] NO (0, default), YES (1).
fields		Root element for the dashboard widget fields.

Note:
See also: Dashboard widget (refer to the relevant property with a matching name).

Dashboard widget field

Element	Type	Description
type	string	(required) Widget field type. Possible values:[1] INTEGER (0), STRING (1), HOST_GROUP (2), HOST (3), ITEM (4), ITEM_PROTOTYPE (5), GRAPH (6), GRAPH_PROTOTYPE (7), MAP (8), SERVICE (9), SLA (10), USER (11), ACTION (12), MEDIA_TYPE (13).
name	string	(required) Widget field name.
value	mixed	(required) Widget field value, depending on the field type.
host	string	Item host.
key	string	Item key.

Note:
See also: Dashboard widget field (refer to the relevant property with a matching name).

Footnotes

[1] API integer values in brackets, for example, ENABLED (0), are mentioned only for reference. For more information, see the linked API object page in the table entry or at the end of each section.

Regular expressions

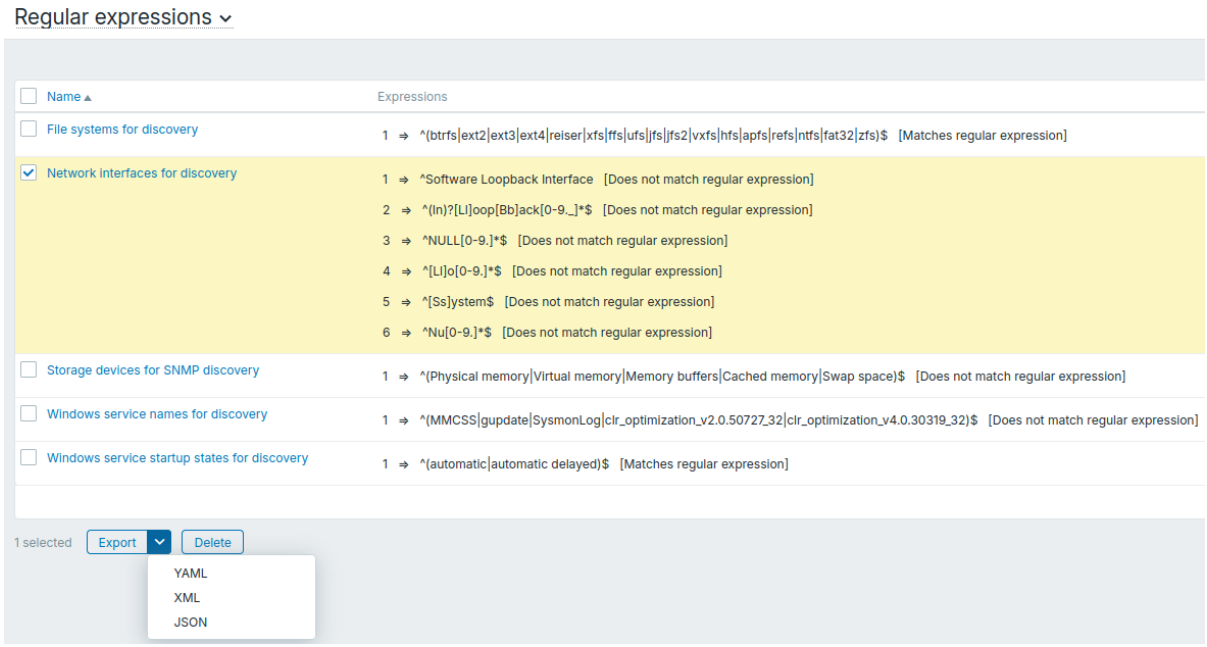
Overview

Regular expressions are exported with all related objects and object relations.

Exporting

To export regular expressions, do the following:

1. Go to *Administration* > *General* > *Regular expressions*.
2. Mark the checkboxes of the regular expressions to export.
3. Click *Export* below the list to export in the default YAML format, or click the arrow to select a different export format.



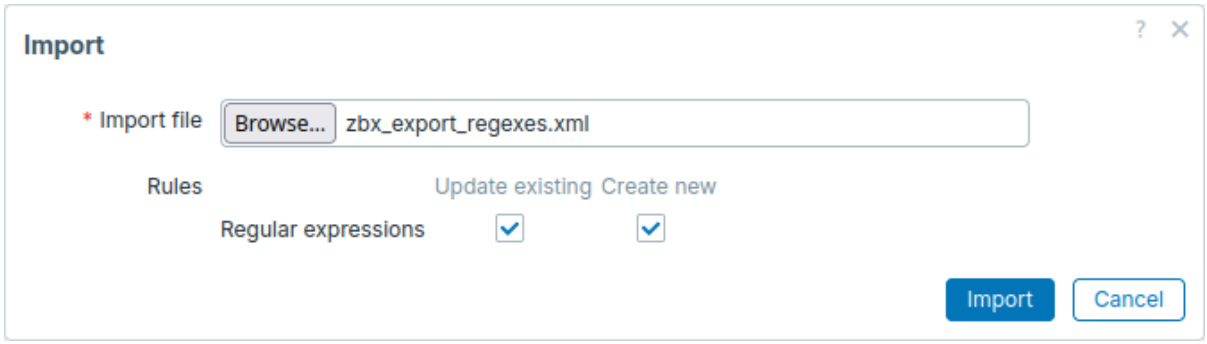
Depending on the selected format, regular expressions are exported to a local file with a default name:

- `zbx_export_regexes.yaml` - in YAML export;
- `zbx_export_regexes.xml` - in XML export;
- `zbx_export_regexes.json` - in JSON export.

Importing

To import regular expressions, do the following:

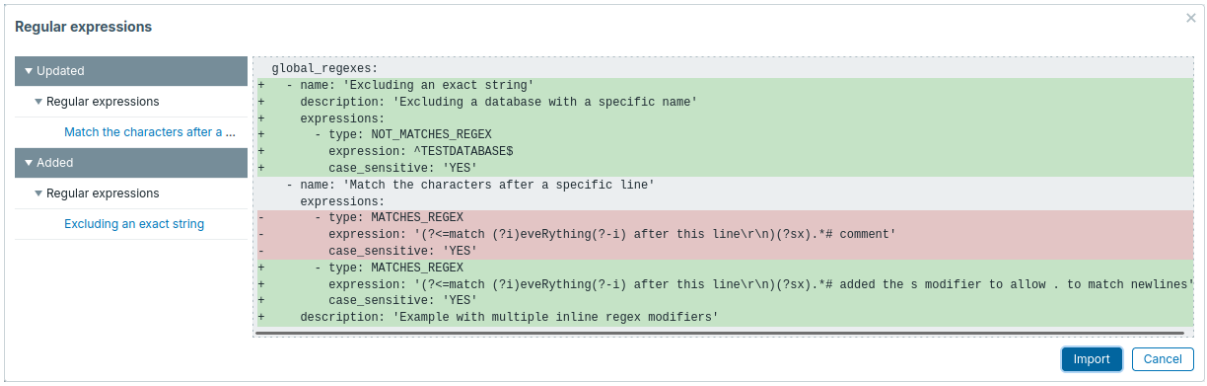
1. Go to *Administration* > *General* > *Regular expressions*.
2. Click *Import* in the upper-right corner.
3. Select the import file.
4. Mark the checkboxes of the import rules you want to apply (see below).
5. Click *Import* in the lower-right corner of the configuration form.



Import rules:

Rule	Description
Update existing	Existing regular expressions will be updated using data from the import file, fully overwriting existing expressions with the same name. Otherwise, they will not be updated.
Create new	New regular expressions will be created using data from the import file. Otherwise, they will not be created.

On the next screen, you will be able to view the content of the regular expressions being imported. If this is a new regular expression, all data will be listed in green. If updating an existing regular expression, new regular expression data will be highlighted in green; removed data will be highlighted in red; unchanged data will be displayed on a gray background.



The menu on the left can be used to navigate through the list of changes. The *Updated* section highlights all changes made to existing regular expressions. The *Added* section lists new regular expressions. Click the arrow to expand or collapse a group of regular expressions.

Review regular expression changes and then click *Import* to perform the regular expressions import. A success or failure message of the import will be displayed in the frontend.

Export format

Export to YAML:

```yaml
zabbix_export:
  version: '8.2'
  global_regexes:
    - name: 'File systems for discovery'
      description: 'Global regular expression used to discover and identify file systems on the monitored
      expressions:
        - type: MATCHES_REGEX
          expression: ^(btrfs|ext2|ext3|ext4|reiser|xfs|ffs|ufs|jfs|jfs2|vxfs|hfs|apfs|refs|ntfs|fat32|zfs
    - name: 'Storage devices for SNMP discovery'
      expressions:
        - type: NOT_MATCHES_REGEX
          expression: '^(Physical memory|Virtual memory|Memory buffers|Cached memory|Swap space)$'
          case_sensitive: 'YES'
```

Exported elements

Exported elements of regular expressions:

Element	Type	Description
global_regexes	object	Root element for regular expressions.
name	string	(required) Name of the regular expression.
description	text	Description of the regular expression.
expressions	object	Root element for regular expressions.
type	string	Type of the regular expression.
expression	string	Regular expression.
exp_delimiter	string	Regular expression delimiter; present when the type is *Contains any substring from list*.
case_sensitive	string	Case-sensitivity option for the regular expression.

Note:
See also: Regular expression object (refer to the relevant property with a matching name).

Regular expressions

Overview Perl Compatible Regular Expressions (PCRE, PCRE2) are supported in Zabbix.

There are two ways of using regular expressions in Zabbix:

- Manually entering a regular expression

- Using a global regular expression created in Zabbix

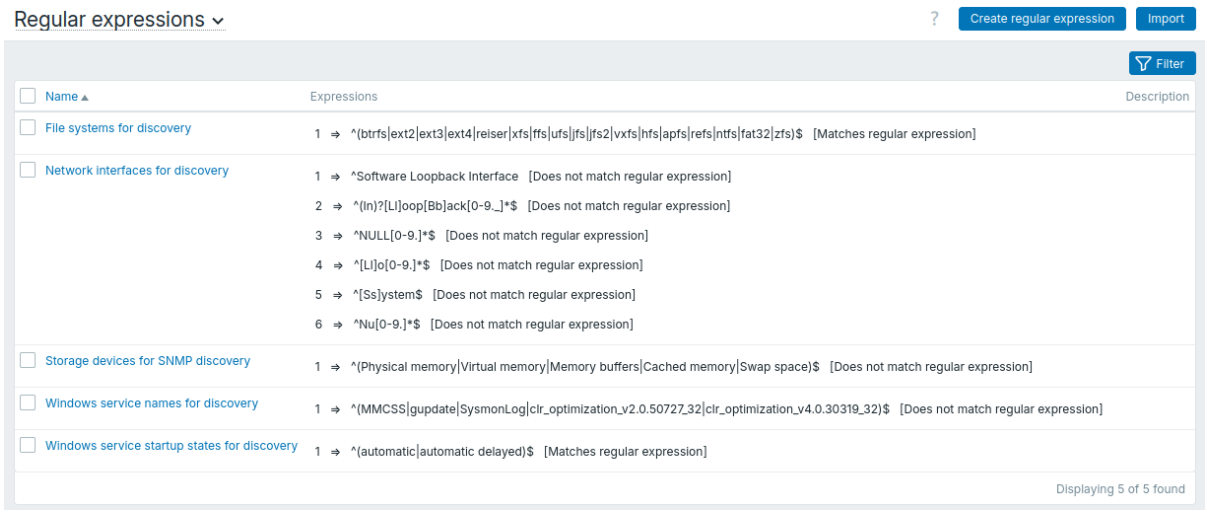
Regular expressions You may manually enter a regular expression in supported places.
 Note that the expression must not start with @, as this symbol is used in Zabbix to reference global regular expressions.

Attention:
It's possible to run out of stack memory when using regular expressions. See the PCRE PERFORMANCE for more information.

Note that in multiline matching, the ^ and $ anchors match at the beginning and end of each line respectively, instead of the beginning and end of the entire string.

See also examples for correct escaping in various contexts.

Global regular expressions Zabbix provides an advanced editor in the frontend for creating and testing complex regular expressions.
 To access it, go to *Administration* > *General* > *Regular expressions*.
 Zabbix includes a number of predefined global regular expressions in its default dataset.

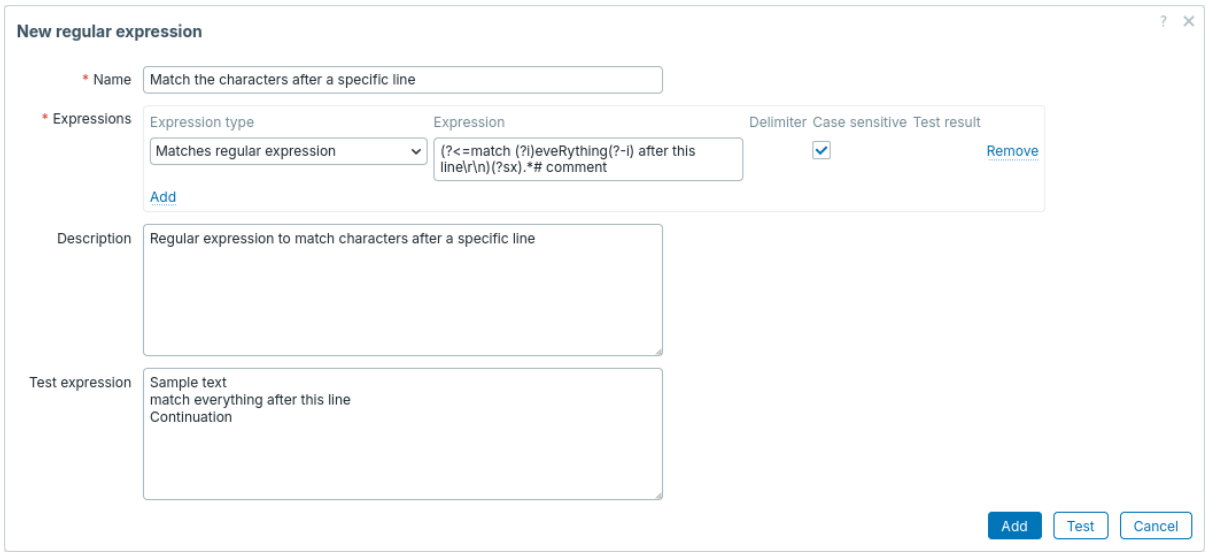


You can filter the regular expressions list by name and description to find the expressions you need.

Configuration

To create a global regular expression:

1. Click *Create regular expression*.
2. Enter the regular expression parameters in the form.



All mandatory input fields are marked with a red asterisk.

Parameter	Description
Name	Set the regular expression name. Any Unicode characters are allowed.

Parameter	Description
Expressions	Click *Add* in the *Expressions* section to add a new subexpression.
Expression type	Select expression type: **Contains string** - matches the substring. **Contains any substring from list** - matches any substring from a delimited list. The delimited list includes a comma `,`, a dot `.` or a forward slash `/`. **Does not contain string** - matches any string except the substring. **Matches regular expression** - matches the regular expression. **Does not match regular expression** - does not match the regular expression.
Expression	Enter substring/regular expression. Maximum length: 2048 characters.
Delimiter	A comma `,`, a dot `.`, or a forward slash `/` to separate text strings in a regular expression. This parameter is active only when *Contains any substring from list* expression type is selected.
Case sensitive	A checkbox to specify whether a regular expression is sensitive to capitalization of letters.
Description	Enter a description for the regular expression.
Test expression	A value to use for testing the regular expression.

A forward slash `/` in the expression is treated literally rather than as a delimiter. This makes it possible to save expressions containing a slash without errors.

Attention:
A custom regular expression name in Zabbix may contain commas, spaces, etc. In cases where this may lead to misinterpretation when referencing (for example, a comma in the parameter of an item key), the entire reference may be enclosed in quotes, like this: *"@My custom regexp for purpose1, purpose2"*.
 Regular expression names must not be quoted in other locations (for example, in LLD rule properties).

Once a regular expression has been created, it can be used in several places in the frontend by referring to its name, prefixed with `@`, for example, `@mycustomregexp`.

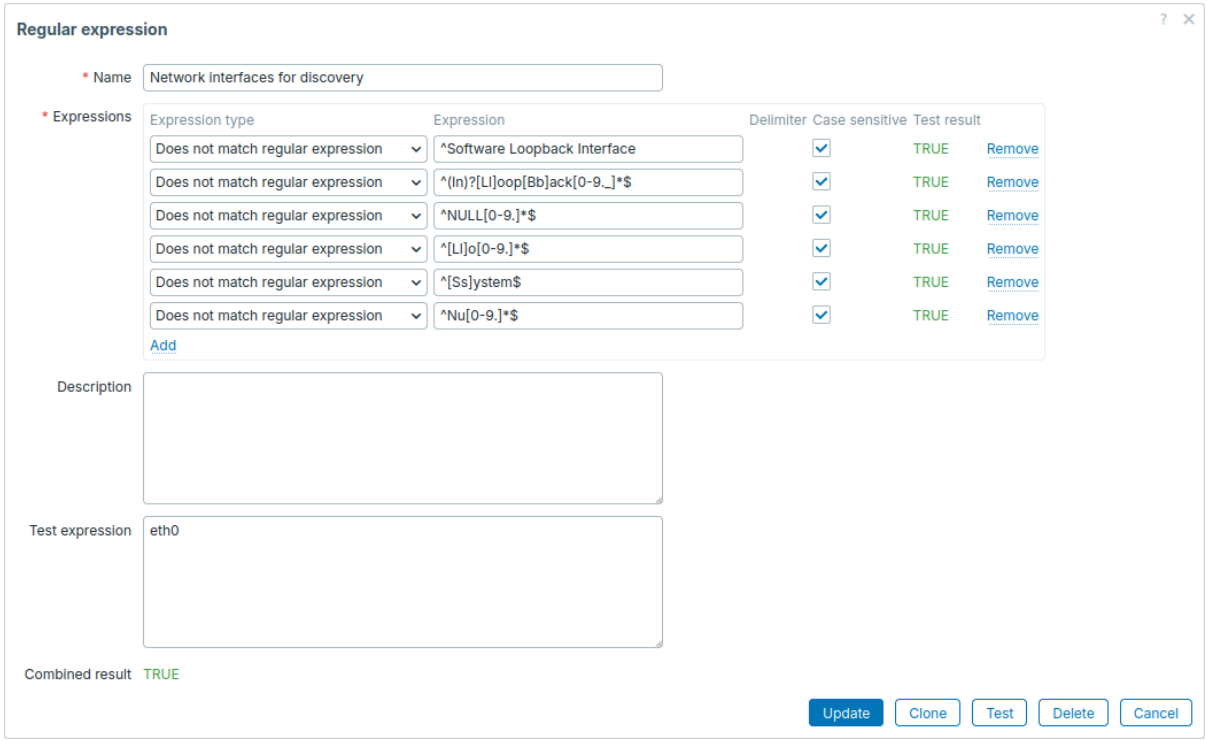
Form buttons

Buttons at the bottom of the form allow you to perform various operations:

Button	Description
Add	Add a regular expression. This button is only available for new regular expressions.
Update	Update the properties of a regular expression.
Clone	Create another regular expression based on the properties of the current regular expression.
Test	Test a regular expression and its subexpressions against a provided value.
Delete	Delete the regular expression.
Cancel	Cancel editing the regular expression properties.

Testing a regular expression

You may test the regular expression and its subexpressions by entering a value in the *Test expression* field and clicking the *Test* button.



In the *Test result* column, the status of each subexpression is displayed. The *Combined result* field displays the overall status of the regular expression. If multiple subexpressions are defined, Zabbix uses the logical AND operator to calculate the *Combined result*; that is, if at least one result is `False`, the *Combined result* is also `False`.

Default global regular expressions

The default dataset includes the following predefined global regular expressions:

Name	Expression	Matches
File systems for discovery	`^(btrfs|ext2|ext3|ext4|reiser|xfs|ffs|ufs|jfs|jfs2|vxfs|hfs|apfs|refs|ntfs|fat32|zfs)$`	"btrfs" or "ext2" or "ext3" or "ext4" or "reiser" or "xfs" or "ffs" or "ufs" or "jfs" or "jfs2" or "vxfs" or "hfs" or "apfs" or "refs" or "ntfs" or "fat32" or "zfs".
Network interfaces for discovery	`^Software Loopback Interface`	Strings starting with "Software Loopback Interface".
	`^lo$`	"lo".
	`^(In)?[Ll]oop[Bb]ack[0-9._]*$`	Strings optionally starting with "In", followed by "L" or "l", followed by "oop", followed by "B" or "b", followed by "ack", and optionally followed by any number of digits, dots, or underscores.
	`^NULL[0-9.]*$`	Strings starting with "NULL" optionally followed by any number of digits or dots.
	`^[Ll]o[0-9.]*$`	Strings starting with "Lo" or "lo" and optionally followed by any number of digits or dots.
	`^[Ss]ystem$`	"System" or "system".
	`^Nu[0-9.]*$`	Strings starting with "Nu" optionally followed by any number of digits or dots.
Storage devices for SNMP discovery	`^(Physical memory|Virtual memory|Memory buffers|Cached memory|Swap space)$`	"Physical memory" or "Virtual memory" or "Memory buffers" or "Cached memory" or "Swap space".
Windows service names for discovery	`^(MMCSS|gupdate|SysmonLog|clr_optimization_v2.0.50727_32|clr_optimization_v4.0.30319_32)$`	"MMCSS" or "gupdate" or "SysmonLog" or strings matching "clr_optimization_v2.0.50727_32" or "clr_optimization_v4.0.30319_32", where each dot can be replaced by any character except a newline.
Windows service startup states for discovery	`^(automatic|automatic delayed)$`	"automatic" or "automatic delayed".

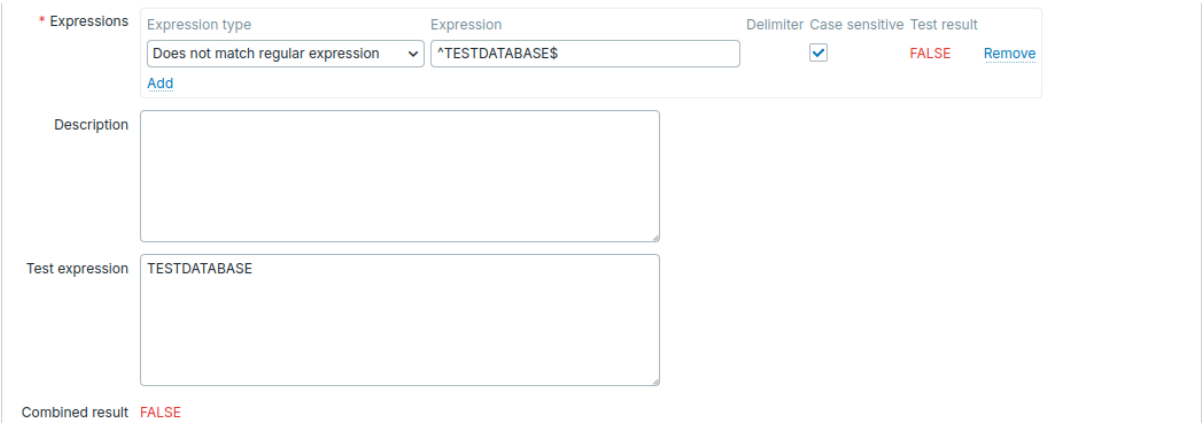
Importing/exporting a regular expression

You can import or export global regular expressions as described on the Regular expressions export/import configuration page.

Examples Example of excluding an exact string

Use the following expression in low-level discovery to discover databases except for a database with a specific name:

```
^TESTDATABASE$
```

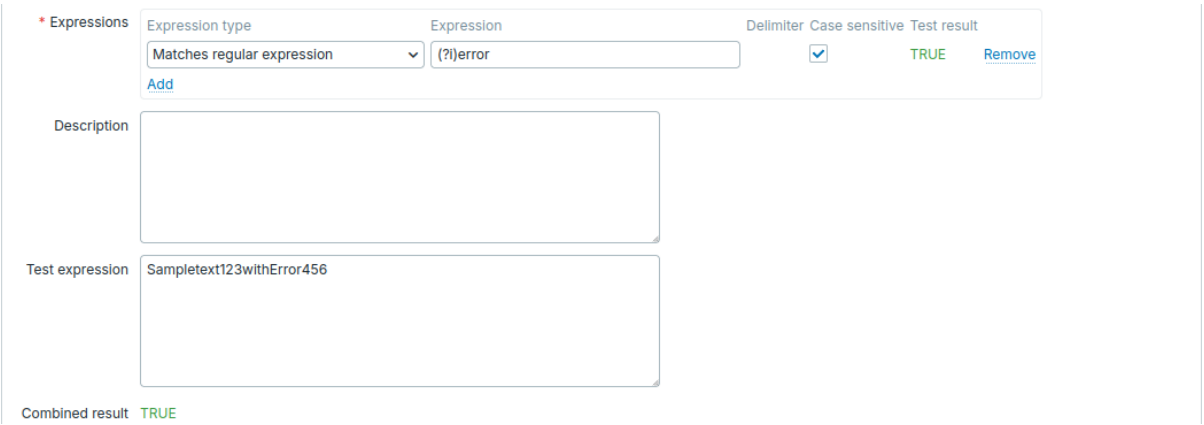


Both *Test result* and *Combined result* display FALSE: the name matches "TESTDATABASE", so the *Does not match regular expression* condition is not satisfied.

Example with an inline regex modifier for a case-insensitive match

Use the following regular expression, which includes the inline modifier `(?i)` to perform a case-insensitive match of the characters "error":

```
(?i)error
```



Both *Test result* and *Combined result* display TRUE: the characters "error" are matched even with *Case sensitive* selected, because the inline modifier `(?i)` makes the match case-insensitive.

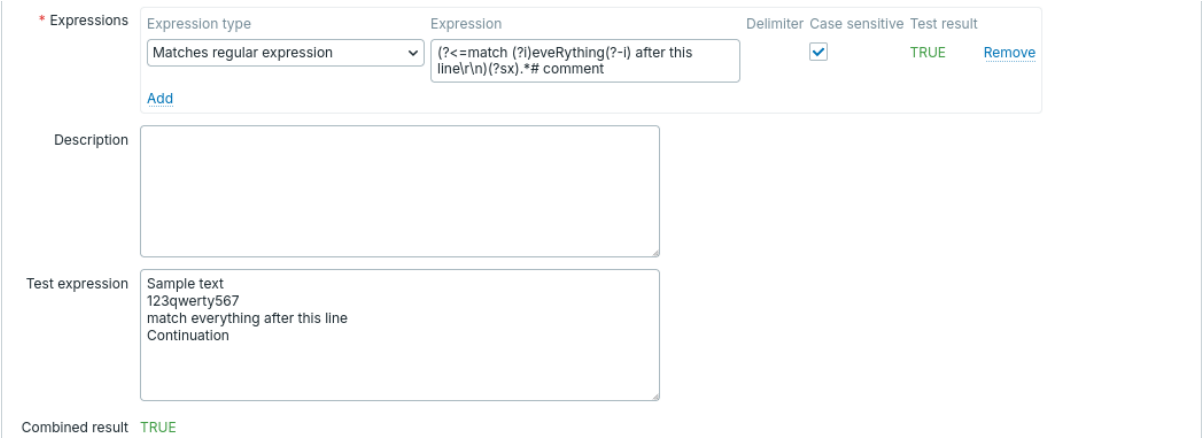
Example with multiple inline regex modifiers

Use the following regular expression, which includes multiple inline modifiers, to match the characters after a specific line:

```
(?<=match (?i)eveRything(?-i) after this line\r\n)(?sx).*# comment
```

where:

- `(?i)` - turns case-insensitive matching on;
- `(?-i)` - turns case-insensitive matching off;
- `(?s)` - enables dotall mode, allowing `.` to match newline characters;
- `(?x)` - enables extended mode, allowing whitespace and comments in the regular expression to be ignored.



Both *Test result* and *Combined result* display TRUE: the characters following the specific line are matched.

Attention:
The g modifier cannot be specified inline. The list of available modifiers can be found in pcresyntax man page. For more information about PCRE syntax, please refer to PCRE HTML documentation.

Regular expression support by location

Location		Regular expression	Global regular expression	Multiline matching	Comments
Agent items					
	eventlog[]	Yes	Yes	Yes	`regexp`, `severity`, `source`, `eventid` parameters
	eventlog.count[]				`regexp`, `severity`, `source`, `eventid` parameters
	log[]				`regexp` parameter
	log.count[]				
	logrt[]		Yes/No		`regexp` parameter supports both, `file_regexp` parameter supports non-global expressions only
	logrt.count[]				
	proc.cpu.util[]		No	No	`cmdline` parameter
	proc.get[]				
	proc.mem[]				
	proc.num[]				
	sensor[]				`device` and `sensor` parameters on Linux 2.4
	system.hw.macaddr[]				`interface` parameter
	system.sw.packages[]				`regexp` parameter
	system.sw.packages.get[]				`regexp` parameter
	vfs.dir.count[]				`regex_incl`, `regex_excl`, `regex_excl_dir` parameters
	vfs.dir.get[]				`regex_incl`, `regex_excl`, `regex_excl_dir` parameters
	vfs.dir.size[]				`regex_incl`, `regex_excl`, `regex_excl_dir` parameters
	vfs.file.regexp[]			Yes	`regexp` parameter
	vfs.file.regmatch[]				
	web.page.regexp[]				
SNMP traps					
	snmptrap[]	Yes	Yes	No	`regexp` parameter

Location		Regular expression	Global regular expression	Multiline matching	Comments
Item value preprocessing		Yes	No	No	`pattern` parameter
Functions for triggers/calculated items					
	count()	Yes	Yes	Yes	`pattern` parameter if `operator` parameter is *regexp* or *iregexp*
	countunique()	Yes	Yes		
	find()	Yes	Yes		
	logeventid()	Yes	Yes	No	`pattern` parameter
	logsource()				
Low-level discovery					
	Filters	Yes	Yes	No	*Regular expression* field
	Overrides	Yes	No		In *matches*, *does not match* options for *Operation* conditions
Action conditions		Yes	No	No	In *matches*, *does not match* options for *Host name* and *Host metadata* autoregistration conditions
Scripts		Yes	Yes	No	*Input validation rule* field
Web monitoring		Yes	No	Yes	*Variables* with a **regex:** prefix *Required string* field
User macro context		Yes	No	No	In macro context with a **regex:** prefix
Macro functions					
	regsub() iregsub()	Yes	No	No	`pattern` parameter
Link indicators in maps		Yes	No	No	*Pattern* field (for text items)
Icon mapping		Yes	Yes	No	*Expression* field
Value mapping		Yes	No	No	*Value* field if mapping type is `regexp`
Agent configuration parameters					
	AllowKeyRegexp DenyKeyRegexp	Yes	No	No	Pattern

Events and resolution

Overview There are several types of events generated in Zabbix:

- trigger events - whenever a trigger changes its status (*OK→PROBLEM→OK*)
- service events - whenever a service changes its status (*OK→PROBLEM→OK*)
- discovery events - when hosts or services are detected
- autoregistration events - when active agents are auto-registered by server
- internal events - when an item/low-level discovery rule becomes unsupported or a trigger goes into an unknown state

Events are time-stamped and can be the basis of actions such as sending notification email, creating tickets, or executing remote commands.

To view event details in the frontend, go to *Monitoring* > *Problems*. There you can click the event date and time to open the event details page, where you can see additional information about the event, including whether it was suppressed because of maintenance or suppressed manually by a user.

Problem resolution Problem events created by triggers are resolved automatically, if the condition that caused the problem no longer exists.

To add flexibility and convenience to problem event management, the following options exist:

- Event acknowledgment and suppression
- Complex event processing and (deprecated) event correlation
- Manual closing of problems
- Maintenance periods - define periods when problems are suppressed
- Tagging - tag problems by common feature

Trigger events Change of trigger state is the most frequent and most important source of events. Each time the trigger changes its state, an event is generated. The event contains details of the trigger state's change - when it happened and what the new state is.

Two types of events are created by triggers - `Problem` and `OK`.

Problem events

A problem event is created:

- when a trigger expression evaluates to TRUE if the trigger is in OK state;
- each time a trigger expression evaluates to TRUE if multiple problem event generation is enabled for the trigger.

OK events

An OK event closes the related problem event(s) and may be created by 3 components:

- triggers - based on 'OK event generation' and 'OK event closes' settings;
- event correlation
- task manager - when an event is manually closed

Triggers have an 'OK event generation' setting that controls how OK events are generated:

- *Expression* - an OK event is generated for a trigger in problem state when its expression evaluates to FALSE. This is the simplest setting, enabled by default.
- *Recovery expression* - an OK event is generated for a trigger in problem state when its expression evaluates to FALSE and the recovery expression evaluates to TRUE. This can be used if trigger recovery criteria is different from problem criteria.
- *None* - an OK event is never generated. This can be used in conjunction with multiple problem event generation to simply send a notification when something happens.

Additionally triggers have an 'OK event closes' setting that controls which problem events are closed:

- *All problems* - an OK event will close all open problems created by the trigger
- *All problems if tag values match* - an OK event will close open problems created by the trigger and having at least one matching tag value. The tag is defined by 'Tag for matching' trigger setting. If there are no problem events to close then OK event is not generated. This is often called trigger level event correlation.

Event correlation

Event correlation (also called global event correlation) is a way to set up custom event closing (resulting in OK event generation) rules.

The rules define how the new problem events are paired with existing problem events allow the new event or the matched events to be closed by generating corresponding OK events.

However, event correlation must be configured very carefully, as it can negatively affect event processing performance or, if misconfigured, close more events than intended (in the worst case even all problem events could be closed). A few configuration tips:

1. always reduce the correlation scope by setting a unique tag for the control event (the event that is paired with old events) and use the 'new event tag' correlation condition
2. don't forget to add a condition based on the old event when using 'close old event' operation, or all existing problems could be closed
3. avoid using common tag names used by different correlation configurations

Task manager

If the 'Allow manual close' setting is enabled for trigger, then it's possible to manually close problem events generated by the trigger. This is done in the frontend when updating a problem. The event is not closed directly - instead a 'close event' task is created, which is handled by the task manager shortly. The task manager will generate a corresponding OK event and the problem event will be closed.

Service events Service events are generated only if service actions for these events are enabled. In this case, each change of service state creates a new event:

- Problem event - when service state is changed from `OK` to `Problem`
- OK event - when service state is changed from `Problem` to `OK`

The event contains details of the service state change - when it happened and what the new state is.

Discovery events Zabbix periodically scans the IP ranges defined in network discovery rules. Frequency of the check is configurable for each rule individually. Once a host or a service is discovered, a discovery event (or several events) are generated.

Zabbix generates the following events:

Event	When generated
Service Up	Every time Zabbix detects active service.
Service Down	Every time Zabbix cannot detect service.
Host Up	If at least one of the services is UP for the IP.
Host Down	If all services are not responding.
Service Discovered	If the service is back after downtime or discovered for the first time.
Service Lost	If the service is lost after being up.
Host Discovered	If host is back after downtime or discovered for the first time.
Host Lost	If host is lost after being up.

Active agent autoregistration events Active agent autoregistration creates events in Zabbix.

If configured, active agent autoregistration event is created when a previously unknown active agent asks for checks or if the host metadata has changed. The server adds a new auto-registered host, using the received IP address and port of the agent.

For more information, see the active agent autoregistration page.

Internal events Internal events happen when:

- an item changes state from 'normal' to 'unsupported'
- an item changes state from 'unsupported' to 'normal'
- a low-level discovery rule changes state from 'normal' to 'unsupported'
- a low-level discovery rule changes state from 'unsupported' to 'normal'
- a trigger changes state from 'normal' to 'unknown'
- a trigger changes state from 'unknown' to 'normal'

The aim of introducing internal events is to allow users to be notified when any internal event takes place, for example, an item becomes unsupported and stops gathering data.

Internal events are only created when internal actions for these events are enabled. To stop generation of internal events (for example, for items becoming unsupported), disable all actions for internal events in *Alerts* > *Actions* > *Internal actions*.

Note:
If internal actions are disabled, while an object is in the 'unsupported' state, recovery event for this object will still be created.
If internal actions are enabled, while an object is in the 'unsupported' state, recovery event for this object will be created, even though 'problem event' has not been created for the object.

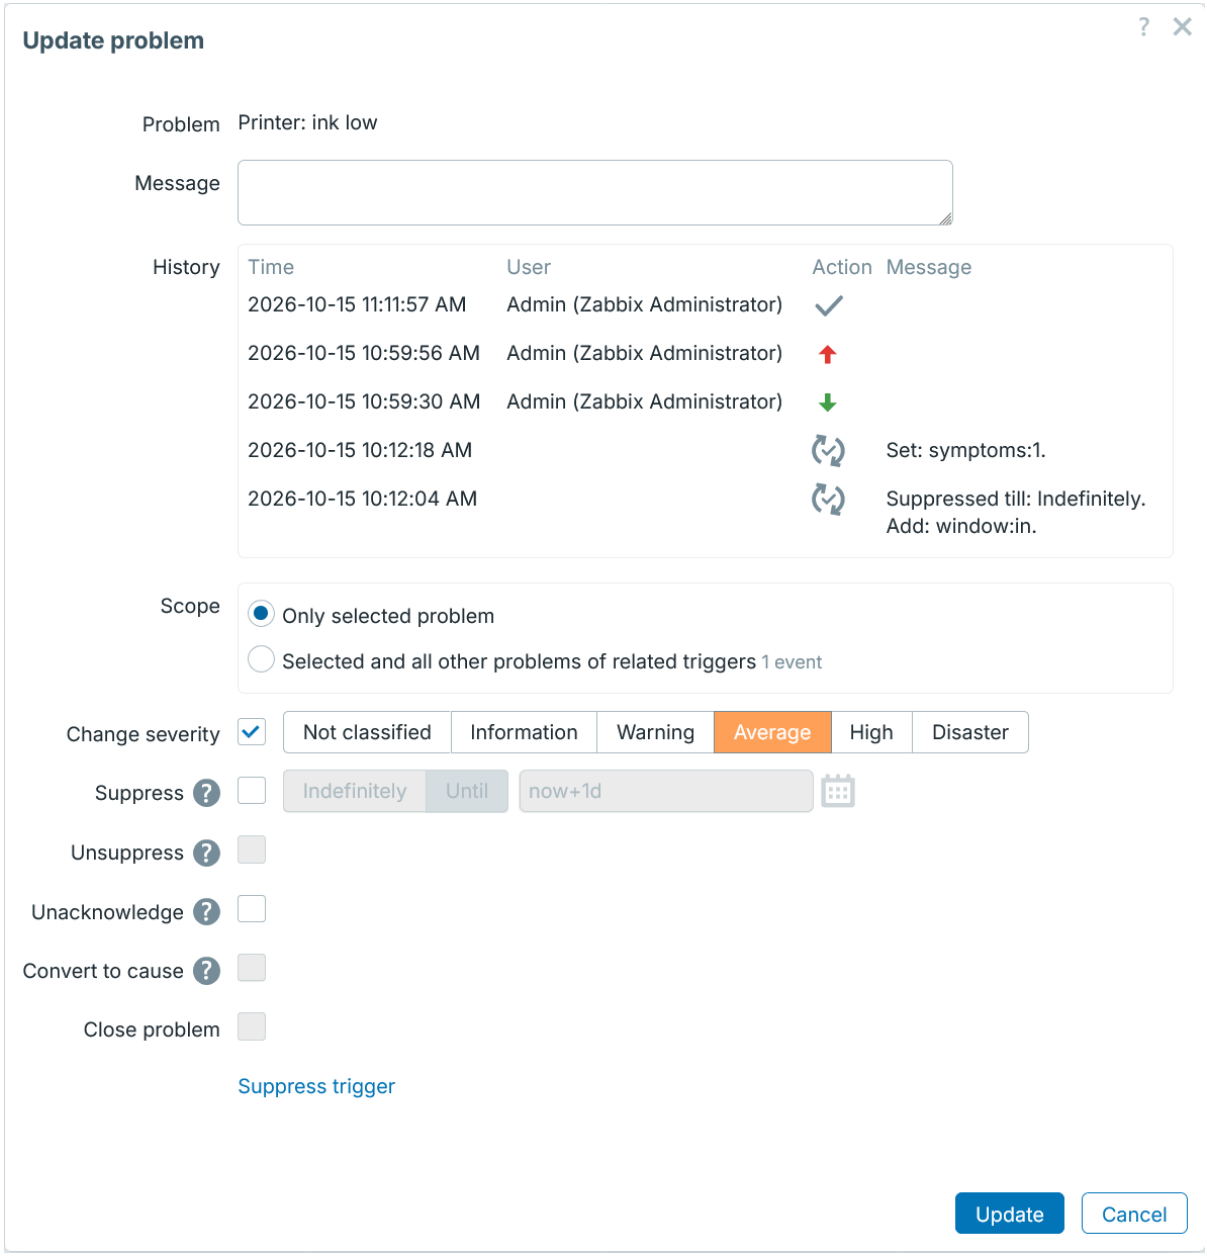
See also: Receiving notification on unsupported items

Problem acknowledgment and suppression

Problem events in Zabbix can be acknowledged or suppressed by users.

Updating problems The problem update popup allows you to:

- Comment on the problem.
- View comments and actions so far.
- Change problem severity.
- Suppress/unsuppress problem.
- Acknowledge/unacknowledge problem.
- Change symptom problem to cause problem.
- Manually close problem.
- Create an ad-hoc maintenance for the trigger.



All mandatory input fields are marked with a red asterisk.

Parameter	Description
Problem	If only one problem is selected, the problem name is displayed. If several problems are selected, *N problems selected* is displayed.
Message	Enter text to comment on the problem (maximum 2048 characters).
History	Previous activities and comments on the problem are listed, along with the time and user details. For the meaning of icons used to denote user actions see the event detail page. Note that history is displayed if only one problem is selected for the update.
Scope	Define the scope of such actions as changing severity, acknowledging or manually closing problems: **Only selected problem** - will affect this event only **Selected and all other problems of related triggers** - in case of acknowledgment/closing problem, will affect this event and all other problems that are not acknowledged/closed so far. If the scope contains problems already acknowledged or closed, these problems will not be acknowledged/closed repeatedly. On the other hand, the number of message and severity change operations are not limited.
Change severity	Mark the checkbox and click the severity button to update problem severity. The checkbox for changing severity is available if read-write permissions exist for at least one of the selected problems. Only those problems that are read-writable will be updated when clicking *Update*. If read-write permissions exist for none of the selected triggers, the checkbox is disabled.
Suppress	Mark the checkbox to suppress the problem: **Indefinitely** - suppress indefinitely **Until** - suppress until a given time. Both absolute and relative time formats are supported, for example: `now+1d` - for one day from now (default) `now/w` - until the end of the current week `2022-05-28 12:00:00` - until absolute date/time Note that a simple period (e. g., `1d`, `1w`) is not supported. Availability of this option depends on the "Suppress problems" user role settings. See also: Problem suppression
Unsuppress	Mark the checkbox to unsuppress the problem. This checkbox is active only if at least one of the selected problems is suppressed. Availability of this option depends on the "Suppress problems" user role settings.
Acknowledge	Mark the checkbox to acknowledge the problem. This checkbox is available if there is at least one unacknowledged problem among the selected. It is not possible to add another acknowledgment for an already acknowledged problem (it is possible to add another comment though).
Unacknowledge	Mark the checkbox to unacknowledge the problem. This checkbox is available if there is at least one acknowledged problem among the selected.
Convert to cause	Mark the checkbox to convert the symptom problem(s) to cause problem(s).
Close problem	Mark the checkbox to manually close the selected problem(s). The checkbox for closing a problem is available if the *Allow manual close* option is checked in trigger configuration for at least one of the selected problems. Only those problems that are allowed to be closed will be closed when you click *Update*. If no problem is manually closeable, the checkbox is disabled. Already closed problems will not be closed repeatedly.
Suppress trigger	Click to start ad-hoc maintenance for the trigger. A maintenance period configuration form will open. In it, a one-time maintenance period of `1h` (configurable in user settings and profile) starting from now will be pre-selected. The respective trigger will be preselected, too. The maintenance name will be pre-filled, in the format `Ad-hoc: <event name>`.

Problem acknowledgment If a user gets notified about a problem event, they can go to Zabbix frontend, open the problem update popup window of that problem using one of the ways listed below and acknowledge the problem. When acknowledging, they can enter their comment for it, saying that they are working on it or whatever else they may feel like saying about it.

This way, if another system user spots the same problem, they immediately see if it has been acknowledged and the comments so far.

This way the workflow of resolving problems with more than one system user can take place in a coordinated way.

Acknowledgment status is also used when defining action operations. You can define, for example, that a notification is sent to a higher level manager only if an event is not acknowledged for some time.

To acknowledge events and comment on them, a user must have at least read permissions to the corresponding triggers. To change problem severity or close problem, a user must have read-write permissions to the corresponding triggers.

There are **several** ways to access the problem update popup window, which allows acknowledging a problem.

- You may select problems in *Monitoring* > *Problems* and then click on *Mass update* below the list
- You can click *Update* in the *Update* column of a problem in:
  - *Dashboards* (*Problems* and *Problems by severity* widgets)
  - *Monitoring* > *Problems*
  - *Monitoring* > *Problems* > *Event details*
- You can click an unresolved problem cell in:
  - *Dashboards* (*Trigger overview* widget)

The popup menu contains an *Update* option that will take you to the problem update window.

Based on acknowledgment information it is possible to configure how the problem count is displayed in the dashboard or maps. To do that, you have to make selections in the *Problem display* option, available in both map configuration and the *Problems by severity* dashboard widget. It is possible to display all problem count, unacknowledged problem count as separated from the total or unacknowledged problem count only.

Based on problem update information (acknowledgment, etc.), it is possible to configure update operations - send a message or execute remote commands.

Problem suppression Problem suppression offers a way of temporarily hiding a problem that can be dealt with later. This is useful for cleaning up the problem list in order to give the highest priority to the most urgent issues. For example, sometimes an issue may arise on the weekend that is not urgent enough to be dealt with immediately, so it can be ”snoozed” until Monday morning.

Problem suppression allows a *single* problem to be hidden, in contrast to suppression through host maintenance, when all problems of the maintenance host are hidden.

Operations for trigger actions will be paused for suppressed problems the same way as it is done with host maintenance.

A problem can be suppressed through the **problem update** window, where suppression is one of the problem update options along with commenting, changing severity, acknowledging, etc.

A problem may also be unsuppressed through the same problem update window.

Display

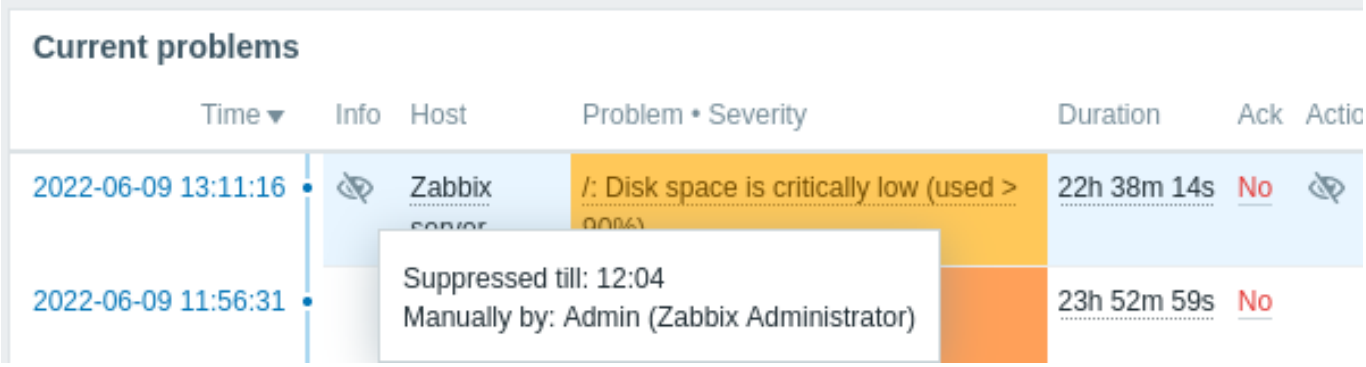
Once suppressed the problem is marked by a blinking  suppression icon in the *Info* column, before being hidden.

If suppression is due to host maintenance, the suppression icon in the *Info* column is shown in orange. If both maintenance and manual suppression apply, manual suppression takes priority in the icon color.

The suppression icon is blinking while the suppression task is in the waiting list. Once the task manager has suppressed the problem, the icon will stop blinking. If the suppression icon keeps blinking for a long time, this may indicate a server problem, for example, if the server is down and the task manager cannot complete the task. The same logic applies to unsuppression. In the short period after the task is submitted and the server has not completed it, the unsuppression icon is blinking.

A suppressed problem may be both hidden or shown, depending on the problem filter/widget settings.

When shown in the problem list, a suppressed problem is marked by the suppression icon and suppression details are shown on mouseover:



Suppression details are also displayed in a popup when positioning the mouse on the suppression icon in the *Actions* column.

Manual closing of problems

Overview While generally problem events are resolved automatically when trigger state goes from `Problem` to `OK`, there may be cases when it is difficult to determine if a problem has been resolved by means of a trigger expression. In such cases, the problem needs to be resolved manually.

For example, *syslog* may report that some kernel parameters need to be tuned for optimal performance. In this case the issue is reported to Linux administrators, they fix it and then close the problem manually.

Problems can be closed manually only for triggers with the *Allow manual close* option enabled.

When a problem is "manually closed", Zabbix generates a new internal task for Zabbix server. Then the *task manager* process executes this task and generates an OK event, therefore closing problem event.

The generated *OK* event includes the full set of event tags resolved for that event (including tags inherited from templates, hosts and triggers). These tags are available in notifications and in macros such as `{EVENT.RECOVERY.TAGS}` and `{EVENT.RECOVERY.TAGSJSON}`.

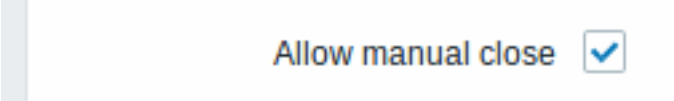
A manually closed problem does not mean that the underlying trigger will never go into a *Problem* state again. The trigger expression is re-evaluated and may result in a problem:

- When new data arrive for any item included in the trigger expression (note that the values discarded by a throttling preprocessing step are not considered as received and will not cause trigger expression to be re-evaluated);
- When date and time and/or nodata() functions are used in the expression.

Configuration Two steps are required to close a problem manually.

Trigger configuration

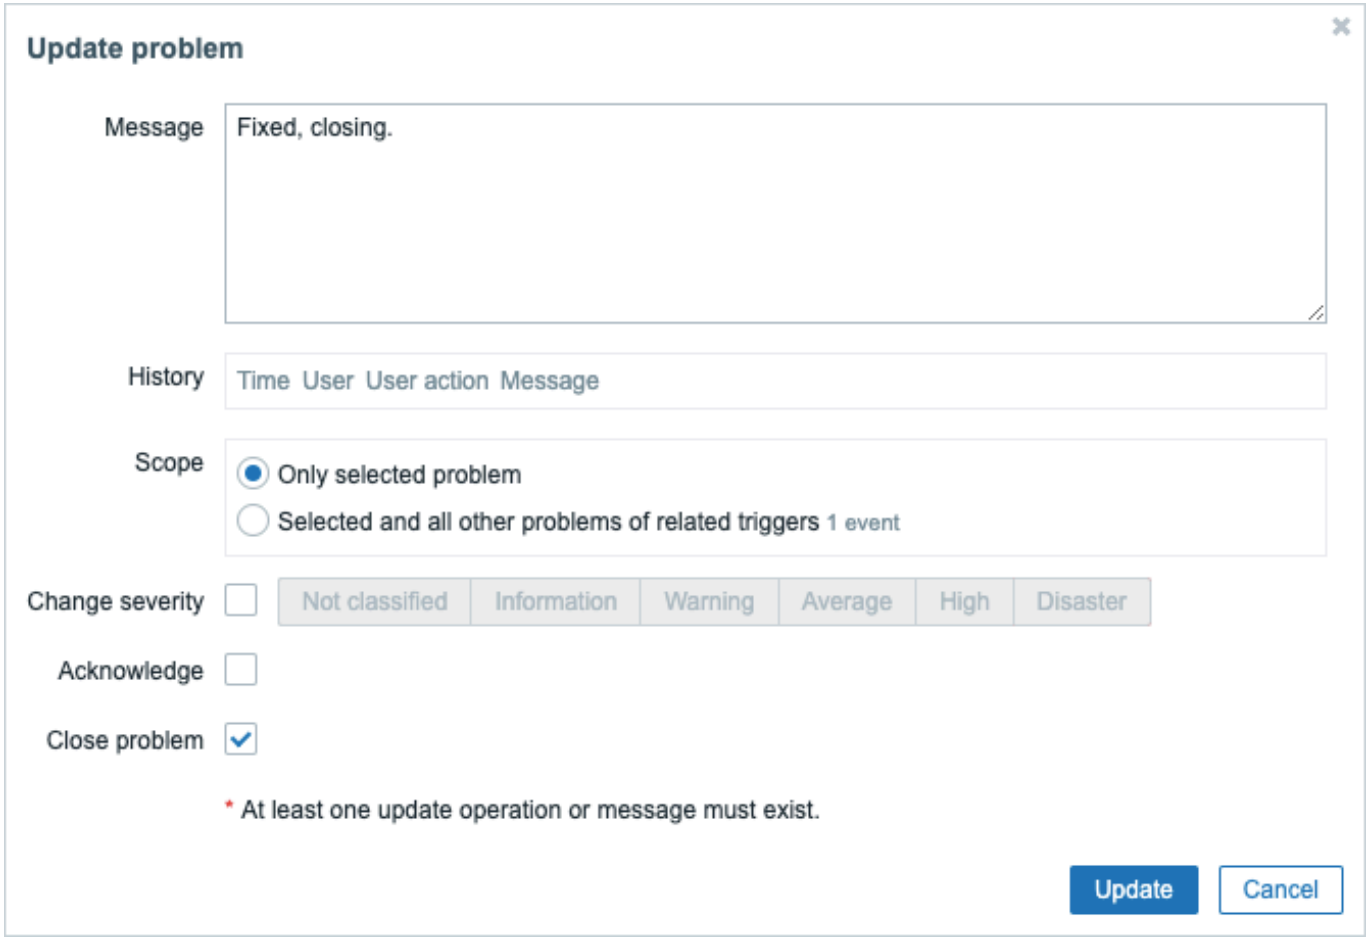
In trigger configuration, enable the *Allow manual close* option.



Problem update window

If a problem arises for a trigger with the *Manual close* flag, you can open the problem update popup window of that problem and close the problem manually.

To close the problem, check the *Close problem* option in the form and click *Update*.



All mandatory input fields are marked with a red asterisk.

The request is processed by Zabbix server. Normally it will take a few seconds to close the problem. During that process *CLOSING* is displayed in *Monitoring* > *Problems* as the status of the problem.

Verification It can be verified that a problem has been closed manually:

- in event details, available through *Monitoring* > *Problems*;
- by using the `{EVENT.UPDATE.HISTORY}` macro in notification messages that will provide this information.

Maintenance periods

Overview

Maintenance is used to suppress problems during predefined time periods.

You can define maintenance periods for these entities:

- Host groups
- Hosts
- Triggers
- Event names
- Tags

If maintenance is defined for specific tags or event names then maintenance will be activated only for the corresponding triggers; all other triggers of the host or host group will not be in maintenance.

There are two maintenance types: *With data collection* and *No data collection*.

During *With data collection* maintenance, triggers are processed as usual and events are created when required. However, if the *Pause operations for suppressed problems* option is checked in action configuration, problem escalations are paused for entities in maintenance. In this case, escalation steps that may include sending notifications or remote commands will be ignored for as long as the maintenance period lasts. Note that problem recovery and update operations are not suppressed during maintenance, only escalations. Recovery notifications are not sent if the problem started during maintenance.

For example, if escalation steps are scheduled at 0, 30 and 60 minutes after a problem start, and there is a half-hour long maintenance lasting from 10 minutes to 40 minutes after a real problem arises, steps two and three will be executed a half-hour later, or at 60 minutes and 90 minutes (providing the problem still exists). Similarly, if a problem arises during the maintenance, the escalation will start after the maintenance.

To receive problem notifications during the maintenance normally (without delay), you have to uncheck the *Pause operations for suppressed problems* option in action configuration.

Note:
If at least one host (used in the trigger expression) is not in maintenance mode, Zabbix will send a problem notification.

Zabbix server must be running during maintenance. Maintenances are recalculated every minute or as soon as the configuration cache is reloaded if there are changes to the maintenance period.

Timer processes check if host status must be changed to/from maintenance at 0 seconds of every minute. Additionally, every second the timer process checks if any maintenances must be started/stopped based on whether there are changes to the maintenance periods after the configuration update. Thus the speed of starting/stopping maintenance periods depends on the configuration update interval (10 seconds by default). Note that maintenance period changes do not include *Active since/Active till* settings. Also, if an entity (host group/host/trigger/etc) is added to an existing active maintenance period, the changes will only be activated by the timer process at the start of next minute.

Note that when a host enters maintenance, Zabbix server timer processes will read all open problems to check if it is required to suppress those. This may have a performance impact if there are many open problems. Zabbix server will also read all open problems upon startup, even if there are no maintenances configured at the time.

Note that the Zabbix server (or proxy) always collects data regardless of the maintenance type (including *No data collection* maintenance). The data is later ignored by the server if *No data collection* is set.

When *No data collection* maintenance ends, triggers using `nodata()` function will not fire before the next check during the period they are checking.

If a log item is added while a host is in maintenance and the maintenance ends, only new logfile entries since the end of the maintenance will be gathered.

If a timestamped value is sent for a host that is in *No data collection* maintenance (e.g., using Zabbix sender), then this value will be dropped; however it is possible to send a timestamped value in for an expired maintenance period, and it will be accepted.

If maintenance period or maintenance option (host group, host, trigger, etc) are changed by the user, the changes will only take effect after configuration cache synchronization.

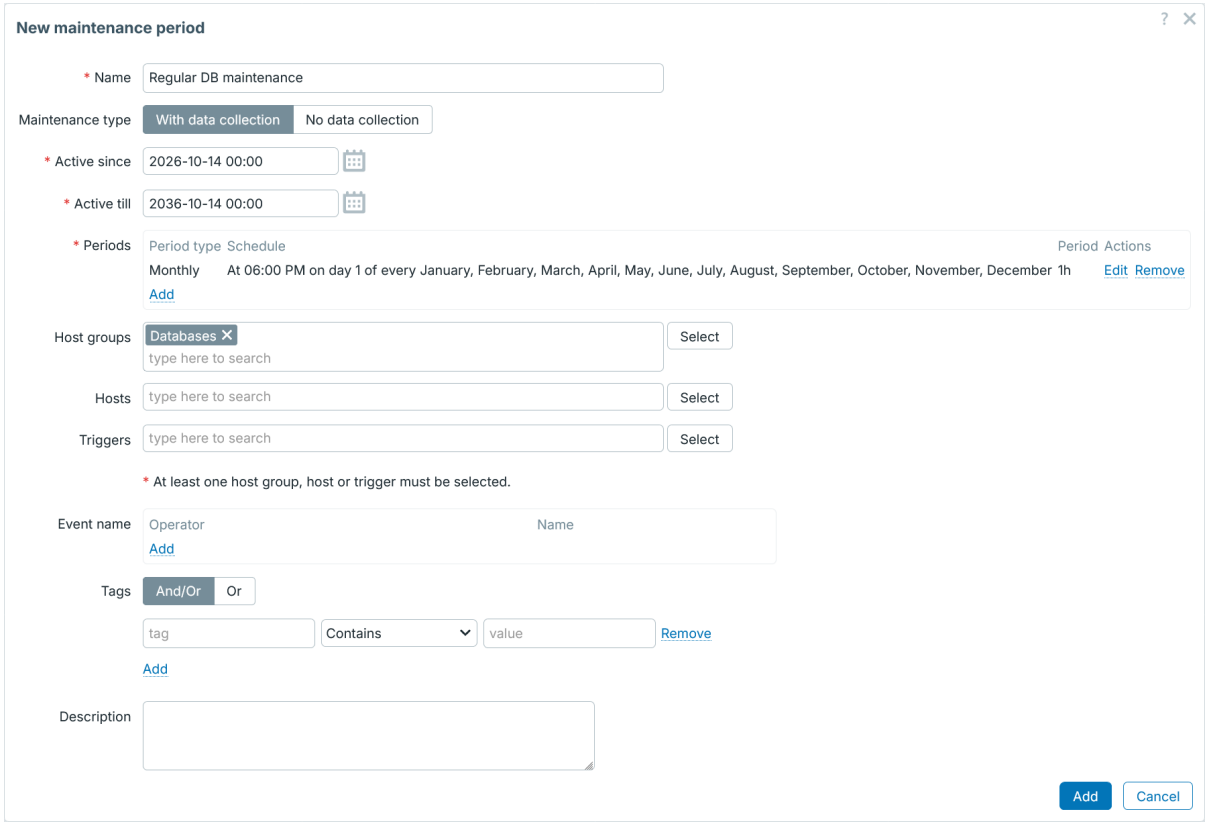
Configuration

To start a maintenance period immediately:

1. Go to *Monitoring* > *Problems*.
2. Click on the problem name.
3. Select maintenance option from the event menu (host, trigger, event name, tags).
4. Modify maintenance parameters in the form.

To configure a planned maintenance period:

1. Go to *Data collection* > *Maintenance*.
2. Click *Create maintenance period* (or the name of an existing maintenance period).
3. Enter maintenance parameters in the form.



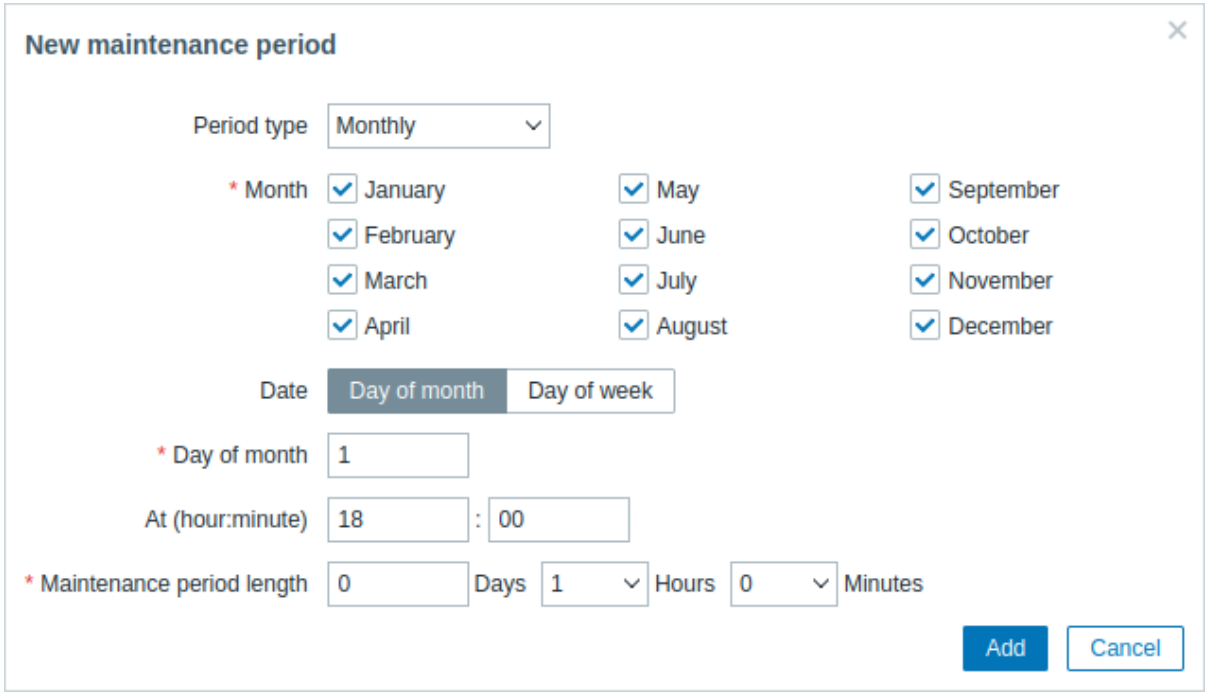
All mandatory input fields are marked with a red asterisk.

Parameter	Description
Name	Name of the maintenance period.
Maintenance type	Two types of maintenance can be set: **With data collection** - data is collected by the server during maintenance, and triggers are processed; **No data collection** - data may still be collected, but it is not saved to the database during maintenance, and triggers (including the `nodata()` function) do not fire. See Effect of maintenance periods for how each type affects the availability report.
Active since	The date and time when executing maintenance periods becomes active. *Note:* Setting this time alone does not activate a maintenance period; maintenance periods must be configured in *Periods* (see below).
Active till	The date and time when executing maintenance periods stops being active.
Periods	This block allows you to define the exact days and hours when the maintenance takes place. Clicking Add opens a popup window with a flexible *Maintenance period* form where you can define maintenance schedule. See Maintenance periods for a detailed description.
Host groups	Select host groups that the maintenance will be activated for. The maintenance will be activated for all hosts from the specified host group(s). This field is auto-complete, so starting to type in it will display a dropdown of all available host groups. Specifying a parent host group implicitly selects all nested host groups. Thus the maintenance will also be activated on hosts from nested groups.
Hosts	Select hosts that the maintenance will be activated for. This field is auto-complete, so starting to type in it will display a dropdown of all available hosts.
Triggers	Select triggers that the maintenance will be activated for. This field is auto-complete, so starting to type in it will display a dropdown of all available triggers.
Event name	Specify event names that the maintenance will be activated for. Several conditions can be set. There are two operators available for each condition: **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive) **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive)

Parameter	Description
Tags	Specify tags that the maintenance will be activated for. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-sensitive); **Equals** - include the specified tag names and values (case-sensitive). **Does not equal** - exclude the specified tag names and values (case-sensitive) **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive) There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met. Tags can be specified only if the *With data collection* maintenance type is selected.
Description	Description of maintenance period.

Maintenance periods

The maintenance period window is for scheduling time for a recurring or a one-time maintenance. The form is dynamic with available fields changing based on the *Period type* selected.



Period type	Description
One time only	Configure a one time only maintenance period: *Date* - date and time of the maintenance period; *Maintenance period length* - for how long the maintenance will be active.

Period type	Description
Daily	Configure a daily maintenance period: *Every day(s)* - maintenance frequency (1 - *(default)* every day, 2 - every two days, etc.); *At (hour:minute)* - time of the day when maintenance starts; *Maintenance period length* - for how long the maintenance will be active. When *Every day(s)* parameter is greater than ”1”, the starting day is the day that the *Active since* time falls on. Examples: - if *Active since* is set to ”2021-01-01 12:00”, *Every day(s)* is set to ”2”, and *At (hour:minute)* is set to ”23:00”, then the first maintenance period will start on January 1 at 23:00, while the second maintenance period will start on January 3 at 23:00; - if *Active since* is set to ”2021-01-01 12:00”, *Every day(s)* is set to ”2”, and *At (hour:minute)* is set to ”01:00”, then the first maintenance period will start on January 3 at 01:00, while the second maintenance period will start on January 5 at 01:00.
Weekly	Configure a weekly maintenance period: *Every week(s)* - maintenance frequency (1 - *(default)* every week, 2 - every two weeks, etc.); *Day of week* - on which day the maintenance should take place; *At (hour:minute)* - time of the day when maintenance starts; *Maintenance period length* - for how long the maintenance will be active. When *Every week(s)* parameter is greater than ”1”, the starting week is the week that the *Active since* time falls on. For examples, see parameter *Daily* description above.
Monthly	Configure a monthly maintenance period: *Month* - select all months during which the regular maintenance is carried out; *Date:* ***Day of month*** - select this option if the maintenance should take place on the same date each month (for example, every 1st day of the month), and then select the required day in the field *Day of month* that appears; *Date:* ***Day of week*** - select this option if the maintenance should take place only on certain days (for example, every first Monday of the month), then select (in the drop-down) the required week of the month (first, second, third, fourth, or last), and then mark the checkboxes for maintenance day(s); *At (hour:minute)* - time of the day when maintenance starts; *Maintenance period length* - for how long the maintenance will be active.

Attention:
When creating a maintenance period, the time zone of the user who creates it is used. However, when recurring maintenance periods (*Daily*, *Weekly*, *Monthly*) are scheduled, the time zone of the Zabbix server is used. To ensure predictable behavior of recurring maintenance periods, it is required to use a common time zone for all parts of Zabbix.

When done, click *Add* to add the maintenance period to the *Periods* section.

Note that Daylight Saving Time (DST) changes do not affect how long the maintenance will be. For example, let’s say that we have a two-hour maintenance configured that usually starts at 01:00 and finishes at 03:00:

- if after one hour of maintenance (at 02:00) a DST change happens and current time changes from 02:00 to 03:00, the maintenance will continue for one more hour (till 04:00);
- if after two hours of maintenance (at 03:00) a DST change happens and current time changes from 03:00 to 02:00, the maintenance will stop, because two hours have passed;
- if a maintenance period starts during the hour that is skipped by a DST change, then the maintenance will not start.

If a maintenance period is set to ”1 day” (the actual period of the maintenance is 24 hours, since Zabbix calculates days in hours), starts at 00:00 and finishes at 00:00 the next day:

- the maintenance will stop at 01:00 the next day if current time changes forward one hour;
- the maintenance will stop at 23:00 that day if current time changes back one hour.

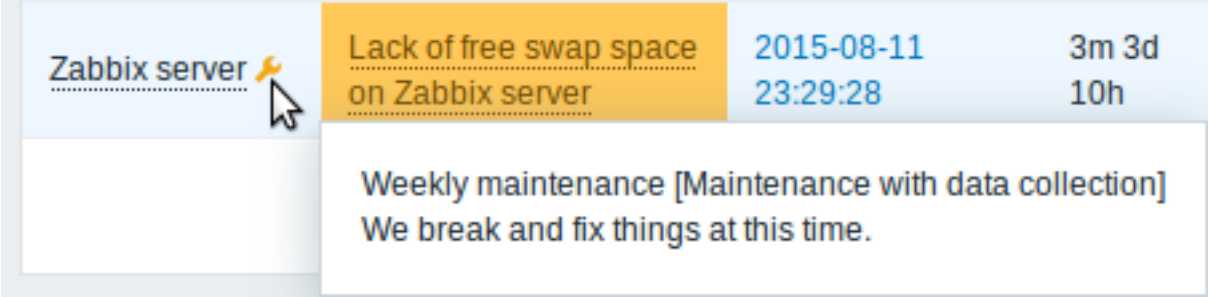
Display

Displaying hosts in maintenance

An orange wrench icon  next to the host name indicates that this host is in maintenance in:

- *Dashboards*
- *Monitoring > Problems*

- *Inventory* > *Hosts* > *Host inventory details*
- *Data collection* > *Hosts* (See *Status* column)



Maintenance details are displayed when the mouse pointer is positioned over the icon.

Additionally, hosts in maintenance get an orange background in *Monitoring* > *Maps*.

Displaying suppressed problems

Normally, problems for hosts in maintenance are suppressed, that is, not displayed in the frontend.

Suppressed problems may also be hidden because they were manually suppressed.

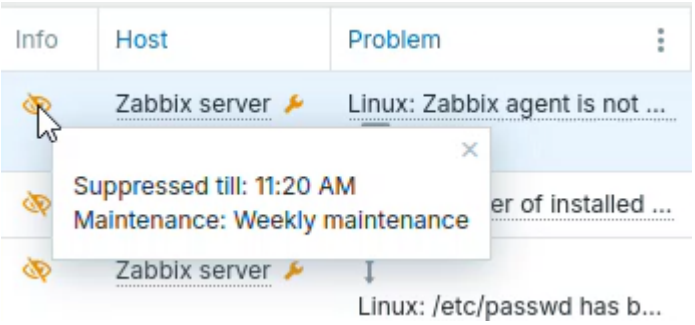
It is possible to display suppressed problems by selecting the *Show suppressed problems* option in these locations:

- *Dashboards* (in *Problem hosts*, *Problems*, *Problems by severity*, *Trigger overview* widget configuration)
- *Monitoring* > *Problems* (in the filter)
- *Monitoring* > *Maps* (in map configuration)
- Global notifications (in user profile configuration)

When suppressed problems are displayed, the following icon is shown: . If suppression is due to host maintenance, the following icons are shown:

- - problem has been suppressed due to host maintenance.
- - problem has been unsuppressed after host maintenance.

Hovering over the icon displays more details. In event details, you can check whether the event was suppressed because of maintenance or suppressed manually. If both maintenance and manual suppression apply, manual suppression takes priority in the icon color.



Calculation of queues during maintenance

Queues displayed in the Zabbix frontend (Administration > Queue) are calculated by the Zabbix server. They do not include items in *No data collection* maintenance, and queue length is always zero for these items, even when their values are delayed. Delayed items in *With data collection* maintenance are still counted in the queue.

The Zabbix proxy is not aware of maintenance periods because there is no synchronization of maintenance configuration between the Zabbix server and proxy. Internal checks calculated on Zabbix proxies (for example, `zabbix[queue,,]` and `zabbix[stats,,,queue,,]`) report delayed items regardless of the maintenance status on the Zabbix server.

As a result, different queue lengths may be reported for the same items in *No data collection* maintenance by the Zabbix frontend and by internal checks on Zabbix proxies.

Tagging

Tags have multiple purposes, most notably, to mark events. When entities are tagged, any new event related to a tagged entity will inherit its tags. For example:

- with tagged templates - any host problem (created by triggers from the template) will inherit the template tags.
- with tagged hosts - any host problem will inherit the host tags.
- with tagged items/web scenarios - any item/web scenario problem will inherit the item/web scenario tags.
- with tagged triggers - any problem created by the trigger will inherit the trigger tags.

A problem event inherits all tags from the whole chain of entities - templates, hosts, items/web scenarios, triggers. Identical `tag:value` combinations (after resolved macros) are merged into one, thus avoiding duplication.

Recovery events generated by a manual close also include the resolved event tags inherited from templates, hosts, items/web scenarios and triggers. These tags are available in notifications and via macros such as `{EVENT.RECOVERY.TAGS}` and `{EVENT.RECOVERY.TAGSJSON}`.

Custom event tags offer more flexibility. For example:

- event correlation can be configured based on event tags.
- action conditions can be configured based on event tags.
- item problems can be grouped based on event tags.
- problem tags can be used to map problems to services.

Entities may be tagged with the same tag name but different tag values (for example, `component:memory` and `component:storage`). Similarly, an entity can have a tag without a value and the same tag with a value (for example, `database` and `database:postgresql`). Such tags are not considered duplicates.

Tags consist of a tag name and a tag value. When tagging entities, you can use just the name or pair it with a value (for example, `mysql`, `jira`, `target:mysql`, `service:jira`, etc.).

Tags can be defined for various entities:

- Templates
- Hosts
- Items
- Web scenarios
- Triggers
- Services
- Template items and triggers
- Host, item, and trigger prototypes

In the *Templates*, *Hosts*, *Items*, *Triggers*, *Web scenarios* lists and their prototypes the *Tags* column displays both regular and inherited tags. When hovering over or clicking an *inherited* tag a JavaScript tooltip appears with the text "Inherited tag". If an inherited tag also exists as a regular tag, a different tooltip text is shown depending on the list where it appears (for example, "Inherited and template tag" in the templates list or "Inherited and host tag" in the hosts list). In other places of the frontend regular and inherited tags are displayed together (duplicates are removed), without icons and without additional tooltip text.

> **Note:**
> Refer to the official Zabbix guidelines for general recommendations on defining tags, along with specific guidance for templates, items, triggers, and low-level discovery rules.

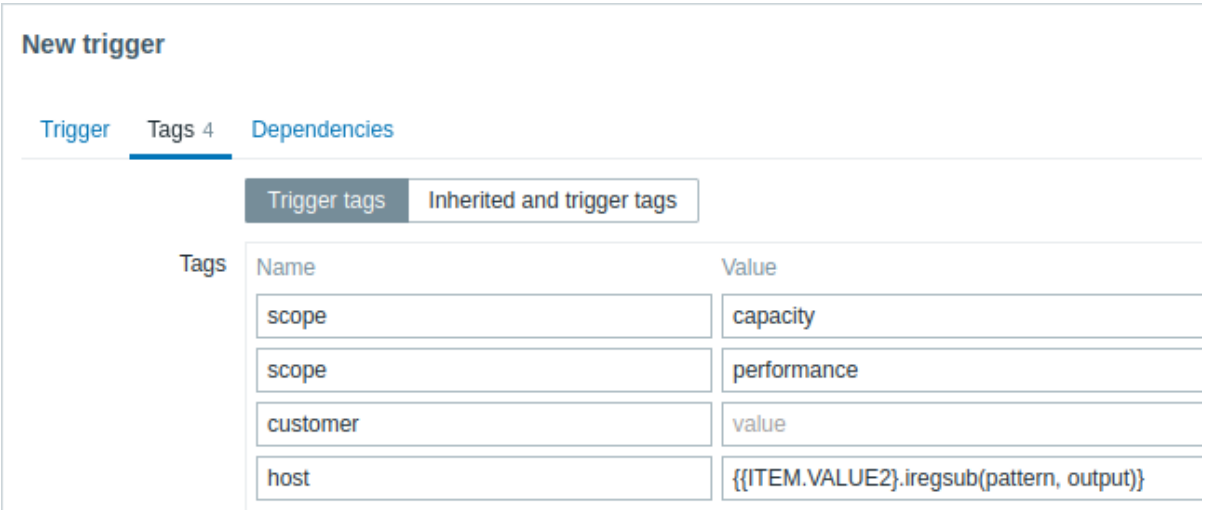
Use cases Some common use cases for tagging are as follows:

1. Mark trigger events:
   - Define a trigger tag (for example, `scope:performance`).
   - Problems created by this trigger will have the trigger tag.
2. Mark template-inherited problems:
   - Define a template tag (for example, `target:mysql`).
   - Problems created by triggers from this template will have the template tag.
3. Mark host problems:
   - Define a host tag (for example, `service:jira`).
   - Problems created by triggers from this host will have the host tag.
4. Filter related items:
   - Define an item tag (for example, `component:cpu`).
   - In *Monitoring* > *Latest data*, items can be filtered by the `component:cpu` tag.
5. Use information extracted from the item value as the tag value:
   - Define a tag with a macro as the tag value (for example, `tag-name:{{ITEM.VALUE<N>}.regsub()}`).
   - In *Monitoring* > *Problems*, problems will have the tag value resolved to the data extracted from the item value.
6. Identify problems in a log file and close them separately:

- Define a trigger tag for the log monitoring item trigger that will extract values from the item value using a macro (for example, `service:{{ITEM.VALUE<N>}.regsub()}`).
- In the trigger configuration, set up event correlation:
  - set *PROBLEM event generation mode* to "Multiple";
  - set *OK event closes* to "All problems if tag values match";
  - set the tag for matching.
- Problems created by the log item trigger will have the trigger tag and will be closed individually.

7. Filter notifications:
   - Define trigger tags (for example, `scope:security` for trigger1 and `scope:availability` for trigger2).
   - Use tag filtering in action conditions to receive notifications only on the events that match tag data.
8. Identify problems in notifications:
   - Define trigger tags.
   - Use the {EVENT.TAGS} macro in the problem notification.
   - The problem notification will contain the trigger tags, making it easier to identify which application/service the notification belongs to.
9. Simplify configuration tasks by using template tags:
   - Define a template trigger tag.
   - Triggers created from this template trigger will have its tag.
10. Create triggers with tags from low-level discovery (LLD):
    - Define a trigger prototype tag with an LLD macro in the tag name or value (for example, `scope:{#FSNAME}`).
    - Triggers created from the trigger prototype will have its tag.
11. Match services using service tags:
    - Define service tags.
    - Configure service actions for services with matching tags.
    - Additionally, use service tags to link a service to an SLA for SLA calculations.
12. Link services to problems using service problem tags:
    - Define a problem tag in service configuration (for example, `target:mysql`).
    - Problems with a matching tag will be automatically correlated to the service, and service status will change based on the configured service status calculation rules.
13. Suppress problems when a host is in maintenance mode:
    - Define tags in maintenance period configuration.
    - Problems with the defined tags will be suppressed.
14. Grant access to user groups:
    - Define tags in user group configuration.
    - Users in the user group will be able to view only problems with the defined tags.

Configuration Tags can be defined in a dedicated tab, for example, in trigger configuration:



Macro support Built-in and user macros in tags are resolved at the time of the event. Until the event has occurred, these macros will be shown in Zabbix frontend unresolved.

Low-level discovery macros are resolved during discovery process.

The following macros may be used in trigger tag names and values:

- {ITEM.VALUE}, {ITEM.VALUE.AGE}, {ITEM.VALUE.DATE}, {ITEM.VALUE.TIME}, {ITEM.VALUE.TIMESTAMP}, {ITEM.LASTVALUE}, {ITEM.LASTVALUE.AGE}, {ITEM.LASTVALUE.DATE}, {ITEM.LASTVALUE.TIME}, {ITEM.LASTVALUE.TIMESTAMP}, {HOST.HOST}, {HOST.NAME}, {HOST.CONN}, {HOST.DNS}, {HOST.IP}, {HOST.PORT}, and {HOST.ID} built-in macros

- {INVENTORY.*} built-in macros (for referencing host inventory values from one or several hosts in a trigger expression)
- User macros and user macros with context (the context may include low-level discovery macros)
- Low-level discovery macros (only in trigger prototype tags)

The following macros may be used in template, host, and item/web scenario tag names and values:

- {HOST.HOST}, {HOST.NAME}, {HOST.CONN}, {HOST.DNS}, {HOST.IP}, {HOST.PORT} and {HOST.ID} built-in macros
- {INVENTORY.*} built-in macros
- User macros
- Low-level discovery macros (only in host and item prototype tags)

The following macros may be used in trigger-based notifications:

- {EVENT.TAGS} and {EVENT.RECOVERY.TAGS} built-in macros (these macros will resolve to a comma-separated list of event tags or recovery event tags)
- {EVENT.TAGSJSON} and {EVENT.RECOVERY.TAGSJSON} built-in macros (these macros will resolve to a JSON array containing event tag objects or recovery event tag objects)

Substring extraction in trigger tags

Substring extraction is supported for populating the tag name or tag value, using a macro function. The function applies a regular expression to the value obtained by the supported macro. For example:

```
{{ITEM.VALUE}.regsub(pattern, output)}
{{ITEM.VALUE}.iregsub(pattern, output)}

{{#LLDMACRO}.regsub(pattern, output)}
{{#LLDMACRO}.iregsub(pattern, output)}
```

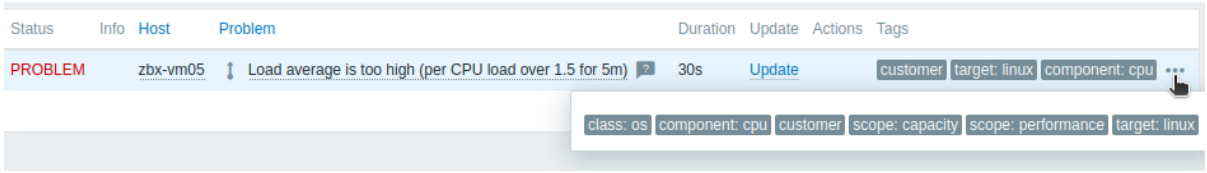
If the tag name or value exceeds 255 characters after macro resolution, it will be truncated to 255 characters.

See also: Using macro functions in low-level discovery macros for event tagging.

Viewing event tags Tags, if defined, can be viewed with new events in:

- *Monitoring* > *Problems*
- *Monitoring* > *Problems* > *Event details*
- *Dashboards* > *Problems* widget

The order and number of displayed tags is determined by the *Tag display priority* and *Show tags* filtering options in *Monitoring* > *Problems* or the *Problems* dashboard widget. Note that a maximum of three tags can be displayed; if there are more tags, hovering over the three dots reveals all tags in a pop-up window.



Complex event processing

Overview Complex event processing (CEP) performs rule-based processing for Zabbix events.

For example, you might want to:

- group related events and mark first event in a series as cause (the following as symptoms)
- specify a time-to-live period for events
- specify a maximum number for certain events
- be notified about similar events by location
- be notified if some event happens too often (or not at all) within a time period

In addition to single event processing, you may define **time windows** for event analysis. Within a time window, all matching events form a single context and can be analyzed as a whole.

A sliding time window moves forward with the time. Older events leave the window (are "evicted"), while new events may enter.

A tumbling time window does not move forward. When the window expires, a new window with the same duration is started. *Cause and symptom grouping* windows are a tumbling time window.

The time window concept presents logical points for executing operations. For example, an event can be transformed or "enriched" when it enters the window - severity can be raised or lowered, tags can be added or adjusted, suppression can be activated, etc.

Similar operations can be executed when the event leaves the window or when a pattern is found.

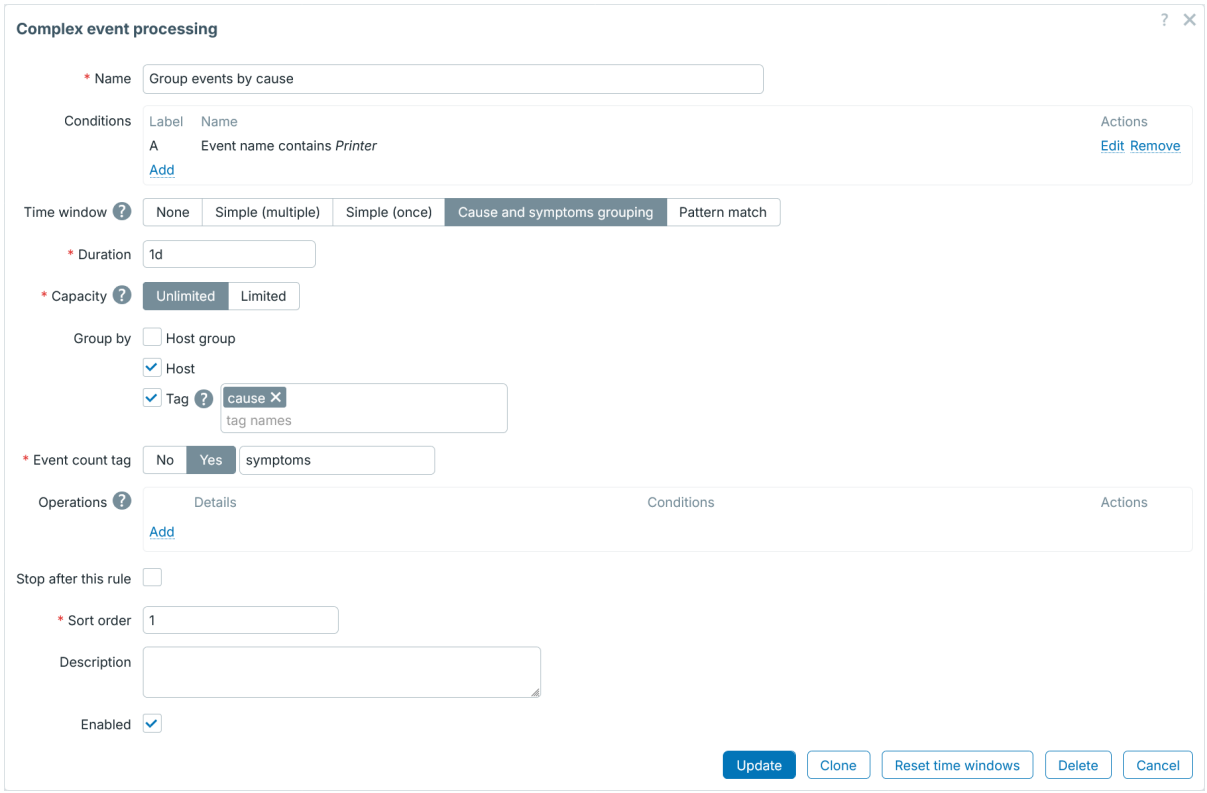
Event stream analysis is defined by CEP rules. A CEP rule consists of a time window (optional), conditions and operations. Operations are tied to an execution point.

Use cases

- Group related events, with the first such event automatically marked as cause (all following as symptoms).
- Enrich event with higher severity/specific tag if too many 401 events happen within a minute (possible bot activity).

Configuration To create a CEP rule in Zabbix frontend, do the following:

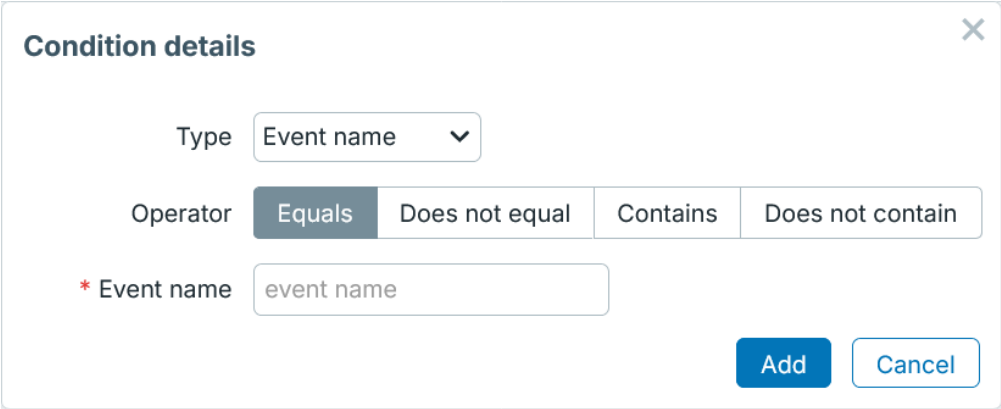
1. Go to *Data collection* > *Event processing*.
2. Click *Create complex event processing* in the upper-right corner of the screen.
3. Enter parameters of the rule in the form.



Parameter	Description
Name	Specify a unique CEP rule name.
Conditions	Specify conditions for associating events with this CEP rule.
Time window	Select the time window type: **None** - process single events with no time window (default) **Simple (multiple)** - process events in a sliding time window (event can be added to multiple windows of this type) **Simple (once)** - process events in a sliding time window (event can be added only to one window of this type) **Cause and symptom grouping** - process first event in time window as cause, consecutive events as symptoms **Pattern match** - analyze sliding time window events based on a JavaScript script. The window is examined every second; events older than the duration are evicted, and the script is then run over the events that remain in it. When the script reports a match, the pattern match operations are executed for every event the window holds; the events stay in the window and are examined again a second later, until they age out. A window can be forcibly closed by the "Close window" operation.

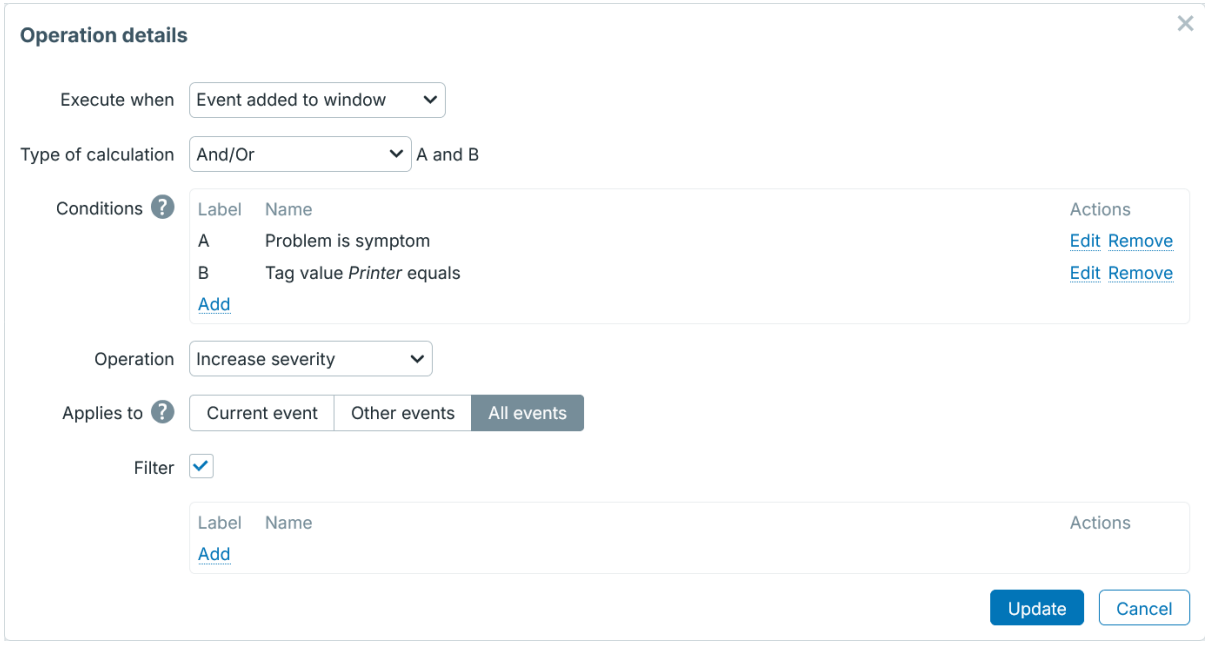
Parameter	Description
Duration	Limit the time window duration (in seconds, 365 days max). The window duration is started by the first matching event. In *Simple* and *Pattern match* windows, the window is moving so older events will leave the window (are evicted) after the duration passes. In *Cause and symptom grouping* windows, window duration is a fixed timeframe for setting cause and symptom relationship. The window is started when the first event matches. When the duration expires, the window is closed and a new one is started. No event is evicted when the duration expires; periods run back to back. Time suffixes are supported (`10m`, `1h`, etc). User macros are supported.
Capacity	Specify the event capacity within the time window: **Unlimited** - event count is not limited **Limited** - event count is limited by the specified number Once the limit is surpassed, the latest event is evicted. User macros are supported.
Group by	Group events by host group, host, and/or event tag. This option is available for CEP rules with a time window. If specified, the rule will create multiple windows based on specified criteria; otherwise events will be grouped into a single window group.
Event count tag	Specify name for a tag displaying symptom event count. Symptom event count will be displayed as `<tag name>:<event count>` (for example: `symptoms:2`). This field is only available if *Cause and symptom grouping* is selected as time window type.
Script	Specify the script for identifying event patterns. This field is only available if *Pattern match* is selected as time window type.
Operations	Click on *Add* to specify operations.
Stop after this rule	Mark the checkbox if no further processing is required after this rule. Otherwise, continue with the next CEP rule according to the sort order.
Sort order	Specify position of this CEP rule in the processing order. If the order priority is the same, the rules will be executed alphabetically.
Description	Add the CEP rule description.

CEP rule condition



Parameter	Description
Type	Select condition type. Event filtering can be based on event name, tag name, tag value, severity, host group, host (visible name), or time period.

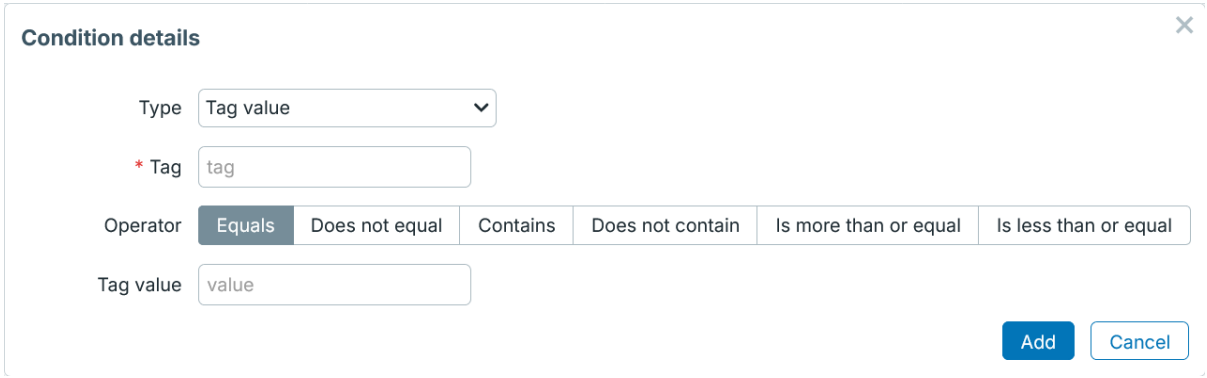
CEP rule operation



Parameter	Description
Execute when	Select execution point: **Event occurred** - when event occurs. This is the only stage where *Discard* operation is supported, since event is not actually created yet and can be discarded. Window-related operations are not available at this point. Changes made to the event in *Event occurred* phase (except closing) are visible to the next CEP rules trying to match this event. For example, if rule A added tag X and rule B has condition "Tag X exists", then it will match rule A and B. Of course, rule A should not be configured to stop after matching. **Event added to window** - when event is added to time window. Operations are executed after event is added to matched CEP-rule windows. If window is at its capacity, then *Event evicted* operations are executed instead. **Event evicted** - when event leaves the window. Operations are executed before event is evicted from a window because of duration or capacity. **Event closed** - when event is closed. **Window closed** - when time window is closed. Operations are executed on all window events when window is closed because of *Close window* operation or as expired cause-symptom window. **Pattern matched** - when pattern is matched. Operations are executed on all window events when the pattern match JavaScript returns `true`. Note that the *Event occurred*, *Event added to window* and *Event evicted* have fixed execution order, while *Window closed* and *Pattern matched* operate independently.
Conditions	Select additional conditions for the operation.

Parameter	Description
Operation	Select operation. **Events**: *Set name* - set event name *Close* - close event *Discard* - discard event. Only available at *Event occurred* stage. *Set severity* - set event severity *Increase severity* - increase event severity *Decrease severity* - decrease event severity *Suppress* - suppress event *Unsuppress* - unsuppress event *Clone first* - clone first event. Useful for producing a summary event based on the first event. Only available at *Event evicted*, *Window closed*, or *Pattern matched* stages. *Clone last* - clone last event. Useful for producing a summary event based on the last event. Only available at *Event evicted*, *Window closed*, or *Pattern matched* stages. For event cloning to work either the source (first/last) event must be resolved or its trigger must have multi-event generation enabled. **Tags**: *Add* - add tag *Set* - set tag and value *Set value* - set tag value *Increase value* - increase tag value *Decrease value* - decrease tag value *Rename* - rename tag *Remove* - remove tag **Window**: *Close (window)* - close event window
Applies to	Select operation scope: **Current event** - apply operation to the current event **Other events** - apply operation to all events except the current one **All events** - apply operation to all events
Filter	Select additional filter for the operation scope. This field is available if *Applies to* is set "Other events" or "All events".

CEP rule operation condition



Parameter	Description
Type	Select condition type. Event filtering can be based on tag name, tag value, problem is opened/closed, problem is symptom/not symptom, first event in time window, last event in time window, problem is suppressed/not suppressed, event is cloned/not cloned.

Event correlation

Overview

Event correlation allows problem events to be correlated with their resolution in a manner that is very precise and flexible.

Event correlation can be defined:

- On trigger level - one trigger may be used to relate separate problems to their solution;
- Globally - problems can be correlated to their solution from a different trigger/polling method using global correlation rules.

Trigger-based event correlation

Overview

Trigger-based event correlation allows separate problems reported by one trigger to be correlated.

While generally an OK event can close all problem events created by one trigger, there are cases when a more detailed approach is needed. For example, when monitoring log files you may want to discover certain problems in a log file and close them individually rather than all together.

This is the case with triggers that have *PROBLEM event generation mode* parameter set to *Multiple*. Such triggers are normally used for log monitoring, trap processing, etc.

It is possible in Zabbix to relate problem events based on tagging. Tags are used to extract values and create identification for problem events. Taking advantage of that, problems can also be closed individually based on matching tag.

In other words, the same trigger can create separate events identified by the event tag. Therefore problem events can be identified one-by-one and closed separately based on the identification by the event tag.

How it works

In log monitoring you may encounter lines similar to these:

```
Line1: Service 1 stopped
Line2: Service 2 stopped
Line3: Service 1 was restarted
Line4: Service 2 was restarted
```

The idea of event correlation is to be able to match the problem event from Line1 to the resolution from Line3 and the problem event from Line2 to the resolution from Line4, and close these problems one by one:

```
Line1: Service 1 stopped
Line3: Service 1 was restarted #problem from Line 1 closed

Line2: Service 2 stopped
Line4: Service 2 was restarted #problem from Line 2 closed
```

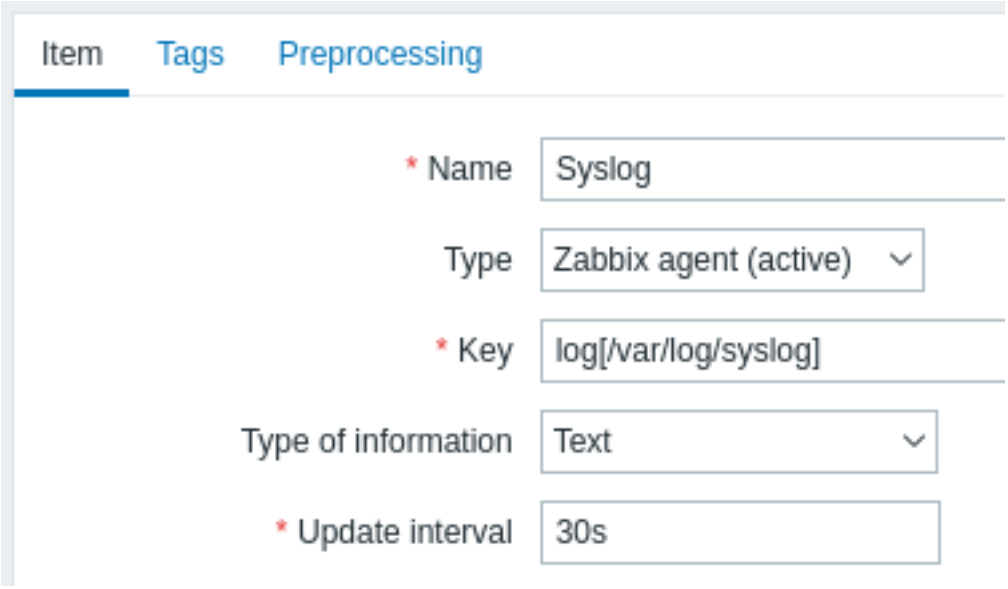
To do this you need to tag these related events as, for example, ”Service 1” and ”Service 2”. That can be done by applying a regular expression to the log line to extract the tag value. Then, when events are created, they are tagged ”Service 1” and ”Service 2” respectively and problem can be matched to the resolution.

Configuration

Item

To begin with, you may want to set up an item that monitors a log file, for example:

```
log[/var/log/syslog]
```



With the item set up, wait a minute for the configuration changes to be picked up and then go to Latest data to make sure that the item has started collecting data.

Trigger

With the item working you need to configure the trigger. It's important to decide what entries in the log file are worth paying attention to. For example, the following trigger expression will search for a string like 'Stopping' to signal potential problems:

```
find(/My host/log[/var/log/syslog],,"regexp","Stopping")=1
```

Attention:
To make sure that each line containing the string "Stopping" is considered a problem also set the *Problem event generation mode* in trigger configuration to 'Multiple'.

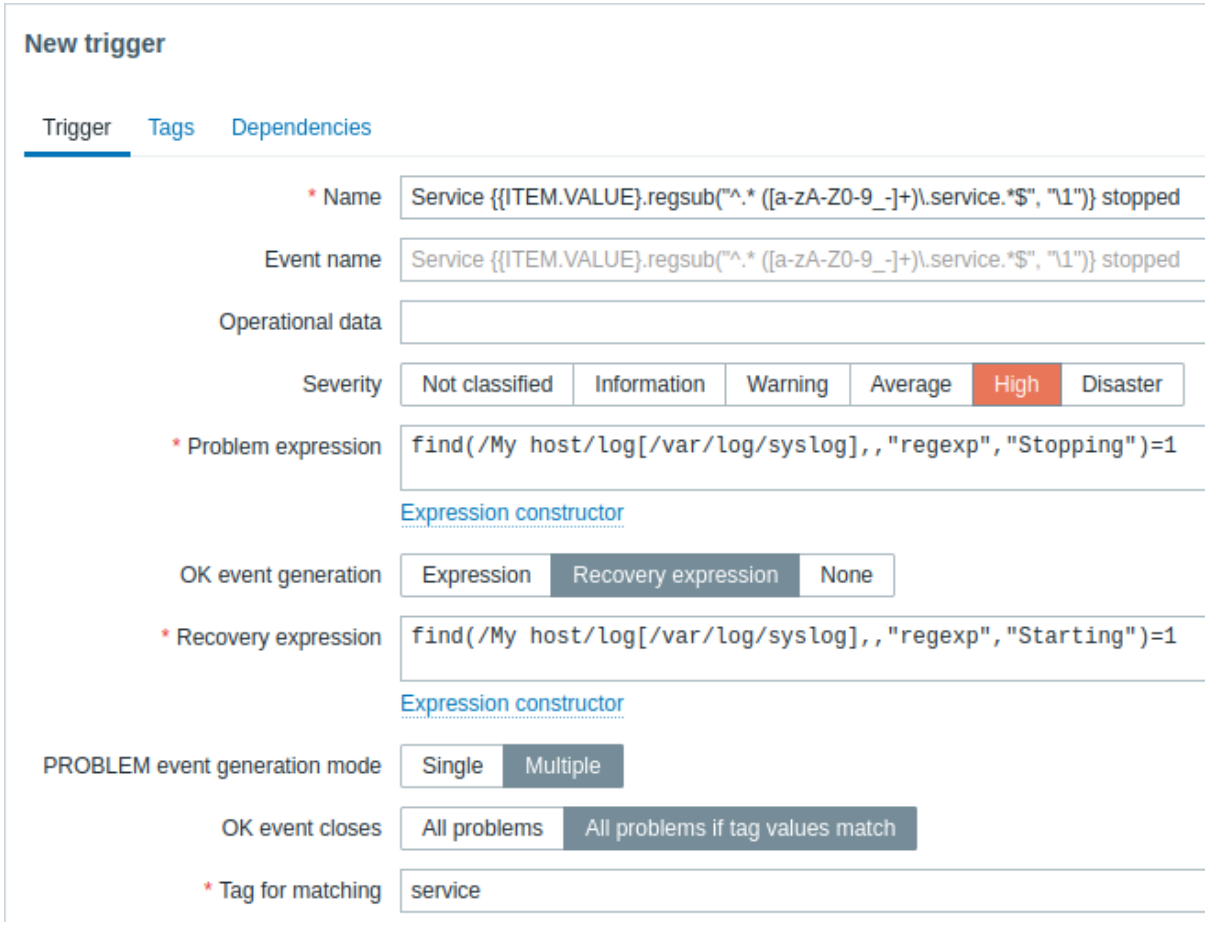
Then define a recovery expression. The following recovery expression will resolve all problems if a log line is found containing the string "Starting":

```
find(/My host/log[/var/log/syslog],,"regexp","Starting")=1
```

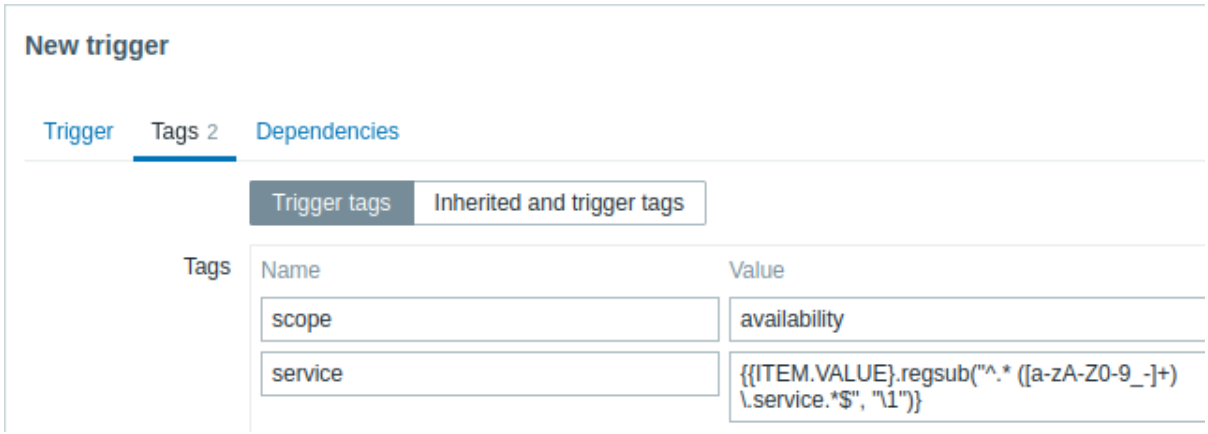
Since we do not want that it's important to make sure somehow that the corresponding root problems are closed, not just all problems. That's where tagging can help.

Problems and resolutions can be matched by specifying a tag in the trigger configuration. The following settings have to be made:

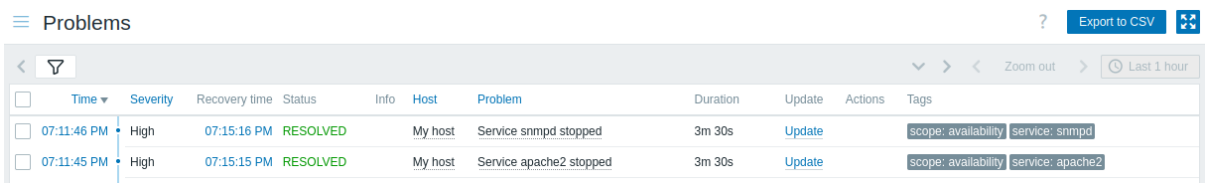
- *Problem event generation mode*: Multiple
- *OK event closes*: All problems if tag values match
- Enter the name of the tag for event matching



- configure the tags to extract tag values from log lines



If configured successfully you will be able to see problem events tagged by application and matched to their resolution in *Monitoring > Problems*.



Warning:
Because misconfiguration is possible, when similar event tags may be created for **unrelated** problems, please review the cases outlined below!

- Indexed macros always refer to the *Expression* field of the trigger configuration, not the *Recovery expression*. For example, in a recovery event, **{ITEM.VALUE1}** will resolve to the latest value of the first item **in the problem expression** at the time of recovery. If the recovery expression is based on a different item and the value of the problem expression item changes by the time of recovery, the events will get different tags and will not be correlated.

- With two applications writing error and recovery messages to the same log file a user may decide to use two *service* tags in the same trigger with different tag values by using separate regular expressions in the tag values to extract the names of, say, service A and service B from the {ITEM.VALUE} macro (e.g., when the message formats differ). However, this may not work as planned if there is no match to the regular expressions. Non-matching regexps will yield empty tag values and a single empty tag value in both problem and OK events is enough to correlate them. So a recovery message from service A may accidentally close an error message from service B.
- Actual tags and tag values only become visible when a trigger fires. If the regular expression used is invalid, it is silently replaced with an *UNKNOWN* string. If the initial problem event with an *UNKNOWN* tag value is missed, there may appear subsequent OK events with the same *UNKNOWN* tag value that may close problem events which they shouldn't have closed.
- If a user uses the {ITEM.VALUE} macro without macro functions as the tag value, the 255-character limitation applies. When log messages are long and the first 255 characters are non-specific, this may also result in similar event tags for unrelated problems.

Global event correlation

Overview

:::noteimportant This feature is deprecated and may be removed in future releases. See Complex event processing for updated event correlation features. :::note

Global event correlation allows reaching out over all metrics monitored by Zabbix and create correlations.

It is possible to correlate events created by completely different triggers and apply the same operations to them all. By creating intelligent correlation rules it is actually possible to save yourself from thousands of repetitive notifications and focus on root causes of a problem!

Global event correlation is a powerful mechanism, which allows you to untie yourself from one-trigger based problem and resolution logic. So far, a single problem event was created by one trigger and we were dependent on that same trigger for the problem resolution. We could not resolve a problem created by one trigger with another trigger. But with event correlation based on event tagging, we can.

For example, a log trigger may report application problems, while a polling trigger may report the application to be up and running. Taking advantage of event tags you can tag the log trigger as *status:down* while tag the polling trigger as *status:up*. Then, in a global correlation rule you can relate these triggers and assign an appropriate operation to this correlation such as closing the old events.

In another use, global correlation can identify similar triggers and apply the same operation to them. What if we could get only one problem report per network port problem? No need to report them all. That is also possible with global event correlation.

Global event correlation is configured in **correlation rules**. A correlation rule defines how the new problem events are paired with existing problem events and what to do in case of a match (close the new event, close matched old events by generating corresponding OK events). If a problem is closed by global correlation, it is reported in the *Info* column of *Monitoring* > *Problems*.

Configuring global correlation rules is available to Super Admin level users only.

> **Attention:**
> Event correlation must be configured very carefully, as it can negatively affect event processing performance or, if misconfigured, close more events than was intended (in the worst case even all problem events could be closed).

To configure global correlation **safely**, observe the following important tips:

- Reduce the correlation scope. Always set a unique tag for the new event that is paired with old events and use the *New event tag name* correlation condition.
- Add an explicit old-event condition when using **Close old events**. Always add at least one *Old event* condition (for example, *Old event tag name*, *Old event tag value*, or *Event tag pair*) when selecting **Close old events**; otherwise, the rule may match and close unrelated existing problem events (in the worst case, all problems). Prefer **Event tag pair** for matching runtime values (host:port, session id, etc.), and further narrow the match by host or host group when possible.
- Avoid using common tag names that may end up being used by different correlation configurations.
- Keep the number of correlation rules limited to the ones you really need.

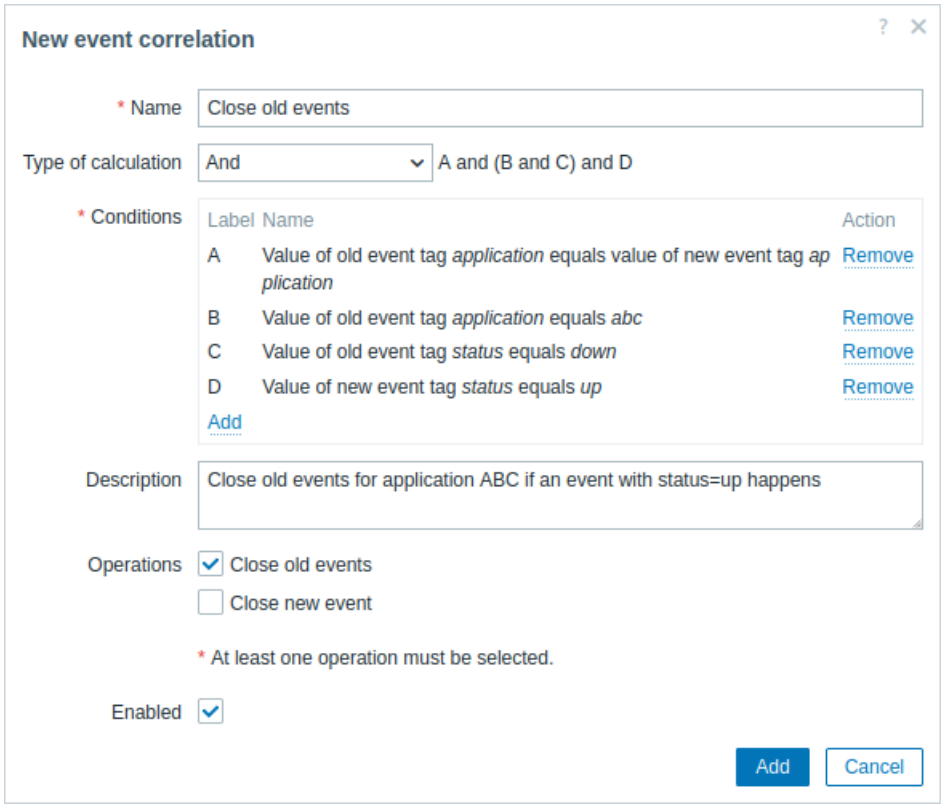
See also: known issues.

Configuration

To configure event correlation rules globally:

1. Go to *Data collection* > *Event processing*.

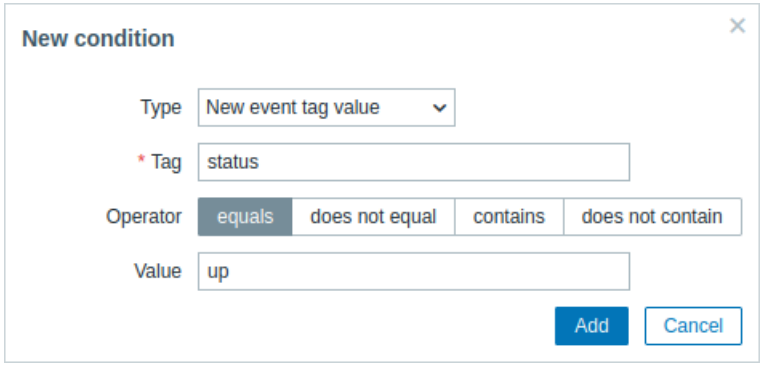
2. Click *Create event correlation* to the right (or the correlation name to edit an existing rule).
3. Enter parameters of the correlation rule in the form.



All mandatory input fields are marked with a red asterisk.

Parameter	Description
Name	Unique correlation rule name.
Type of calculation	The following options of calculating conditions are available: **And** - all conditions must be met **Or** - enough if one condition is met **And/Or** - AND with different condition types and OR with the same condition type **Custom expression** - a user-defined calculation formula for evaluating action conditions. It must include all conditions (represented as uppercase letters A, B, C, ...) and may include spaces, tabs, brackets ( ), **and** (case sensitive), **or** (case sensitive), **not** (case sensitive).
Conditions	List of conditions. See below for details on configuring a condition.
Description	Correlation rule description.
Operations	Mark the checkbox of the operation to perform when event is correlated. The following operations are available: **Close old events** - close old events when a new event happens. Always add a condition based on the old event when using the *Close old events* operation or all existing problems could be closed. **Close new event** - close the new event when it happens. **Warning!** Do not leave old/new event conditions empty when using *Close old events*/*Close new event*.If you select the *Close old events* operation without adding a condition that matches the old event, Zabbix may match all existing old events and close them.Always add an explicit old event condition (for example, *Old event tag name* or *Event tag pair*) when using *Close old events*. For example, a rule that uses only a *New event condition* and the *Close old events operation* will match all old events that meet the (missing) old-event criteria - effectively closing old problems.
Enabled	If you mark this checkbox, the correlation rule will be enabled.

To configure details of a new condition, click Add in the Conditions block. A popup window will open where you can edit the condition details.



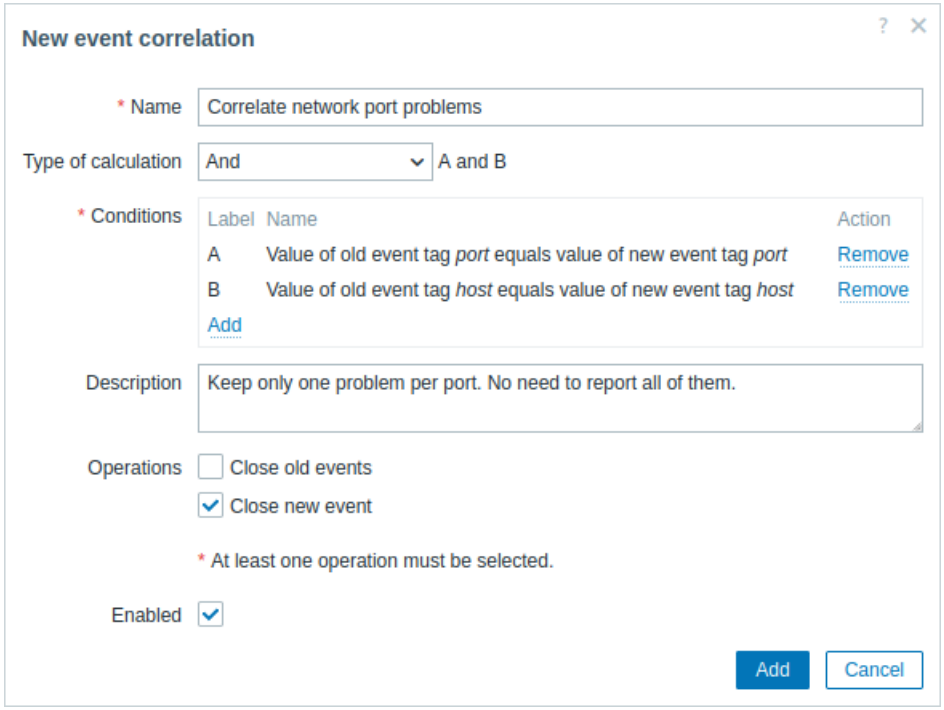
Parameter	Description
New condition	Select a condition for correlating events. *Note* that if no old event condition is specified, all old events may be matched and closed. Similarly if no new event condition is specified, all new events may be matched and closed. The following conditions are available: **Old event tag name** - specify the old event tag name for matching. **New event tag name** - specify the new event tag name for matching. **New event host group** - specify the new event host group for matching. **Event tag pair** - specify new event tag name and old event tag name for matching. In this case there will be a match if the **values** of the tags in both events match. Tag *names* need not match. This option is useful for matching runtime values, which may not be known at the time of configuration (see also Example). **Old event tag value** - specify the old event tag name and value for matching, using the following operators: *equals* - has the old event tag value *does not equal* - does not have the old event tag value *contains* - has the string in the old event tag value *does not contain* - does not have the string in the old event tag value **New event tag value** - specify the new event tag name and value for matching, using the following operators: *equals* - has the new event tag value *does not equal* - does not have the new event tag value *contains* - has the string in the new event tag value *does not contain* - does not have the string in the new event tag value

Warning:
Because misconfiguration is possible, when similar event tags may be created for **unrelated** problems, please review the cases outlined below!

- Actual tags and tag values only become visible when a trigger fires. If the regular expression used is invalid, it is silently replaced with an *UNKNOWN* string. If the initial problem event with an *UNKNOWN* tag value is missed, there may appear subsequent OK events with the same *UNKNOWN* tag value that may close problem events which they shouldn't have closed.
- If a user uses the {ITEM.VALUE} macro without macro functions as the tag value, the 255-character limitation applies. When log messages are long and the first 255 characters are non-specific, this may also result in similar event tags for unrelated problems.

Example

Stop repetitive problem events from the same network port.



This global correlation rule will correlate problems if *host* and *port* tag values exist on the trigger and they are the same in the original event and the new one.

The operation will close new problem events on the same network port, keeping only the original problem open.

Discovery

Overview Discovery provides automatic ways to add hosts and configure monitoring without manual setup. Zabbix supports multiple discovery methods to adapt to different scenarios.

Discovery methods

- Network discovery - automatically detect hosts and services available on your network.
- Active agent autoregistration - allow active agents to automatically register themselves with the server.
- Low-level discovery - automatically create items, triggers, graphs, and hosts for discovered entities like file systems or network interfaces using prototypes and nested discovery rules.

Network discovery

Overview

Zabbix offers automatic network discovery functionality that is effective and very flexible.

With network discovery properly set up you can:

- Speed up Zabbix deployment.
- Simplify administration.
- Use Zabbix in rapidly changing environments without excessive administration.

Zabbix network discovery is based on the following information:

- IP ranges
- Availability of external services (FTP, SSH, WEB, POP3, IMAP, TCP, etc)
- Information received from Zabbix agent (only unencrypted mode is supported)
- Information received from SNMP agent

It does NOT provide discovery of network topology.

Network discovery basically consists of two phases: discovery and actions.

Discovery

Zabbix periodically scans the IP ranges defined in network discovery rules. The frequency of the check is configurable for each rule individually.

Each rule has a set of service checks defined to be performed for the IP range.

Discovery rules are processed by the discovery manager. The discovery manager creates a job per each rule with a list of tasks (network checks). Network checks are performed in parallel by the available discovery workers (the number is configurable in the frontend per each rule). The exception is SNMPv3 checks, which are processed by a single worker. Only checks with the same IP and port are scheduled sequentially because some devices will not accept parallel connections on the same port.

The queue size of network checks is limited to 2000000 or 4 GB of memory approximately. If the queue becomes full then the discovery rule will be skipped and a warning message will be printed in the log. You may use the `zabbix[discovery_queue]` internal item to monitor the number of discovery checks in the queue.

Discovery checks are processed independently from the other checks. If any checks do not find a service (or fail), other checks will still be processed.

Note:
If a discovery rule is changed during execution, then the current discovery execution will be aborted.

Every check of a service and a host (IP) performed by the network discovery module generates a discovery event.

Event	Check of service result
Service Discovered	The service is 'up' after it was 'down' or when discovered for the first time.
Service Up	The service is 'up', after it was already 'up'.
Service Lost	The service is 'down' after it was 'up'.
Service Down	The service is 'down', after it was already 'down'.
Host Discovered	At least one service of a host is 'up' after all services of that host were 'down' or a service is discovered which belongs to a not registered host.
Host Up	At least one service of a host is 'up', after at least one service was already 'up'.
Host Lost	All services of a host are 'down' after at least one was 'up'.
Host Down	All services of a host are 'down', after they were already 'down'.

Actions

Discovery events can be the basis of relevant actions, such as:

- Sending notifications
- Adding/removing hosts
- Enabling/disabling hosts
- Adding hosts to a group
- Removing hosts from a group
- Adding tags to a host
- Removing tags from a host
- Linking a template to hosts/unlinking a template from hosts
- Executing remote scripts

These actions can be configured with respect to the device type, IP, status, uptime/downtime, etc. For full details on configuring actions for network-discovery based events, see action operation and conditions pages.

Since network discovery actions are event-based, they will be triggered both when a discovered host is online and when it is offline. It is highly recommended to add an action condition *Discovery status: up* to avoid such actions as *Add host* being triggered upon *Service Lost*/*Service Down* events. Otherwise, if a discovered host is manually removed, it will still generate *Service Lost*/*Service Down* events and will be recreated during the next discovery cycle.

Note:
Linking templates to a discovered host will fail collectively if any of the linkable templates has a unique entity (e.g., item key) that is the same as a unique entity (e.g., item key) already existing on the host or on another of the linkable templates.

Host creation

A host is added if the *Add host* operation is selected. A host is also added, even if the *Add host* operation is missing, if you select operations resulting in actions on a host. Such operations are:

- Enable host.

- Disable host.
- Add host to a host group.
- Link template to a host.

Created hosts are added to the *Discovered hosts* group (by default, configurable in *Administration* > *General* > *Other*). If you wish hosts to be added to another group, add a *Remove from host groups* operation (specifying "Discovered hosts") and also add an *Add to host groups* operation (specifying another host group), because a host must belong to a host group.

The IP address of the discovered device, together with the discovery source (Zabbix server, Zabbix proxy, or proxy group) and the interface type, is used as the criterion for finding a host in the system. If a host with the same IP address, interface type and discovery source already exists, that host will be the target for performing operations. If the discovery source differs, the discovered entity is treated as a different host and a new host may be created.

If the IP address of the discovered host is changed or the interface is deleted, a new host will be created upon the next discovery.

Host naming

When adding hosts, a host name is the result of reverse DNS lookup or IP address if reverse lookup fails. Lookup is performed from the Zabbix server or Zabbix proxy, depending on which is doing the discovery. If lookup fails on the proxy, it is not retried on the server. If the host with such a name already exists, the next host would get **_2** appended to the name, then **_3** and so on.

It is also possible to override DNS/IP lookup and instead use an item value for host name, for example:

- You may discover multiple servers with Zabbix agent running using a Zabbix agent item for discovery and assign proper names to them automatically, based on the string value returned by this item.
- You may discover multiple SNMP network devices using an SNMP agent item for discovery and assign proper names to them automatically, based on the string value returned by this item.

If the host name has been set using an item value, it is not updated during the following discovery checks. If it is not possible to set host name using an item value, default value (DNS name) is used.

If a host already exists with the discovered IP address and the discovery source (Zabbix server, proxy or proxy group) has not changed, a new host is not created. If the discovery source differs, the discovered entity is treated as distinct and a new host may be created. However, if the discovery action contains operations (link template, add to host group, etc), they are performed on the existing host that matches by IP address, interface type and discovery source.

Host removal

Hosts discovered by a network discovery rule are removed automatically from *Monitoring* > *Discovery* if a discovered entity is not in the rule's IP range any more. Hosts are removed immediately.

Interface creation when adding hosts

When hosts are added as a result of network discovery, they get interfaces created according to these rules:

- The services detected - for example, if an SNMP check succeeded, an SNMP interface will be created.
- If a host responded both to Zabbix agent and SNMP requests, both types of interfaces will be created.
- If uniqueness criteria are Zabbix agent or SNMP-returned data, the first interface found for a host will be created as the default one. Other IP addresses will be ignored, unless the default interface goes down. In case where the default interface is down, the smallest of the other IP addresses will be used as a replacement for the default interface. Action's conditions (such as Host IP) do not impact adding interfaces. *Note* that this will work if all interfaces are discovered by the same discovery rule. If a different discovery rule discovers a different interface of the same host, an additional host will be added.
- If a host responded to agent checks only, it will be created with an agent interface only. If it would start responding to SNMP later, additional SNMP interfaces would be added.
- If 3 separate hosts were initially created, having been discovered by the "IP" uniqueness criteria, and then the discovery rule is modified so that hosts A, B and C have identical uniqueness criteria result, B and C are created as additional interfaces for A, the first host. The individual hosts B and C remain. In *Monitoring* > *Discovery* the added interfaces will be displayed in the "Discovered device" column, in black font and indented, but the "Monitored host" column will only display A, the first created host. "Uptime/Downtime" is not measured for IPs that are considered to be additional interfaces.

Changing proxy setting

The hosts discovered by different proxies are not always treated as different hosts. Discovery and uniqueness checks depend on the proxy group structure: when a proxy runs a discovery rule and creates a host, that host is added to the proxy's parent proxy group, not assigned to the proxy itself. When Zabbix evaluates IP uniqueness during discovery, it checks hosts monitored by the parent proxy group. Hosts monitored by individual proxies within that group (including the proxy that ran the discovery) are ignored for the uniqueness check, which can result in duplicate hosts if multiple proxies monitor overlapping subnets.

While this behavior allows discovery to operate across overlapping IP ranges used by different subnets, changing the proxy assigned to an already monitored subnet is more complicated because proxy changes must be applied consistently to discovered hosts and to the membership of the parent proxy group to avoid duplicates.

For example, the steps to replace proxy in a discovery rule:

1. Disable discovery rule.
2. Sync proxy configuration.
3. Replace the proxy in the discovery rule.
4. Replace the proxy for all hosts discovered by this rule (ensure hosts in the parent proxy group and any hosts monitored by individual proxies in that group are updated to avoid duplicates).
5. Enable discovery rule.

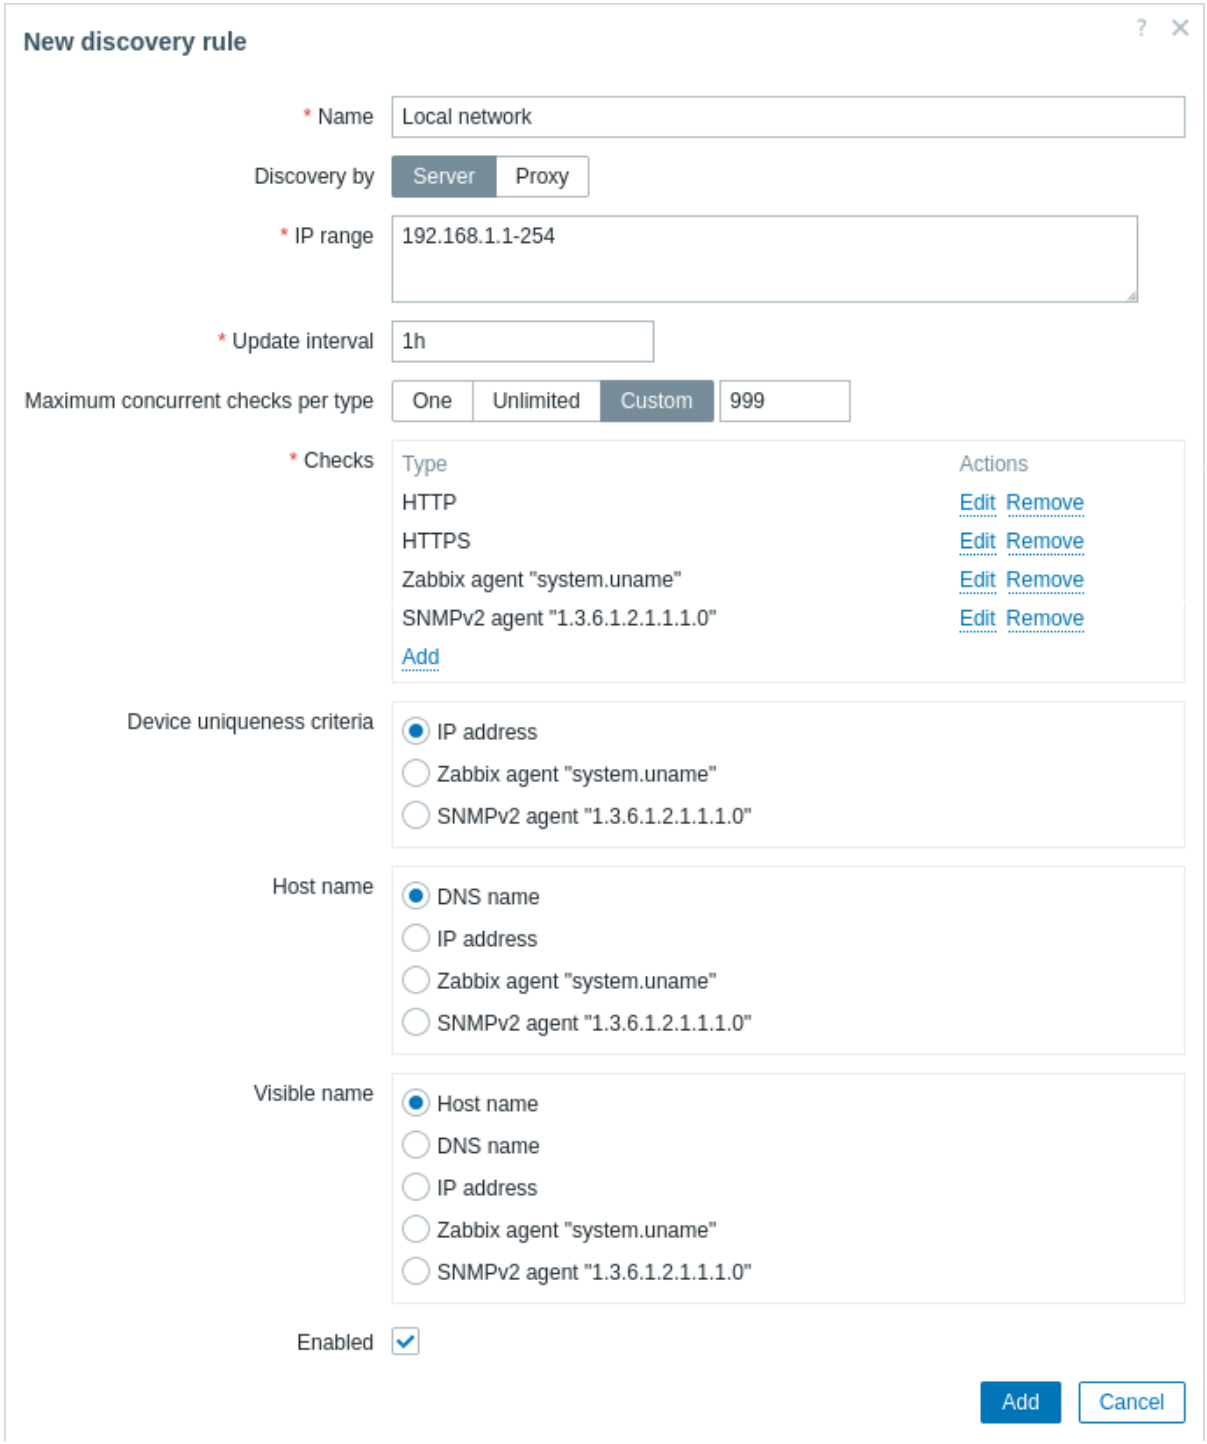
Configuring a network discovery rule

Overview

To configure a network discovery rule used by Zabbix to discover hosts and services:

1. Go to *Data collection* > *Discovery*.
2. Click *Create discovery rule* (or the rule name to edit an existing one).
3. Edit the discovery rule attributes.

Rule attributes



All mandatory input fields are marked with a red asterisk.

Parameter	Description
Name	Unique name of the rule. For example, "Local network".
Discovery by	Discovery is performed by: **Server** - by Zabbix server **Proxy** - by Zabbix proxy (selected in the proxy name field)
IP range	The range of IP addresses for discovery. It may have the following formats: Single IP: 192.168.1.33 Range of IP addresses: 192.168.1-10.1-255. The range is limited by the total number of covered addresses (less than 64K). IP mask: 192.168.4.0/24 supported IP masks: /16 - /30 for IPv4 addresses /112 - /128 for IPv6 addresses List: 192.168.1.1-255, 192.168.2.1-100, 192.168.2.200, 192.168.4.0/24 This field supports spaces, tabulation and multiple lines.

Parameter	Description
Update interval	This parameter defines how often Zabbix will execute the rule. The interval is measured after the execution of previous discovery instance ends so there is no overlap. Time suffixes are supported, e.g., 30s, 1m, 2h, 1d. User macros are supported. *Note* that if a user macro is used and its value is changed (e.g., 1w → 1h), the next check will be executed according to the previous value (far in the future with the example values).
Maximum concurrent checks per type	Set the maximum number of discovery threads (workers) per service check for parallel processing of discovery checks: **One** - one thread **Unlimited** - unlimited number of threads (but no more than in the StartDiscoverers parameter) **Custom** - set a custom number of threads (0-999). SNMPv1 and SNMPv2 asynchronous service checks are processed concurrently by available workers. All discovery rules with SNMPv3 asynchronous service checks are processed by one worker due to the peculiarities of the libsnmp implementation; therefore, increasing the number of workers will not increase the discovery speed.
Checks	Zabbix will use this list of checks for discovery. Click Add to configure a new check in a popup window. Supported checks: SSH, LDAP, SMTP, FTP, HTTP, HTTPS, POP, NNTP, IMAP, TCP, Telnet, Zabbix agent, SNMPv1 agent, SNMPv2 agent, SNMPv3 agent, ICMP ping. A protocol-based discovery uses the **net.tcp.service[]** functionality to test each host, except for SNMP which queries an SNMP OID. Zabbix agent is tested by querying an item in unencrypted mode. Please see agent items for more details. The 'Ports' parameter may be one of following: Single port: 22 Range of ports: 22-45 List: 22-45,55,60-70 Since Zabbix 7.0, all service checks are performed asynchronously, except for LDAP checks. Since Zabbix 7.0 HTTP/HTTPs checking is done via libcurl. If Zabbix server/proxy is compiled without libcurl, then HTTP checks will work like in previous versions (i.e. as TCP checks), but HTTPS checks will not work.
Device uniqueness criteria	Uniqueness criteria may be: **IP address** - do not process multiple single-IP devices. If a device with the same IP already exists it will be considered already discovered and a new host will not be added. **<discovery check>** - either Zabbix agent or SNMP agent check. Note that uniqueness criteria used during discovery is not the same as host identification in the system when performing actions. Uniqueness criteria during discovery define whether two or more discovered devices are the same (or different), whereas only the IP address is the criterion for host identification in Zabbix (see Host creation).
Host name	Set the technical host name of a created host using: **DNS name** - DNS name (default) **IP address** - IP address **<discovery check>** - received string value of the discovery check (e.g., Zabbix agent, SNMP agent check) See also: Host naming.
Visible name	Set the visible host name of a created host using: **Host name** - technical host name (default) **DNS name** - DNS name **IP address** - IP address **<discovery check>** - received string value of the discovery check (e.g., Zabbix agent, SNMP agent check) See also: Host naming.
Enabled	With the check-box marked the rule is active and will be executed by Zabbix server. If unmarked, the rule is not active. It won't be executed.

Exceeding file descriptor limit

In case of large number of concurrent checks, it is possible to exhaust the file descriptor limit for the discovery manager.

The number of file descriptors required for detection equates to the number of discovery workers * 1000. By default, there are 5 discovery workers, while the soft limit of the system is approximately 1024.

If this limit is approached, Zabbix will reduce the default number of concurrent checks per type for each worker and write a warning to the log file. But, if the user has set a higher value for *Maximum concurrent checks per type* than the value calculated by Zabbix, Zabbix will use the user-defined value for one worker.

A real life scenario

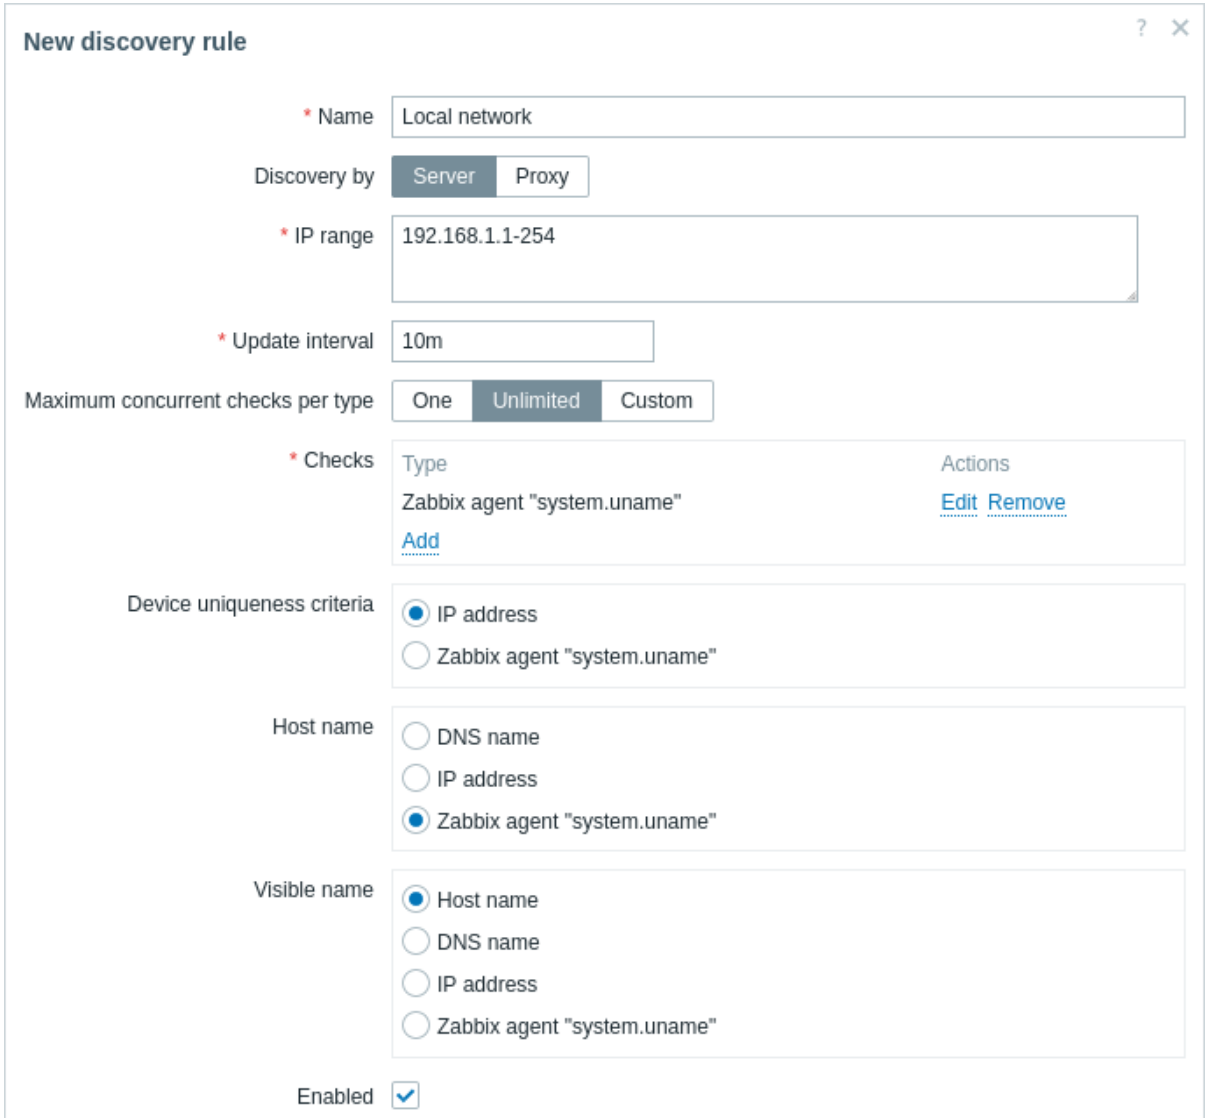
In this example, we would like to set up network discovery for the local network having an IP range of 192.168.1.1-192.168.1.254.

In our scenario we want to:

- discover those hosts that have Zabbix agent running
- run discovery every 10 minutes
- add a host to monitoring if the host uptime is more than 1 hour
- remove hosts if the host downtime is more than 24 hours
- add Linux hosts to the ”Linux servers” group
- add Windows hosts to the ”Windows servers” group
- use the template *Linux* for Linux hosts
- use the template *Windows* for Windows hosts

Step 1

Defining a network discovery rule for our IP range.



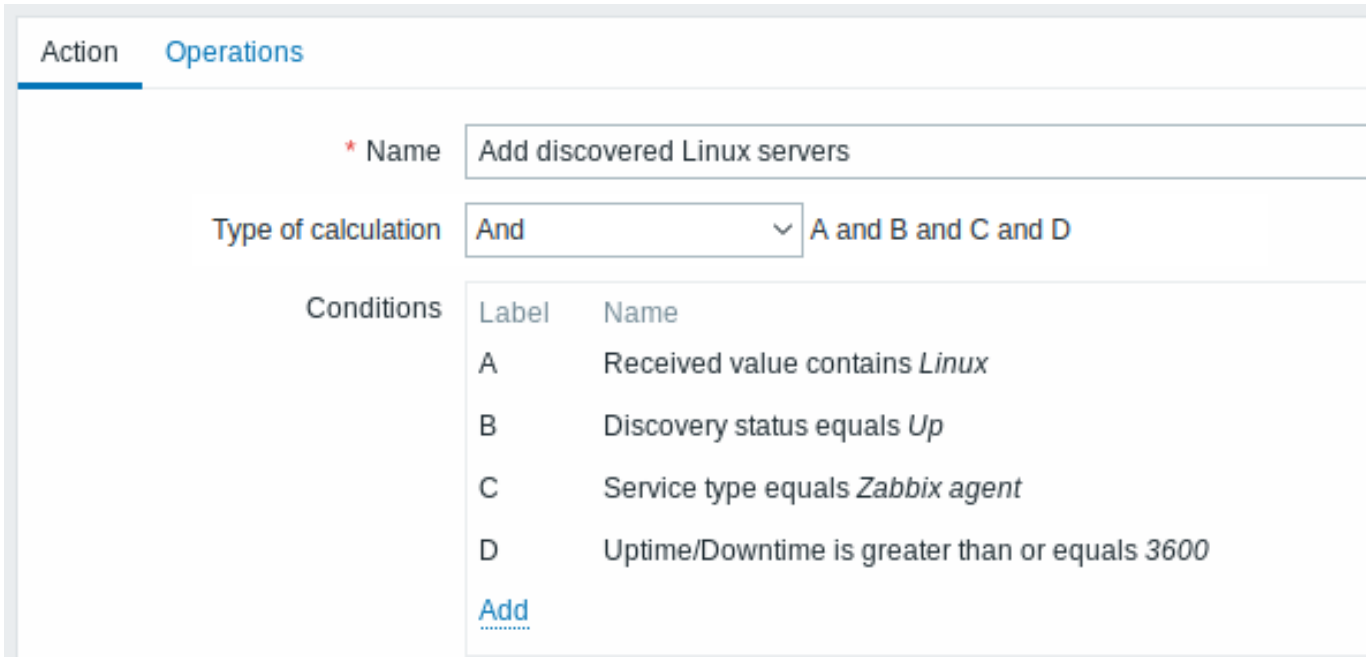
Zabbix will try to discover hosts in the IP range of 192.168.1.1-192.168.1.254 by connecting to Zabbix agents and getting the value from the **system.uname** key. The value received from the agent can be used to name the hosts and also to apply different actions for different operating systems. For example, link Windows servers to the template *Windows*, Linux servers to the template *Linux*.

The rule will be executed every 10 minutes.

When this rule is added, Zabbix will automatically start the discovery and generation of the discovery-based events for further processing.

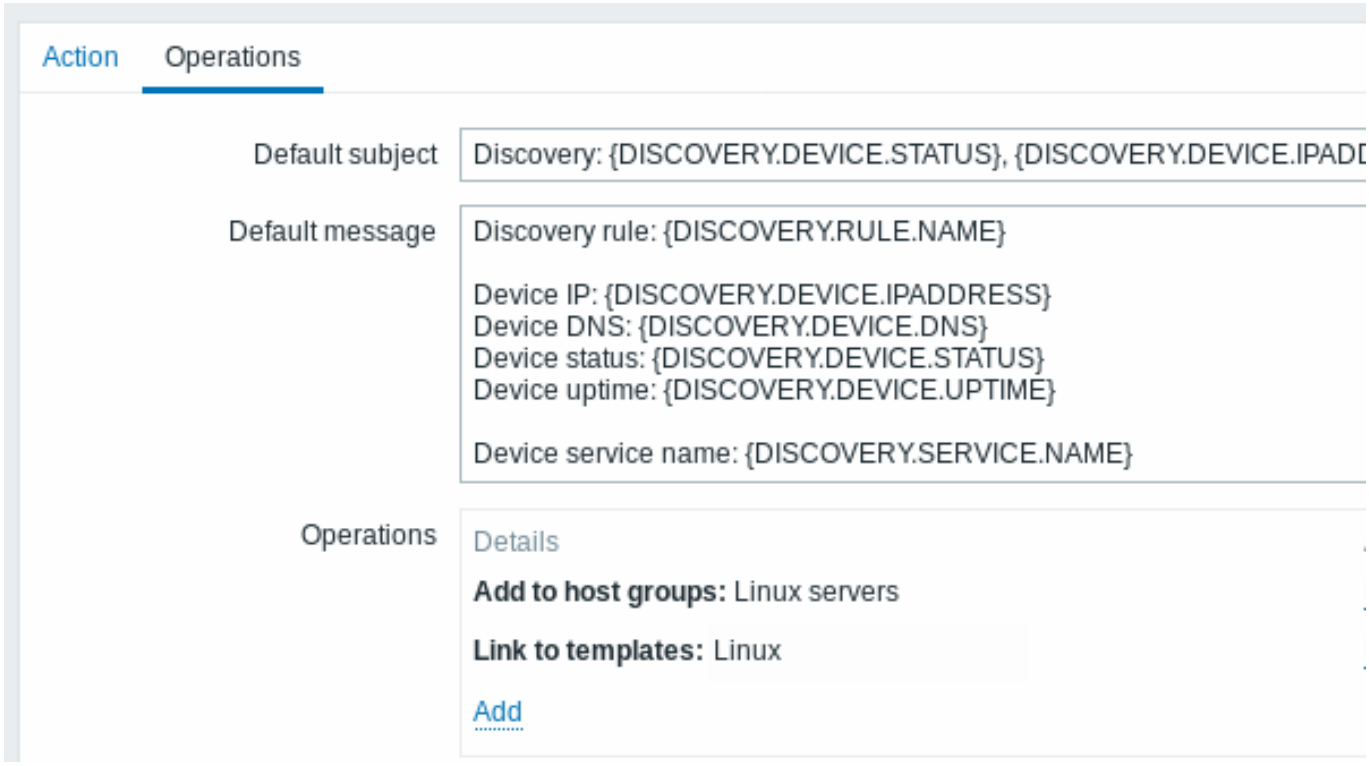
Step 2

Defining a discovery action for adding the discovered Linux servers to the respective group/template.



The action will be activated if:

- the ”Zabbix agent” service is ”up”
- the value of system.uname (the Zabbix agent key we used in rule definition) contains ”Linux”
- Uptime is 1 hour (3600 seconds) or more

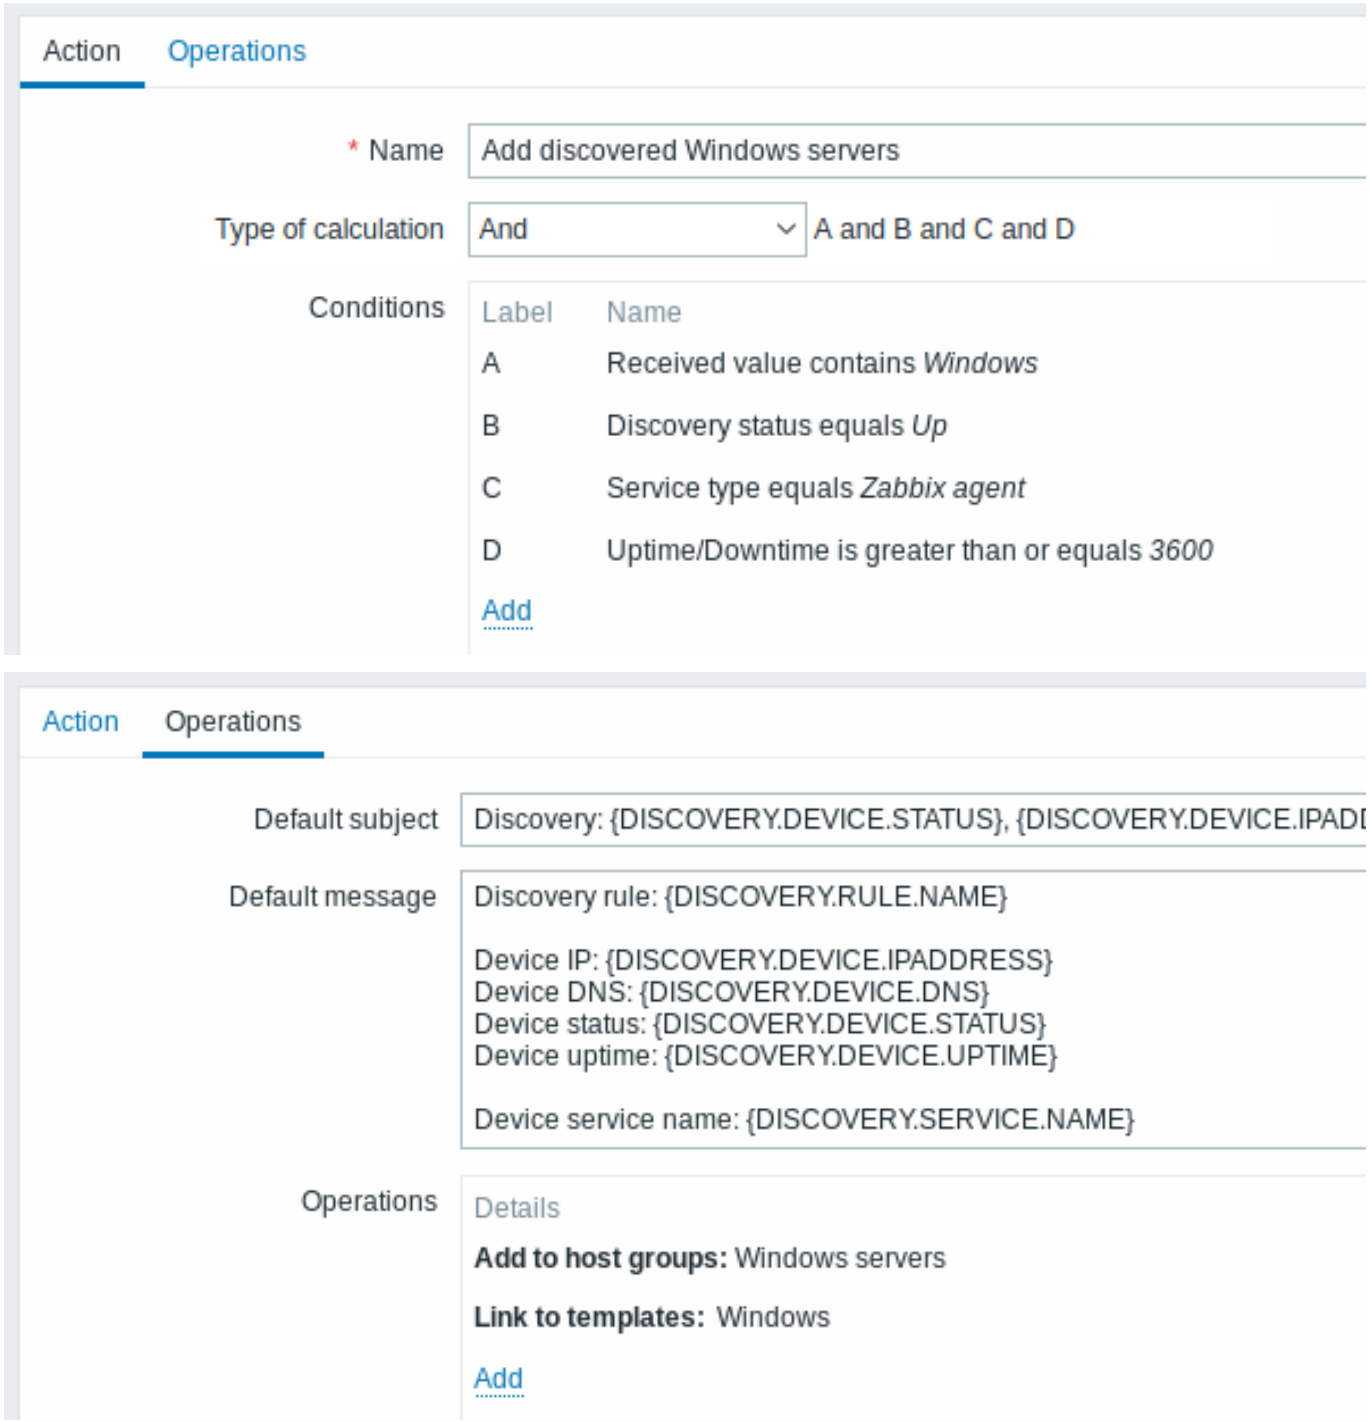


The action will execute the following operations:

- add the discovered host to the ”Linux servers” group (and also add host if it wasn’t added previously)
- link host to the *Linux* template. Zabbix will automatically start monitoring the host using items and triggers from the ”Linux” template.

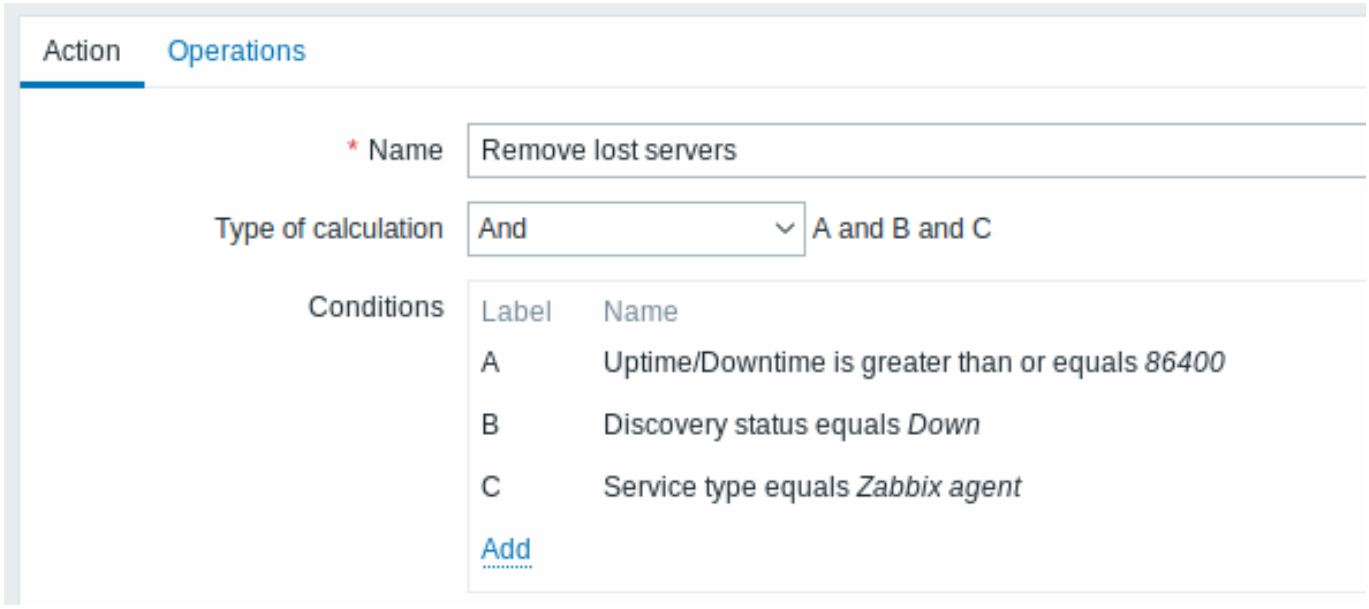
Step 3

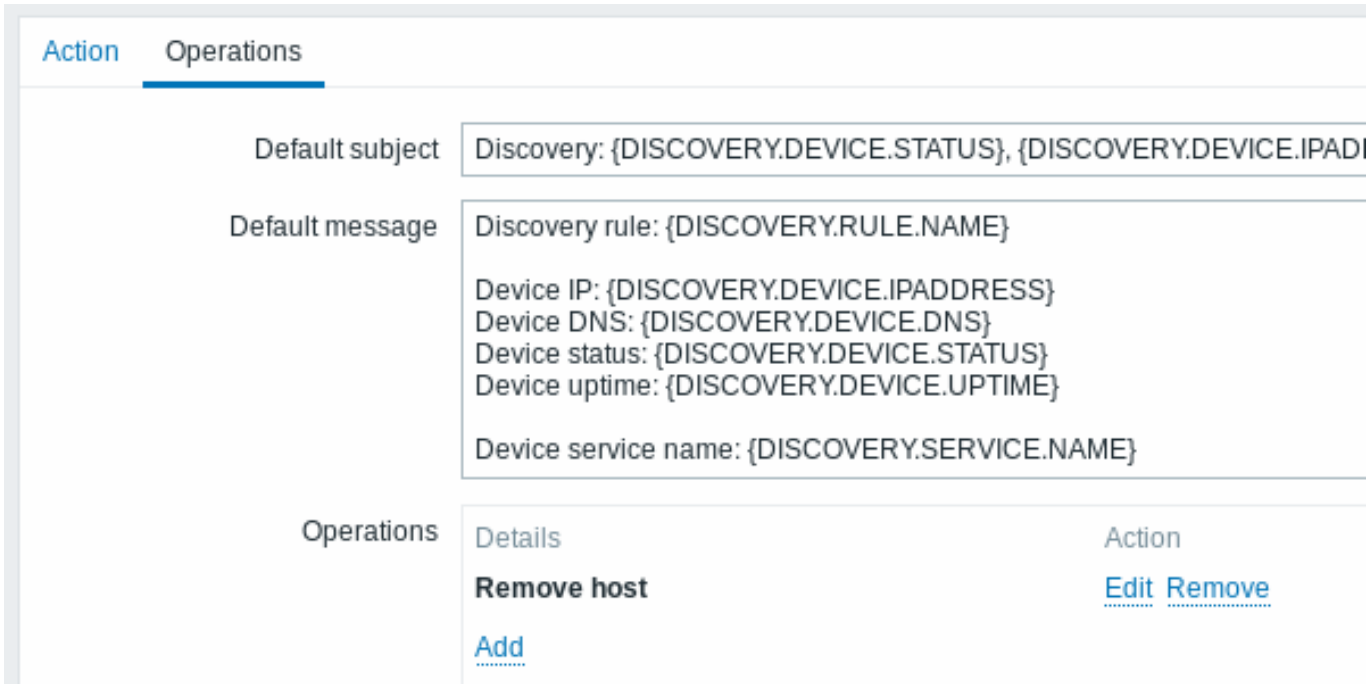
Defining a discovery action for adding the discovered Windows servers to the respective group/template.



Step 4

Defining a discovery action for removing lost servers.





A server will be removed if ”Zabbix agent” service is ’down’ for more than 24 hours (86400 seconds).

Active agent autoregistration

Overview

It is possible to allow active Zabbix agent autoregistration, after which the server can start monitoring them. This way new hosts can be added for monitoring without configuring them manually on the server.

Autoregistration can happen when a previously unknown active agent asks for checks.

The feature might be very handy for automatic monitoring of new cloud nodes. As soon as you have a new node in the Zabbix Cloud will automatically start the collection of performance and availability data of the host.

Active agent autoregistration also supports the monitoring of added hosts with passive checks. When the active agent asks for checks, providing it has the `ListenIP` or `ListenPort` configuration parameters defined in the configuration file, these are sent along to the server. If multiple IP addresses are specified, the first one is sent to the server.

Server, when adding the new autoregistered host, uses the received IP address and port to configure the agent. If no IP address value is received, the one used for the incoming connection is used. If no port value is received, 10050 is used.

It is possible to specify that the host should be autoregistered with a DNS name as the default agent interface.

Autoregistration is rerun:

- if host metadata information changes:
  - due to HostMetadata changed and agent restarted
  - due to value returned by HostMetadataItem changed
- for manually created hosts with metadata missing
- if a host is manually changed to be monitored by another Zabbix proxy
- if autoregistration for the same host comes from a new Zabbix proxy

The active agent autoregistration heartbeat for Zabbix server and Zabbix proxy is 120 seconds. So in case a discovered host is deleted, the autoregistration will be rerun in 120 seconds.

Configuration

Specify server

Make sure you have the Zabbix server identified in the agent configuration file (`zabbix_agentd.conf`):

```
ServerActive=10.0.0.1
```

Unless you specifically define a *Hostname* in `zabbix_agentd.conf`, the system hostname of agent location will be used by server for naming the host. The system hostname in Linux can be obtained by running the `hostname` command.

If *Hostname* is defined in Zabbix agent configuration as a comma-delimited list of hosts, hosts will be created for all listed hostnames.

Restart the agent after making any changes to the configuration file.

Action for active agent autoregistration

When server receives an autoregistration request from an agent it calls an action. An action of event source "Autoregistration" must be configured for agent autoregistration.

Note:
Setting up network discovery is not required to have active agents autoregister.

In the Zabbix frontend:

1. Go to *Alerts* > *Actions*.
2. Select *Autoregistration actions*.
3. Click *Create action*.
4. In the *Action* tab, give your action a name.
5. Optionally specify conditions. You can do a substring match or regular expression match in the conditions for host name/host metadata. If you are going to use the *Host metadata* condition type, see Using host metadata.
6. In the *Operations* tab, add relevant operations, such as *Add host*, *Add to host group* (for example, *Discovered hosts*), *Link templates*, etc.

Note:
If the hosts that will be autoregistering are likely to be supported for active monitoring only (such as hosts that are firewalled from your Zabbix server) then you might want to create a specific template like *Template_Linux-active* to link to.

Created hosts are added to the *Discovered hosts* group (by default, configurable in *Administration* > *General* > *Other*). If you wish hosts to be added to another group, add a *Remove from host group* operation (specifying "Discovered hosts") and also add an *Add to host group* operation (specifying another host group), because a host must belong to a host group.

Secure autoregistration

A secure way of autoregistration is possible by configuring PSK-based authentication with encrypted connections.

The level of encryption is configured globally in *Administration* > *General* > *Autoregistration*. It is possible to select no encryption, TLS encryption with PSK authentication or both (so that some hosts may register without encryption while others through encryption).

Authentication by PSK is verified by Zabbix server before adding a host. If successful, the host is added and *Connections from/to host* are set to 'PSK' only with identity/pre-shared key the same as in the global autoregistration setting.

Attention:
To ensure security of autoregistration on installations using proxies, encryption between Zabbix server and proxy should be enabled.

Using DNS as default interface

HostInterface and HostInterfaceItem configuration parameters allow you to specify a custom value for the host interface during autoregistration.

More specifically, they are useful if the host should be autoregistered with a DNS name as the default agent interface rather than its IP address. In that case the DNS name should be specified or returned as the value of either HostInterface or HostInterfaceItem parameters. If the value of one of these parameters changes—for example, from an IP address to a DNS name or vice versa—the default interface of the autoregistered host will be updated accordingly. This update is applied to the existing host, not by creating a new one. To send the new value, the agent must be restarted so that it re-initiates the autoregistration process.

> **Note:**
> If HostInterface or HostInterfaceItem parameters are not configured, the listen_dns parameter is used instead. This value is determined by performing a reverse DNS lookup of the agent's IP address. If reverse DNS is not properly configured or returns an invalid name, it may result in incorrect or failed autoregistration due to an invalid interface value.

Using host metadata

When agent is sending an autoregistration request to the server it sends its hostname. In some cases (for example, Amazon cloud nodes) a hostname is not enough for Zabbix server to differentiate discovered hosts. Host metadata can be optionally used to send other information from an agent to the server.

Host metadata is configured in the agent configuration file (`zabbix_agentd.conf`). There are 2 ways of specifying host metadata in the configuration file:

```
HostMetadata
HostMetadataItem
```

See the description of the options in the link above.

The HostMetadataItem parameter may return up to 65535 UTF-8 code points. A longer value will be truncated.

Note that on MySQL, the effective maximum length in characters will be less if the returned value contains multibyte characters. For example, a value containing 3-byte characters only will be limited to 21844 characters in total, while a value containing 4-byte characters only will be limited to 16383 symbols.

> **Attention:**
> An autoregistration attempt happens every time an active agent sends a request to refresh active checks to the server. The delay between requests is specified in the RefreshActiveChecks parameter of the agent. The first request is sent immediately after the agent is restarted.

Examples

Autoregistration by OS using HostMetadata

Say you would like the hosts to be autoregistered by the Zabbix server. You have active Zabbix agents (see "Configuration" section above) on your network. There are Windows hosts and Linux hosts on your network and you have "Linux by Zabbix agent" and "Windows by Zabbix agent" templates available in your Zabbix frontend. So at host registration, you would like the appropriate Linux/Windows template to be applied to the host being registered. By default, only the hostname is sent to the server at autoregistration, which might not be enough. In order to make sure the proper template is applied to the host you should use host metadata.

Frontend configuration

The first thing to do is to configure the frontend. Create 2 actions. The first action:

- Name: Linux host autoregistration
- Conditions: Host metadata contains *Linux*
- Operations: Link templates: Linux by Zabbix agent

> **Note:**
> You can skip an "Add host" operation in this case. Linking a template to a host requires adding the host first so the server will do that automatically.

The second action:

- Name: Windows host autoregistration
- Conditions: Host metadata contains *Windows*
- Operations: Link templates: Windows by Zabbix agent

Agent configuration

Now you need to configure the agents. Add the next line to the agent configuration files:

```
HostMetadataItem=system.uname
```

This way you make sure host metadata will contain ”Linux” or ”Windows” depending on the host an agent is running on. An example of host metadata in this case:

```
Linux: Linux server3 3.2.0-4-686-pae #1 SMP Debian 3.2.41-2 i686 GNU/Linux
Windows: Windows WIN-0PXGGSTYNHO 6.0.6001 Windows Server 2008 Service Pack 1 Intel IA-32
```

Do not forget to restart the agent after making any changes to the configuration file.

Using HostMetadata to control autoregistration and templates

Step 1 - Protect registration with HostMetadata

Using host metadata to allow some basic protection against unwanted hosts registering.

Frontend configuration

Create an action in the frontend, using some hard-to-guess secret code to disallow unwanted hosts:

- Name: Autoregistration action Linux
- Conditions:
  - Type of calculation: AND
  - Condition (A): Host metadata contains //Linux//
  - Condition (B): Host metadata contains //21df83bf21bf0be663090bb8d4128558ab9b95fba66a6dbf834f8b91ae5e08ae//
- Operations:
  - Send message to users: Admin via all media
  - Add to host groups: Linux servers
  - Link templates: Linux by Zabbix agent

Please note that this method alone does not provide strong protection because data is transmitted in plain text. Configuration cache reload is required for changes to have an immediate effect.

Agent configuration

Add the next line to the agent configuration file:

```
HostMetadata=Linux    21df83bf21bf0be663090bb8d4128558ab9b95fba66a6dbf834f8b91ae5e08ae
```

where ”Linux” is a platform, and the rest of the string is the hard-to-guess secret text.

Do not forget to restart the agent after making any changes to the configuration file.

Step 2 - Add template to registered host

It is possible to add additional templates for an already registered host. In this case the MySQL by Zabbix agent template will be linked only to hosts whose HostMetadata contains the token `MySQL`.

Frontend configuration

Update the action in the frontend:

- Name: Autoregistration action Linux
- Conditions:
  - Type of calculation: AND
  - Condition (A): Host metadata contains Linux
  - Condition (B): Host metadata contains 21df83bf21bf0be663090bb8d4128558ab9b95fba66a6dbf834f8b91ae5e08ae
  - Condition (C): Host metadata contains MySQL
- Operations:
  - Send message to users: Admin via all media
  - Add to host groups: Linux servers
  - Link templates: Linux by Zabbix agent
  - Link templates: MySQL by Zabbix Agent

Agent configuration

Update the next line in the agent configuration file:

```
HostMetadata=MySQL on Linux 21df83bf21bf0be663090bb8d4128558ab9b95fba66a6dbf834f8b91ae5e08ae
```

Do not forget to restart the agent after making any changes to the configuration file.

Low-level discovery

Overview Low-level discovery (LLD) automatically discovers entities (e.g., file systems, network interfaces) on a host and creates the corresponding items, triggers, and graphs without manual setup for each.

To use LLD, you create a **discovery rule** that collects JSON data describing entities, and **prototypes** (item, trigger, and graph) for each entity.

The JSON returned by the rule must be an array of objects, where each object represents one discovered entity using key-value pairs. For example, a discovery rule with `net.if.discovery` might return:

```
[
  {"{#IFNAME}": "lo"},
  {"{#IFNAME}": "eth0"}
]
```

A corresponding set of items, triggers, and graphs is then created from the prototypes—one per interface (`lo` and `eth0`).

If objects contain keys in the `{#MACRO}` format, those macros are used directly in prototypes. If not, or if additional/custom macros are required, macros can be defined manually and mapped to JSON fields using JSONPath.

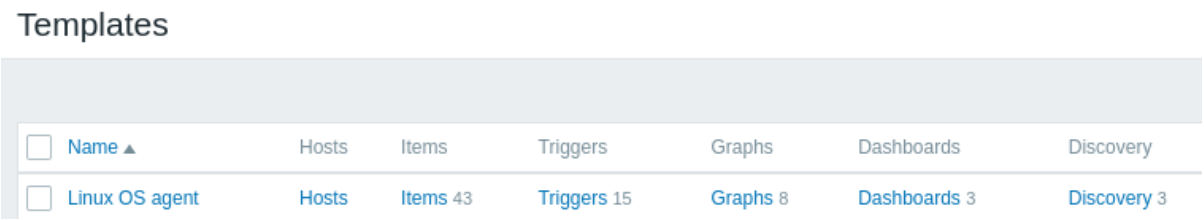
LLD can also create hosts (e.g., for virtual machines discovered on a hypervisor) and supports nested discovery rules for multi-level discovery.

Entities that are no longer discovered can be automatically disabled or deleted.

Configuring low-level discovery We will illustrate low-level discovery based on an example of file system discovery.

To configure the discovery, do the following:

1. Go to *Data collection* > *Templates* or *Hosts*.
2. Click *Discovery* in the row of an appropriate template/host.



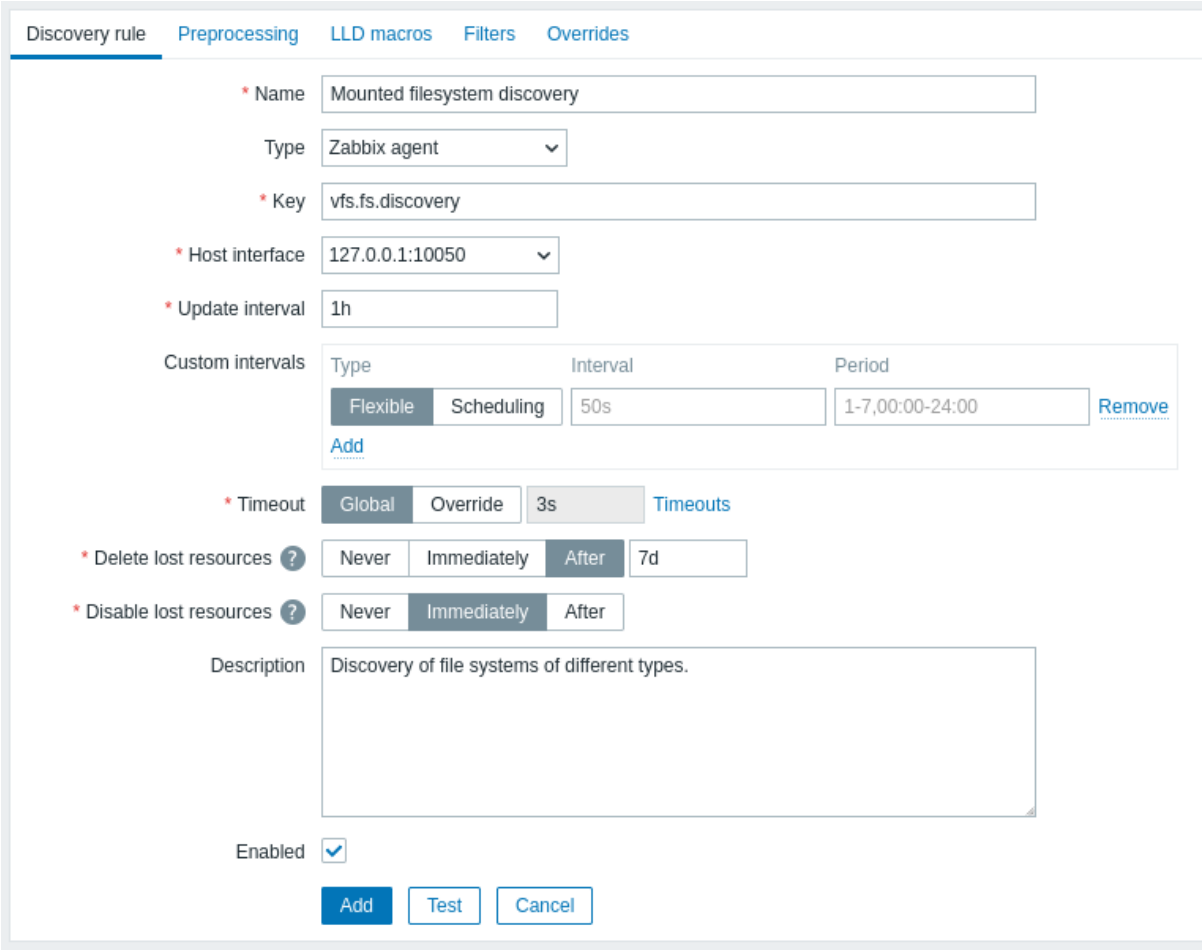
3. Click *Create discovery rule* in the upper-right corner of the screen.
4. Fill in the discovery rule form with the required details.

Discovery rule

The discovery rule form contains five tabs, representing, from left to right, the data flow during discovery:

- *Discovery rule* - specify, most importantly, the built-in item or custom script to retrieve discovery data.
- *Preprocessing* - apply some preprocessing to the discovered data.
- *LLD macros* - allow the extraction of macro values for use in discovered items, triggers, etc.
- *Filters* - allow filtering of discovered values.
- *Overrides* - allow the modification of items, triggers, graphs or host prototypes when applying to specific discovered objects.

The **Discovery rule** tab contains the item key to use for discovery (as well as some general discovery rule attributes):



All mandatory input fields are marked with a red asterisk.

Parameter	Description
Name	Name of discovery rule.
Type	The type of check to perform discovery. In this example we are using a *Zabbix agent* item type. The discovery rule can also be a dependent item, depending on a regular item. It cannot depend on another discovery rule. For a dependent item, select the respective type (*Dependent item*) and specify the master item in the 'Master item' field. The master item must exist.
Key	Enter the discovery item key (up to 2048 characters). For example, you may use the built-in "vfs.fs.discovery" item key to return a JSON string with the list of file systems present on the computer, their types and mount options. Note that another option for filesystem discovery is using discovery results by the "vfs.fs.get" agent key (see example).
Update interval	This field specifies how often Zabbix performs discovery. In the beginning, when you are just setting up file system discovery, you might wish to set it to a small interval, but once you know it works you can set it to 30 minutes or more, because file systems usually do not change very often. Time suffixes are supported, e.g., 30s, 1m, 2h, 1d. User macros are supported. *Note*: The update interval can only be set to '0' if custom intervals exist with a non-zero value. If set to '0', and a custom interval (flexible or scheduled) exists with a non-zero value, the item will be polled during the custom interval duration. New discovery rules will be checked within 60 seconds of their creation, unless they have Scheduling or Flexible update interval and the *Update interval* is set to 0. *Note* that for an existing discovery rule the discovery can be performed immediately by pushing the *Execute now* button.
Custom intervals	You can create custom rules for checking the item: **Flexible** - create an exception to the *Update interval* (interval with different frequency) **Scheduling** - create a custom polling schedule. For detailed information see Custom intervals.

Parameter	Description
Timeout	Set the discovery check timeout. Select the timeout option: **Global** - proxy/global timeout is used (displayed in the grayed out *Timeout* field); **Override** - custom timeout is used (set in the *Timeout* field; allowed range: 1 - 600s). Time suffixes, e.g., 30s, 1m, and user macros are supported. Clicking the *Timeouts* link allows you to configure proxy timeouts or global timeouts (if a proxy is not used). Note that the *Timeouts* link is visible only to users of *Super admin* type with permissions to *Administration* > *General* or *Administration* > *Proxies* frontend sections.
Delete lost resources	Specify how soon the discovered entity will be deleted once its discovery status becomes "Not discovered anymore": *Never* - it will not be deleted; *Immediately* - it will be deleted immediately; *After* - it will be deleted after the specified time period. The value must be greater than *Disable lost resources* value. Time suffixes are supported, e.g., 2h, 1d. User macros are supported. *Note*: Using "Immediately" is not recommended, since just wrongly editing the filter may end up in the entity being deleted with all the historical data. Note that manually disabled resources will not be deleted by low-level discovery.
Disable lost resources	Specify how soon the discovered entity will be disabled once its discovery status becomes "Not discovered anymore": *Never* - it will not be disabled; *Immediately* - it will be disabled immediately; *After* - it will be disabled after the specified time period. The value should be greater than the discovery rule update interval. Note that automatically disabled resources will become enabled again, if re-discovered by low-level discovery. Manually disabled resources will not become enabled again if re-discovered. This field is not displayed if *Delete lost resources* is set to "Immediately". Time suffixes are supported, e.g., 2h, 1d. User macros are supported.
Description	Enter a description.
Enabled	If checked, the rule will be processed.

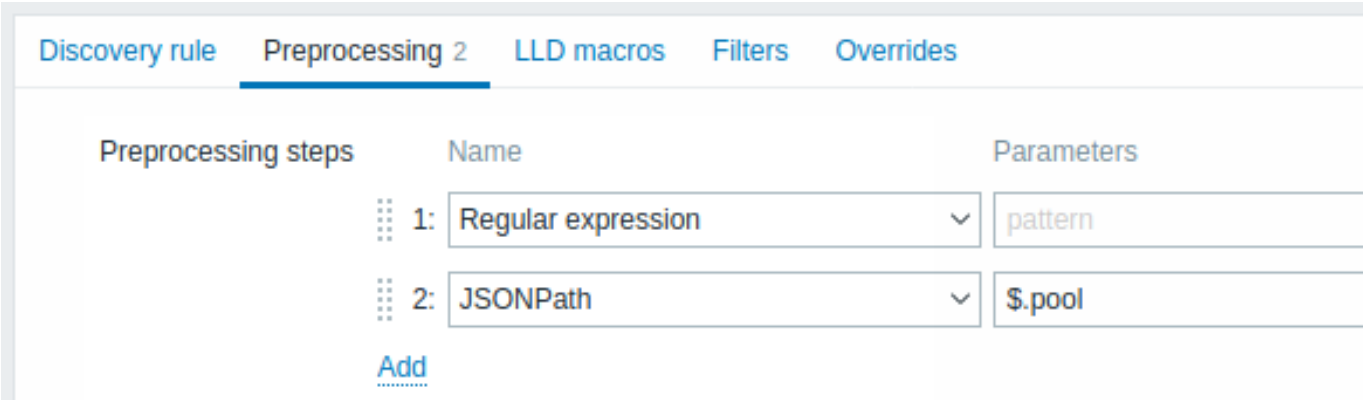
Note:
Discovery rule history is not preserved.

Preprocessing

The **Preprocessing** tab allows you to define transformation rules to apply to the result of discovery. One or several transformations are possible in this step. Transformations are executed in the order in which they are defined. All preprocessing is done by Zabbix server.

See also:

- Preprocessing details
- Preprocessing testing



Type		
	Transformation	Description
Text		
	Regular expression	Match the received value to the <pattern> regular expression and replace value with the extracted <output>. The regular expression supports extraction of maximum 10 captured groups with the \N sequence. Parameters: **pattern** - regular expression **output** - output formatting template. An \N (where N=1…9) escape sequence is replaced with the Nth matched group. A \0 escape sequence is replaced with the matched text. If you mark the *Custom on fail* checkbox, it is possible to specify custom error-handling options: either to discard the value, set a specified value or set a specified error message.
	Replace	Find the search string and replace it with another (or nothing). All occurrences of the search string will be replaced. Parameters: **search string** - the string to find and replace, case-sensitive (required) **replacement** - the string to replace the search string with. The replacement string may also be empty, effectively allowing the search string to be deleted when found. It is possible to use escape sequences to search for or replace line breaks, carriage return, tabs and spaces "\n \r \t \s"; backslash can be escaped as "\\" and escape sequences can be escaped as "\\n". Escaping of line breaks, carriage return, tabs is automatically done during low-level discovery.
Structured data		
	JSONPath	Extract value or fragment from JSON data using JSONPath functionality. If you mark the *Custom on fail* checkbox, it is possible to specify custom error-handling options: either to discard the value, set a specified value or set a specified error message.

	XML XPath	Extract value or fragment from XML data using XPath functionality.
		For this option to work, Zabbix server must be compiled with libxml support.
		Examples:
		`number(/document/item/value)` will extract 10 from `<document><item><value>10</value></item></document>`
		`number(/document/item/@attribute)` will extract 10 from `<document><item attribute="10"></item></document>`
		`/document/item` will extract `<item><value>10</value></item>` from `<document><item><value>10</value></item></document>`
		Note that namespaces are not supported.
		If you mark the *Custom on fail* checkbox, it is possible to specify custom error-handling options: either to discard the value, set a specified value or set a specified error message.

	CSV to JSON	Convert CSV file data into JSON format. For more information, see: CSV to JSON preprocessing.
	XML to JSON	Convert data in XML format to JSON. For more information, see: Serialization rules. If you mark the *Custom on fail* checkbox, it is possible to specify custom error-handling options: either to discard the value, set a specified value or set a specified error message.
SNMP		
	SNMP walk value	Extract value by the specified OID/MIB name and apply formatting options: **Unchanged** - return Hex-STRING as unescaped hex string (*note* that display hints are still applied); **UTF-8 from Hex-STRING** - convert Hex-STRING to UTF-8 string; **MAC from Hex-STRING** - convert Hex-STRING to MAC address string (which will have `' '` replaced by `':'`); **Integer from BITS** - convert the first 8 bytes of a bit string expressed as a sequence of hex characters (e.g., "1A 2B 3C 4D") into a 64-bit unsigned integer. In bit strings longer than 8 bytes, consequent bytes will be ignored. If you mark the *Custom on fail* checkbox, it is possible to specify custom error-handling options: either to discard the value, set a specified value or set a specified error message.
	SNMP walk to JSON	Convert SNMP values to JSON. Specify a field name in the JSON and the corresponding SNMP OID path. Field values will be populated by values in the specified SNMP OID path. You may use this preprocessing step for SNMP OID discovery. Similar value formatting options as in the *SNMP walk value* step are available. If you mark the *Custom on fail* checkbox, it is possible to specify custom error-handling options: either to discard the value, set a specified value or set a specified error message.

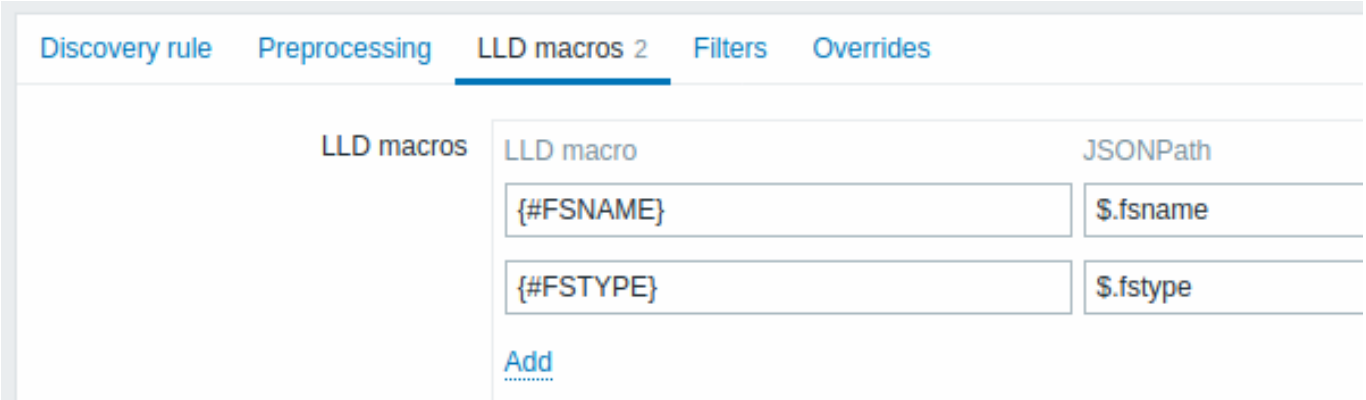
Type	
SNMP get value	Apply formatting options to the SNMP get value: **UTF-8 from Hex-STRING** - convert Hex-STRING to UTF-8 string; **MAC from Hex-STRING** - convert Hex-STRING to MAC address string (which will have `' '` replaced by `':'`); **Integer from BITS** - convert the first 8 bytes of a bit string expressed as a sequence of hex characters (e.g., "1A 2B 3C 4D") into a 64-bit unsigned integer. In bit strings longer than 8 bytes, consequent bytes will be ignored. If you mark the *Custom on fail* checkbox, it is possible to specify custom error-handling options: either to discard the value, set a specified value or set a specified error message.
Custom scripts	
JavaScript	Enter JavaScript code in the modal editor that opens when clicking in the parameter field or on the pencil icon next to it. Note that available JavaScript length depends on the database used. For more information, see: Javascript preprocessing
Validation	
Does not match regular expression	Specify a regular expression that a value must not match. E.g., `Error:(.*?)\.` If you mark the *Custom on fail* checkbox, it is possible to specify custom error-handling options: either to discard the value, set a specified value or set a specified error message.
Check for error in JSON	Check for an application-level error message located at JSONPath. Stop processing if succeeded and message is not empty; otherwise, continue processing with the value that was before this preprocessing step. Note that these external service errors are reported to user as is, without adding preprocessing step information. E.g., `$.errors`. If a JSON like `{"errors":"e1"}` is received, the next preprocessing step will not be executed. If you mark the *Custom on fail* checkbox, it is possible to specify custom error-handling options: either to discard the value, set a specified value or set a specified error message.
Check for error in XML	Check for an application-level error message located at Xpath. Stop processing if succeeded and message is not empty; otherwise, continue processing with the value that was before this preprocessing step. Note that these external service errors are reported to user as is, without adding preprocessing step information. No error will be reported in case of failing to parse invalid XML. If you mark the *Custom on fail* checkbox, it is possible to specify custom error-handling options: either to discard the value, set a specified value or set a specified error message.
Matches regular expression	Specify a regular expression that a value must match. If you mark the *Custom on fail* checkbox, it is possible to specify custom error-handling options: either to discard the value, set a specified value or set a specified error message.
Throttling	
Discard unchanged with heartbeat	Discard a value if it has not changed within the defined time period (in seconds). Positive integer values are supported to specify the seconds (minimum - 1 second). Time suffixes can be used in this field (e.g., `30s`, `1m`, `2h`, `1d`). User macros and low-level discovery macros can be used in this field. Only one throttling option can be specified for a discovery item. E.g., `1m`. If identical text is passed into this rule twice within 60 seconds, it will be discarded. *Note*: Changing item prototypes does not reset throttling. Throttling is reset only when preprocessing steps are changed.
Prometheus	
Prometheus to JSON	Convert required Prometheus metrics to JSON. See Prometheus checks for more details.

Note that if the discovery rule has been applied to the host via template then the content of this tab is read-only.

Custom macros

The **LLD macros** tab allows you to specify custom LLD macros when the returned JSON does not already contain keys in the `{#MACRO}` format.

For example, the `vfs.fs.get` agent item key returns JSON filesystem data without any predefined LLD macros.
Here, you define the macros manually and map them to values in the JSON using JSONPath expressions:



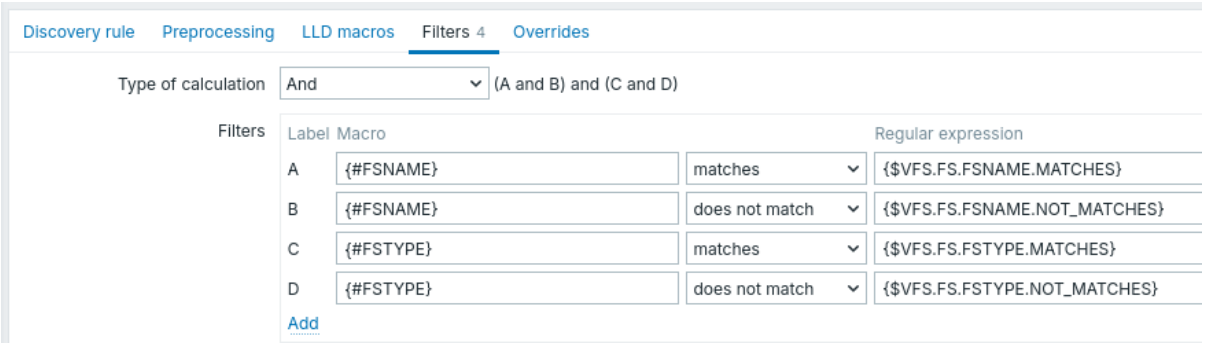
Parameter	Description
LLD macro	Name of the LLD macro in the format `{#MACRO}`.
JSONPath	Expression that defines which value to extract from each discovered entity in the JSON array. For example, `$.foo` extracts `"bar"` and `"baz"` from `[{"foo":"bar"}, {"foo":"baz"}]`. JSONPath supports both dot and bracket notation. Bracket notation must be used for keys containing special characters or Unicode, e.g., `$['unicode + special chars #1']['unicode + special chars #2']`.

The discovery rule must return a JSON array at the root (`$.`).
For backward compatibility, a JSON object containing a single `"data"` array is also accepted. In this case, the `"data"` array is automatically unwrapped:

```
{"data":[{"foo":"bar"}, {"foo":"baz"}]}
```

Filter

A filter can be used to generate real items, triggers, and graphs only for entities that match the criteria. The **Filters** tab contains discovery rule filter definitions that allow you to filter discovery values:



Parameter	Description
Type of calculation	The following options for calculating filters are available: **And** - all filters must be passed; **Or** - enough if one filter is passed; **And/Or** - uses *And* with different macro names and *Or* with the same macro name; **Custom expression** - offers the possibility to define a custom calculation of filters. The formula must include all filters in the list. Limited to 255 symbols.

Parameter	Description
Filters	The following filter condition operators are available: *matches*, *does not match*, *exists*, *does not exist*. *Matches* and *does not match* operators expect a Perl Compatible Regular Expression (PCRE). You can enter a regular expression or reference a global regular expression in the *Regular expression* field. *Exists* and *does not exist* operators allow filtering of entities based on the presence or absence of the specified LLD macro in the response. Note that if a macro from the filter is missing in the response, the found entity will be ignored, unless a "does not exist" condition is specified for this macro. A warning will be displayed, if the absence of a macro affects the expression result. For example, if {#B} is missing in: {#A} matches 1 and {#B} matches 2 - will give a warning {#A} matches 1 or {#B} matches 2 - no warning.

For example, if you are only interested in the C:, D:, and E: file systems, you can put {#FSNAME} into the *Macro* field and ^C|^D|^E into the *Regular expression* field.
 Filtering is also possible by file system types using the {#FSTYPE} macro (e.g., ^ext|^reiserfs) and by drive types (supported only by Windows agent) using the {#FSDRIVETYPE} macro (e.g., fixed).

In order to test a regular expression you can use grep -E, for example:

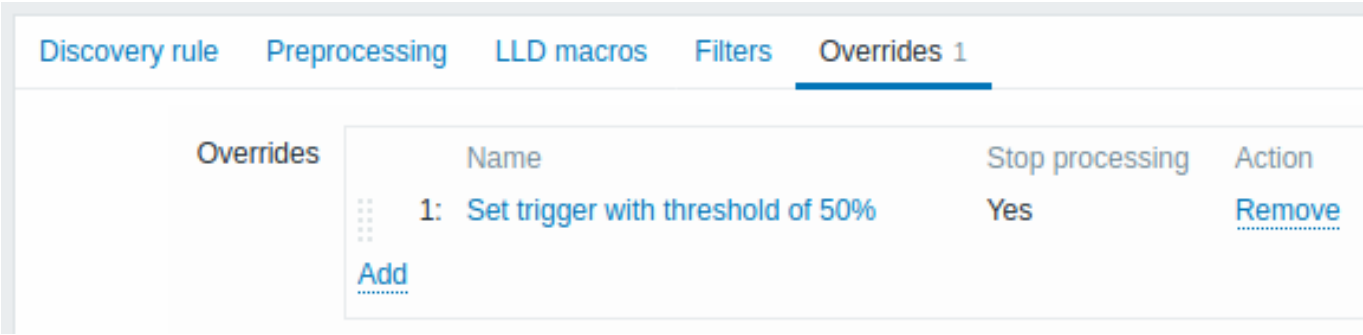
```
for f in ext2 nfs reiserfs smbfs;
do echo $f | grep -E '^ext|^reiserfs' || echo "SKIP: $f";
done
```

Warning:
A mistake or a typo in the regular expression used in the LLD rule (for example, an incorrect "File systems for discovery" regular expression) may cause deletion of thousands of configuration elements, historical values, and events for many hosts.

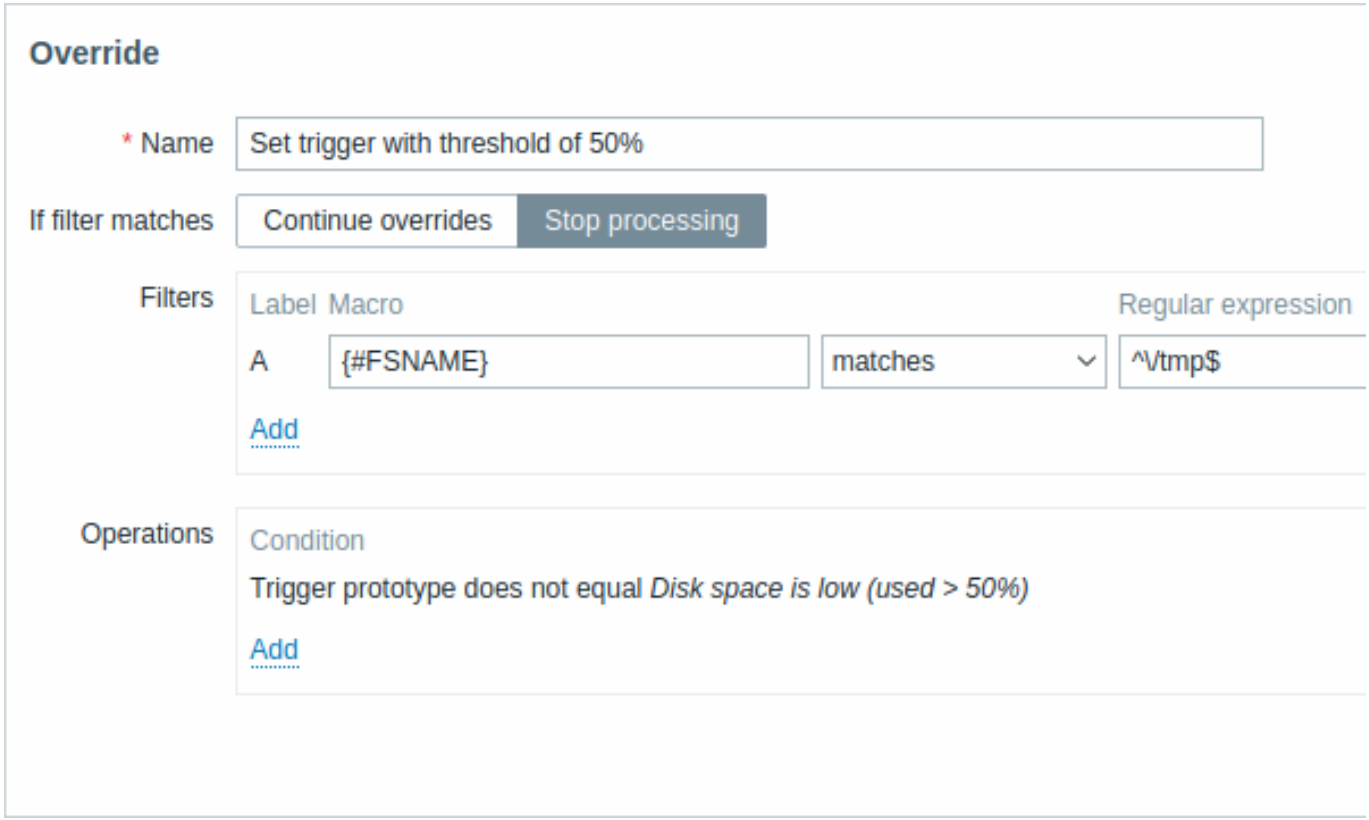
Attention:
Zabbix database in MySQL must be created as case-sensitive if file system names that differ only by case are to be discovered correctly.

Override

The **Overrides** tab allows setting rules to modify the list of item, trigger, graph, host, and discovery prototypes or their attributes for discovered objects that meet given criteria.



Overrides (if any) are displayed in a reorderable drag-and-drop list and executed in the order in which they are defined. To configure details of a new override, click Add in the *Overrides* block. To edit an existing override, click the override name. A popup window will open, allowing you to edit the override rule details.

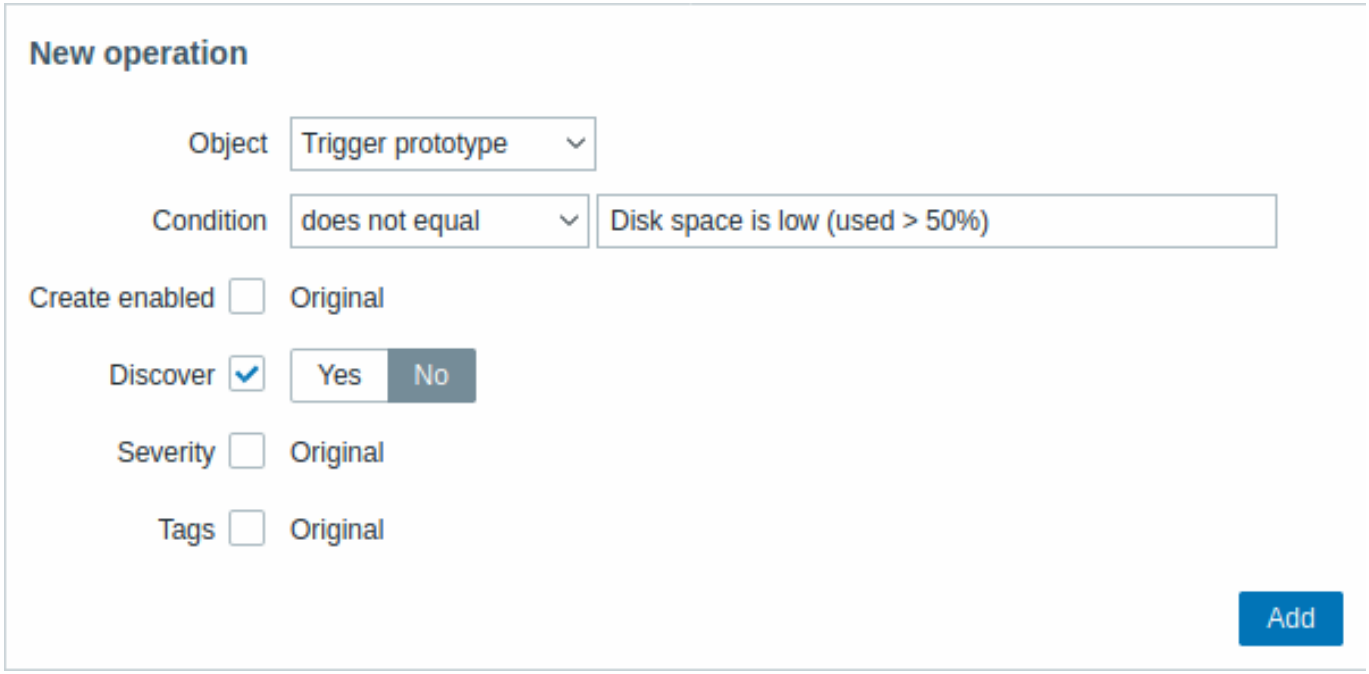


All mandatory parameters are marked with red asterisks.

Parameter	Description
Name	A unique (per LLD rule) override name.
If filter matches	Defines whether next overrides should be processed when filter conditions are met: **Continue overrides** - subsequent overrides will be processed. **Stop processing** - operations from preceding (if any) and this override will be executed, subsequent overrides will be ignored for matched LLD rows.
Filters	Determines to which discovered entities the override should be applied. Override filters are processed after discovery rule filters and have the same functionality.
Operations	Override operations are displayed with these details: **Condition** - an object type and a condition to be met for the object name; for example: Trigger prototype does not equal *Disk space is low (used > 50%)*. **Actions** - links for editing and removing an operation are displayed.

Configuring an operation

To configure details of a new operation, click Add in the Operations block. To edit an existing operation, click Edit next to the operation. A popup window where you can edit the operation details will open.



Parameter	Description
Object	Five types of objects are available: Item prototype Trigger prototype Graph prototype Host prototype Discovery prototype
Condition	Allows filtering entities to which the operation should be applied.
Operator	Supported operators: **equals** - apply to this prototype **does not equal** - apply to all prototypes, except this **contains** - apply, if prototype name contains this string **does not contain** - apply, if prototype name does not contain this string **matches** - apply, if prototype name matches regular expression **does not match** - apply, if prototype name does not match regular expression
Pattern	A regular expression or a string to search for.
Object: *Item prototype*	
Create enabled	When the checkbox is marked, the buttons will appear, allowing you to override original item prototype settings: *Yes* - the item will be added in an enabled state. *No* - the item will be added to a discovered entity but in a disabled state.
Discover	When the checkbox is marked, the buttons will appear, allowing you to override original item prototype settings: *Yes* - the item will be added. *No* - the item will not be added.
Update interval	When the checkbox is marked, two options will appear, allowing you to set different interval for the item: *Delay* - Item update interval. User macros, LLD macros, and time suffixes (e.g., 30s, 1m, 2h, 1d) are supported. Should be set to 0 if *Custom interval* is used. *Custom interval* - click Add to specify flexible/scheduling intervals. For detailed information see Custom intervals.

Parameter		Description
	History	When the checkbox is marked, the buttons will appear, allowing you to set different history storage periods for the item: *Do not store* - if selected, the history will not be stored. *Store up to* - if selected, an input field for specifying storage period will appear to the right. User macros and LLD macros are supported.
	Trends	When the checkbox is marked, the buttons will appear, allowing you to set different trend storage periods for the item: *Do not store* - if selected, the trends will not be stored. *Store up to* - if selected, an input field for specifying storage period will appear to the right. User macros and LLD macros are supported.
	Tags	When the checkbox is marked, a new block will appear, allowing you to specify tag-value pairs. User macros and LLD macros are supported. These tags will be appended to the tags specified in the item prototype, even if the tag names match.
Object: *Trigger prototype*		
	Create enabled	When the checkbox is marked, the buttons will appear, allowing you to override original trigger prototype settings: *Yes* - the trigger will be added in an enabled state. *No* - the trigger will be added to a discovered entity, but in a disabled state.
	Discover	When the checkbox is marked, the buttons will appear, allowing you to override original trigger prototype settings: *Yes* - the trigger will be added. *No* - the trigger will not be added.
	Severity	When the checkbox is marked, trigger severity buttons will appear, allowing you to modify trigger severity.
	Tags	When the checkbox is marked, a new block will appear, allowing you to specify tag-value pairs. User macros and LLD macros are supported. These tags will be appended to the tags specified in the trigger prototype, even if the tag names match.
Object: *Graph prototype*		
	Discover	When the checkbox is marked, the buttons will appear, allowing you to override original graph prototype settings: *Yes* - the graph will be added. *No* - the graph will not be added.
Object: *Host prototype*		
	Create enabled	When the checkbox is marked, the buttons will appear, allowing you to override original host prototype settings: *Yes* - the host will be created in an enabled state. *No* - the host will be created in a disabled state.
	Discover	When the checkbox is marked, the buttons will appear, allowing you to original host prototype settings: *Yes* - the host will be discovered. *No* - the host will not be discovered.
	Link templates	When the checkbox is marked, an input field for specifying templates will appear. Start typing the template name or click *Select* next to the field and select templates from the list in a popup window. Templates from this override are added to all templates already linked to the host prototype.

Parameter		Description
	Tags	When the checkbox is marked, a new block will appear, allowing you to specify tag-value pairs. User macros and LLD macros are supported. These tags will be appended to the tags specified in the host prototype, even if the tag names match.
	Host inventory	When the checkbox is marked, the buttons will appear, allowing you to select different inventory mode for the host prototype: *Disabled* - do not populate host inventory *Manual* - provide details manually *Automated* - auto-fill host inventory data based on collected metrics.
Object: *Discovery prototype*		
	Create enabled	When the checkbox is marked, the buttons will appear, allowing you to override original discovery prototype settings: *Yes* - the discovery rule will be added in an enabled state. *No* - the discovery rule will be added in a disabled state.
	Discover	When the checkbox is marked, the buttons will appear, allowing you to override original discovery prototype settings: *Yes* - the discovery rule will be added. *No* - the discovery rule will not be added.
	Update interval	When the checkbox is marked, two options will appear, allowing you to set different interval for the discovery rule: *Delay* - Rule update interval. User macros, LLD macros, and time suffixes (e.g., 30s, 1m, 2h, 1d) are supported. Should be set to 0 if *Custom interval* is used. *Custom interval* - click Add to specify flexible/scheduling intervals. For detailed information see Custom intervals.

Form buttons

Buttons at the bottom of the form allow you to perform several operations.

Add	Add a discovery rule. This button is only available for new discovery rules.
Update	Update the properties of a discovery rule. This button is only available for existing discovery rules.
Clone	Create another discovery rule based on the properties of the current discovery rule.
Execute now	Perform discovery based on the discovery rule immediately. The discovery rule must already exist. See more details. *Note* that when performing discovery immediately, configuration cache is not updated, thus the result will not reflect very recent changes to discovery rule configuration.
Test	Test the discovery rule configuration. Use this button to verify the configuration settings (such as connectivity and parameter correctness) without permanently applying any changes.
Delete	Delete the discovery rule.
Cancel	Cancel the editing of discovery rule properties.

Discovered entities The screenshots below illustrate how discovered items, triggers, and graphs look like in the host's configuration. Discovered entities are prefixed with an orange link to a discovery rule they come from.

Items

All hosts / Zabbix server Enabled ZBX SNMP IPMI JMX Items 140 Triggers 77 Graphs 14 Discovery rules 6 Web scenarios

Name ▲	Triggers	Key	Interval
Mounted filesystem discovery: Get filesystems: FS [/]: Get data		vfs.fs.dependent[/,data]	
Mounted filesystem discovery: FS [/]: Get data: FS [/]: Inodes: Free, in %	Triggers 2	vfs.fs.dependent.inode[/,pfree]	
Mounted filesystem discovery: FS [/]: Get data: FS [/]: Option: Read-only	Triggers 1	vfs.fs.dependent[/,readonly]	
Mounted filesystem discovery: FS [/]: Get data: FS [/]: Space: Available		vfs.fs.dependent.size[/,free]	
Mounted filesystem discovery: FS [/]: Get data: FS [/]: Space: Total		vfs.fs.dependent.size[/,total]	
Mounted filesystem discovery: FS [/]: Get data: FS [/]: Space: Used		vfs.fs.dependent.size[/,used]	
Mounted filesystem discovery: FS [/]: Get data: FS [/]: Space: Used, in %	Triggers 2	vfs.fs.dependent.size[/,pused]	

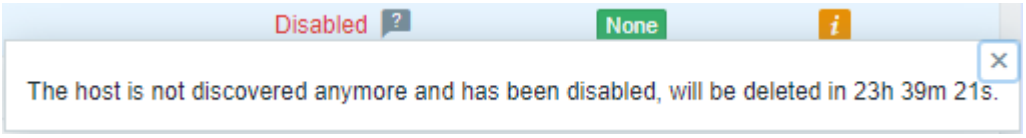
Note that discovered entities will not be created in case there are already existing entities with the same uniqueness criteria, for example, an item with the same key or graph with the same name. An error message is displayed in this case in the frontend that the low-level discovery rule could not create certain entities. The discovery rule itself, however, will not turn unsupported because some entity could not be created and had to be skipped. The discovery rule will go on creating/updating other entities.

If a discovered entity (host, file system, interface, etc) stops being discovered (or does not pass the filter anymore) the entities that were created based on it may be automatically disabled and eventually deleted.

Lost resources may be automatically disabled based on the value of the *Disable lost resources* parameter. This affects lost hosts, items and triggers.

Lost resources may be automatically deleted based on the value of the *Delete lost resources* parameter. This affects lost hosts, host groups, items, triggers, and graphs.

When discovered entities become 'Not discovered anymore', a lifetime indicator is displayed in the entity list. Move your mouse pointer over it and a message will be displayed indicating its status details.



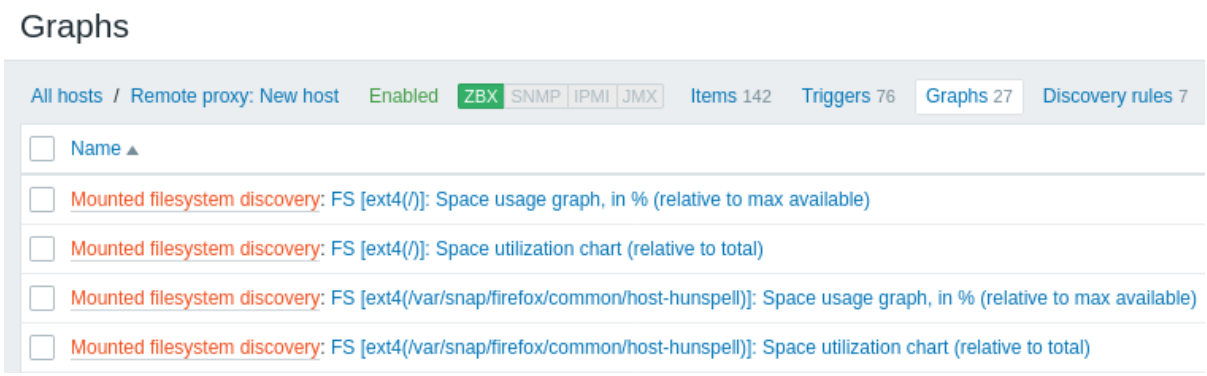
If entities were marked for deletion, but were not deleted at the expected time (disabled discovery rule or item host), they will be deleted the next time the discovery rule is processed.

Entities containing other entities, which are marked for deletion, will not update if changed on the discovery rule level. For example, LLD-based triggers will not update if they contain items that are marked for deletion.

Triggers

All hosts / Zabbix server Enabled ZBX SNMP IPMI JMX Items 140 Triggers 77 Graphs 14 Discovery rules 6

Severity	Value	Name ▲	Operational data
Warning	OK	Mounted filesystem discovery: Linux: FS [/]: Running out of free inodes **Depends on:** Zabbix server: Linux: FS [/]: Running out of free inodes	Free inodes: {ITEM.LASTVALUE1}
Warning	OK	Mounted filesystem discovery: Linux: FS [/]: Space is low **Depends on:** Zabbix server: Linux: FS [/]: Space is critically low	Space used: {{ITEM.LASTVALUE1}.fmtnum(1)}%



Other types of discovery More detail and how-tos on other types of out-of-the-box discovery is available in the following sections:

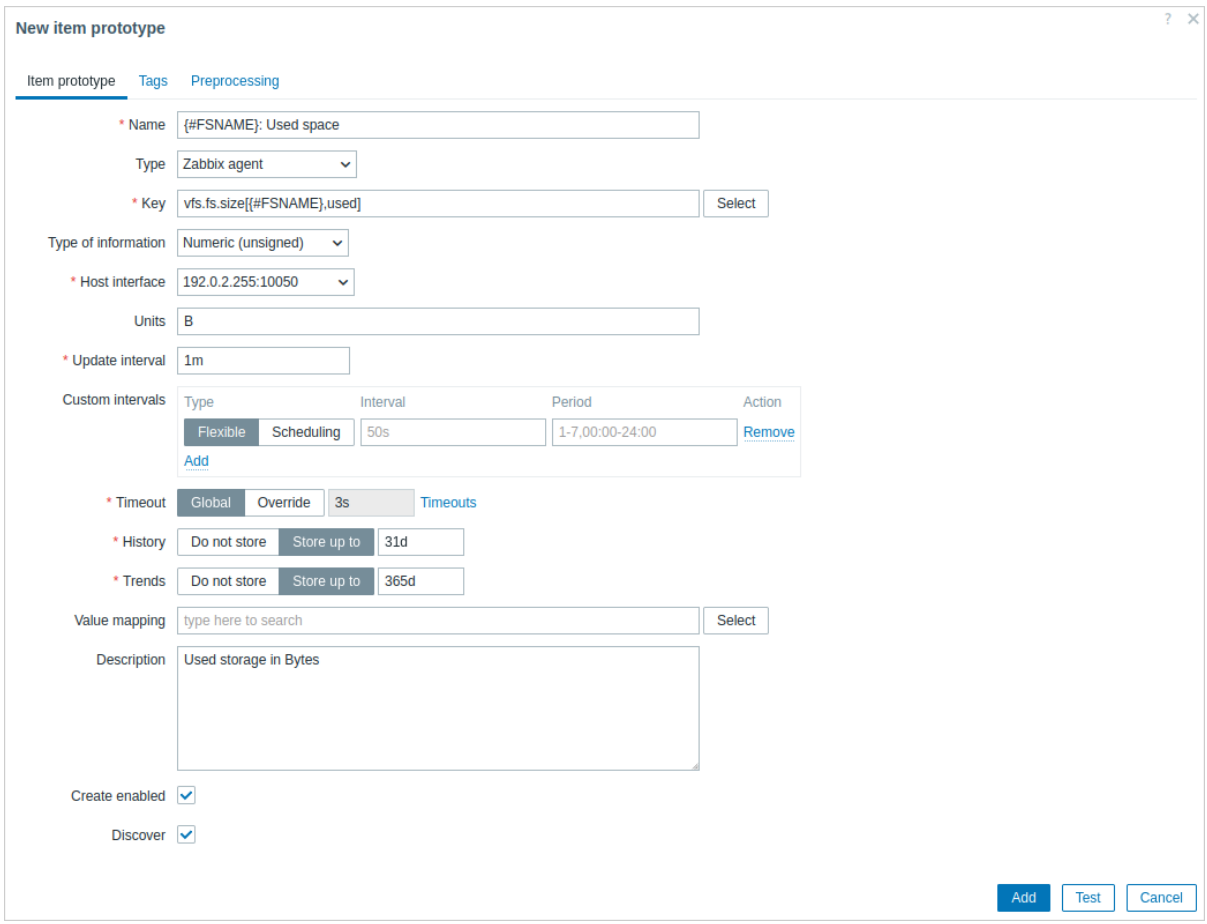
- discovery of network interfaces
- discovery of CPUs and CPU cores
- discovery of SNMP OIDs
- discovery of JMX objects;
- discovery using ODBC SQL queries
- discovery of Windows services
- discovery of host interfaces in Zabbix

For more detail on the JSON format for discovery items and an example of how to implement your own file system discoverer as a Perl script, see creating custom LLD rules.

Item prototypes

Once a rule is created, go to the items for that rule and click *Create item prototype* to create an item prototype.

Note how the `{#FSNAME}` macro is used where a file system name is required. The use of a low-level discovery macro is mandatory in the item key to make sure that the discovery is processed correctly. When the discovery rule is processed, this macro will be substituted with the discovered file system.



Note:
When item key parameters are not used, place the LLD macro inside the item key brackets `[...]` (for example, `v_[{#MACRO}]`).

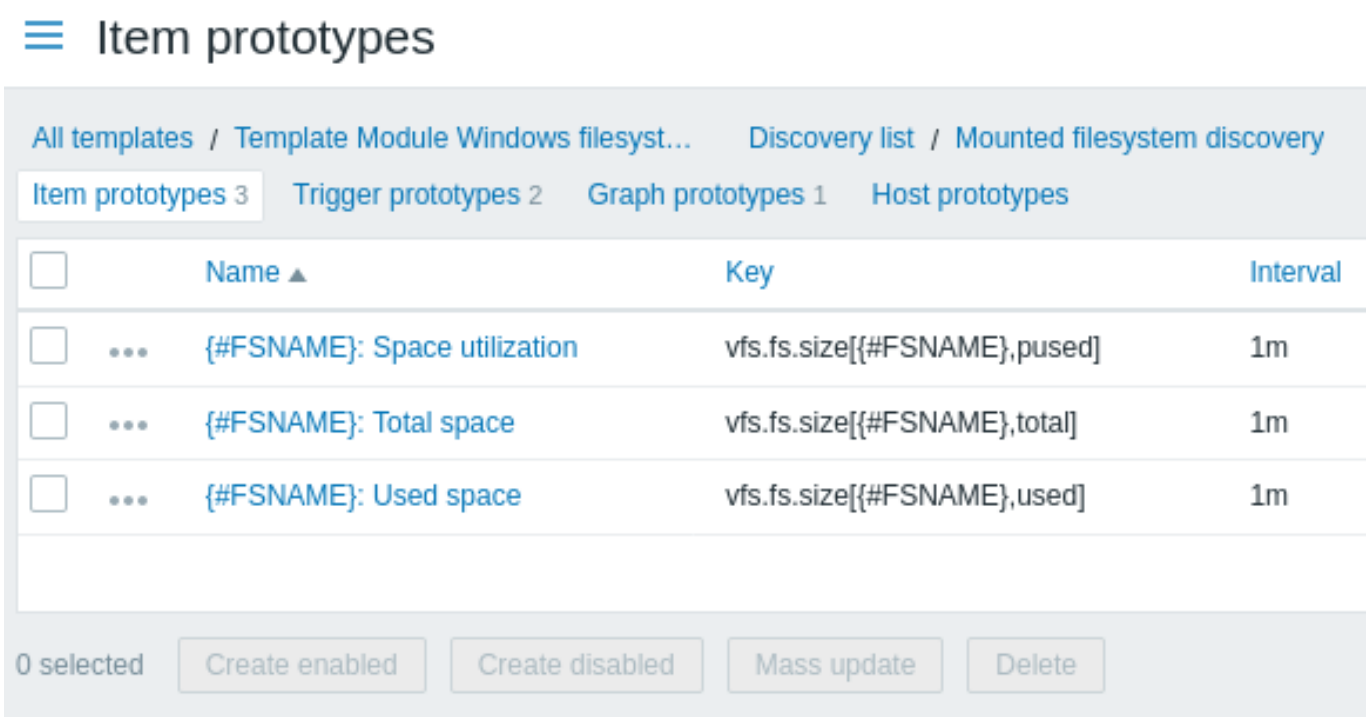
Low-level discovery macros and user macros are supported in item prototype configuration and item value preprocessing parameters.Note that when used in update intervals, a single macro has to fill the whole field. Multiple macros in one field or macros mixed with text are not supported.

Note:
Context-specific escaping of low-level discovery macros is performed for safe use in regular expression and XPath preprocessing parameters.

Attributes that are specific for item prototypes:

Parameter	Description
Create enabled	If checked the item will be added in an enabled state. If unchecked, the item will be added to a discovered entity, but in a disabled state.
Discover	If checked (default) the item will be added to a discovered entity. If unchecked, the item will not be added to a discovered entity, unless this setting is overridden in the discovery rule.

We can create several item prototypes for each file system metric we are interested in:



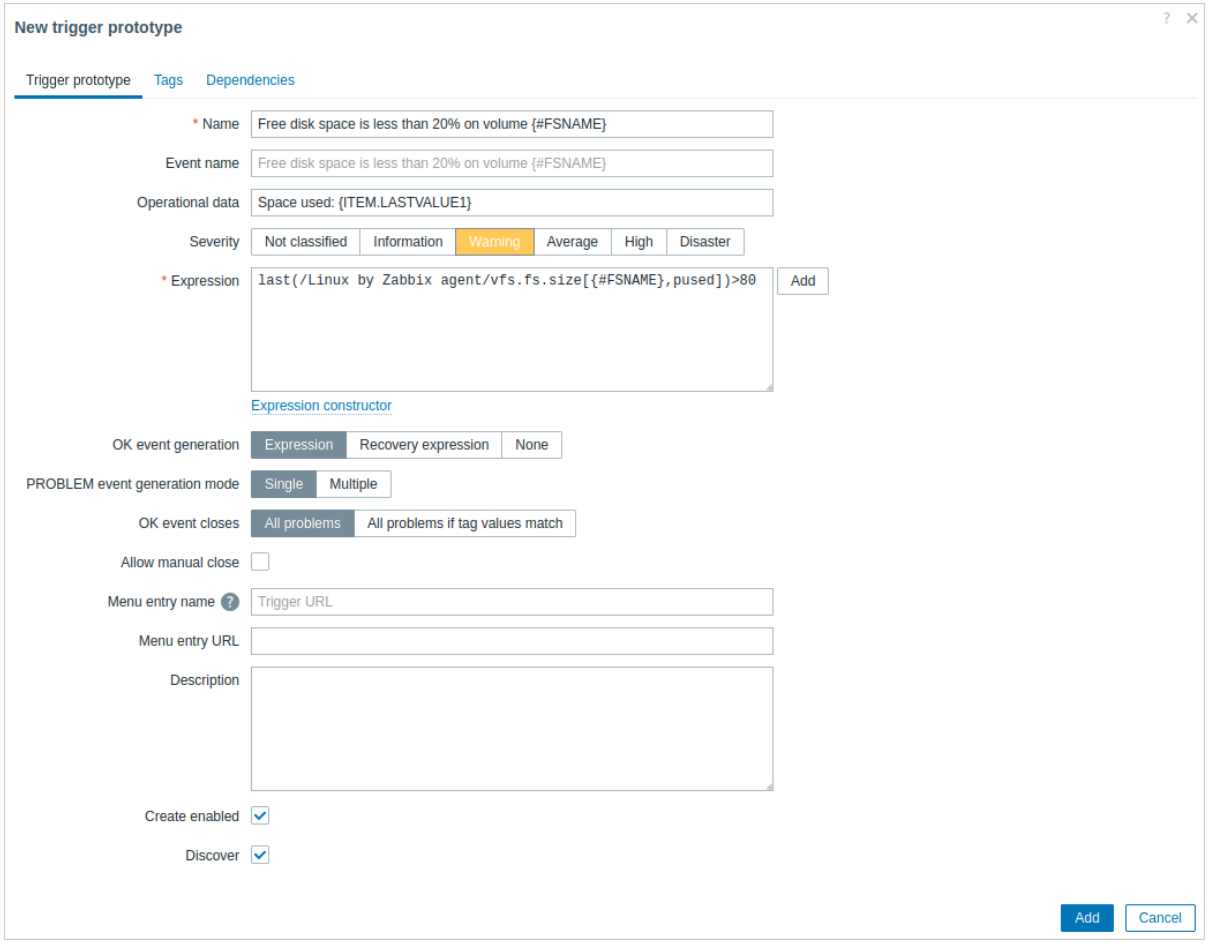
Click the three-dot icon to open the menu for the specific item prototype with these options:

- *Create trigger prototype* - create a trigger prototype based on this item prototype
- *Trigger prototypes* - click to see a list with links to already-configured trigger prototypes of this item prototype
- *Create dependent item* - create a dependent item for this item prototype

Mass update option is available if you want to update properties of several item prototypes at once.

Trigger prototypes

We create trigger prototypes in a similar way as item prototypes:

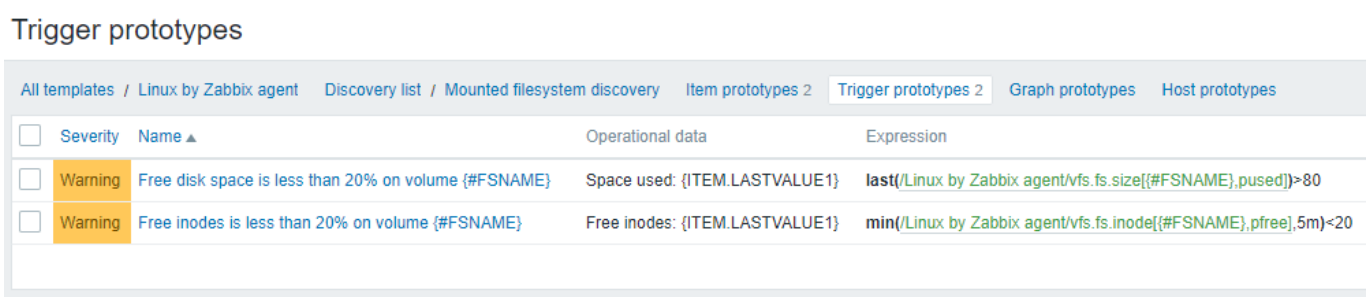


Attributes that are specific for trigger prototypes:

Parameter	Description
Create enabled	If checked the trigger will be added in an enabled state. If unchecked, the trigger will be added to a discovered entity, but in a disabled state.
Discover	If checked (default) the trigger will be added to a discovered entity. If unchecked, the trigger will not be added to a discovered entity, unless this setting is overridden in the discovery rule.

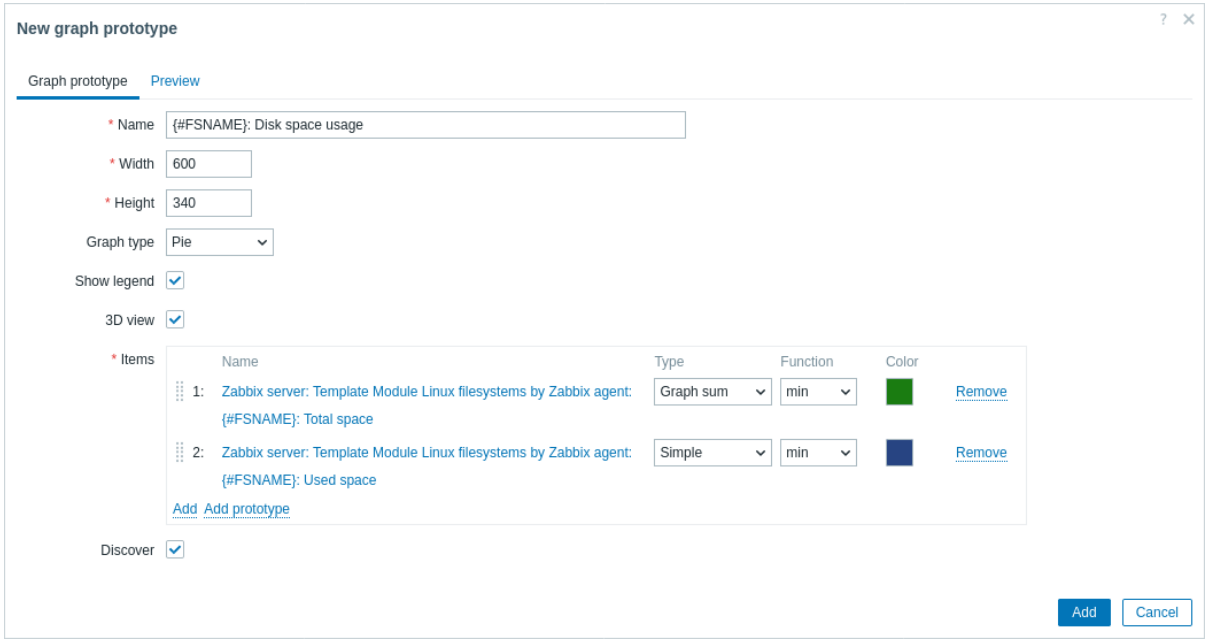
When real triggers are created from the prototypes, there may be a need to be flexible as to what constant ('20' in our example) is used for comparison in the expression. See how user macros with context can be useful to accomplish such flexibility.

You can define dependencies between trigger prototypes. To do that, go to the *Dependencies* tab. A trigger prototype may depend on another trigger prototype from the same low-level discovery (LLD) rule or on a regular trigger. A trigger prototype may not depend on a trigger prototype from a different LLD rule or on a trigger created from trigger prototype. Host trigger prototype cannot depend on a trigger from a template.



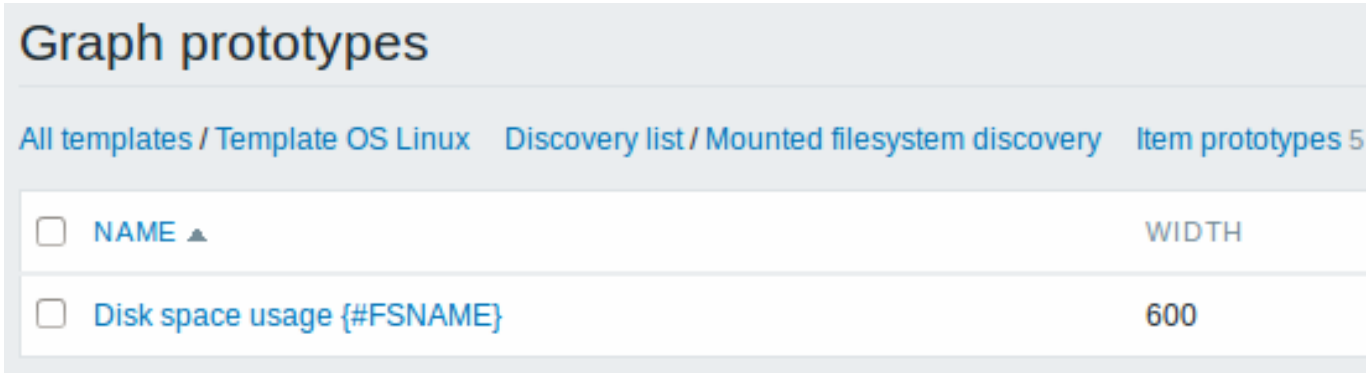
Graph prototypes

We can create graph prototypes, too:

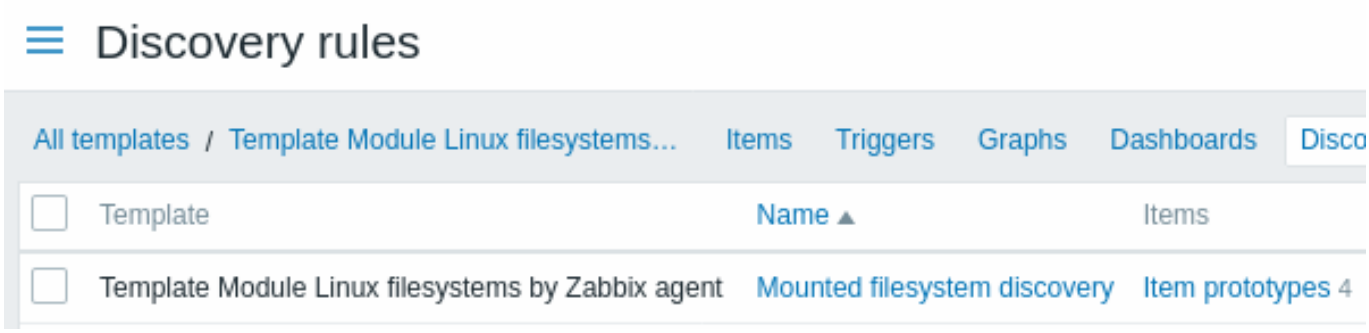


Attributes that are specific for graph prototypes:

Parameter	Description
Discover	If checked (default) the graph will be added to a discovered entity. If unchecked, the graph will not be added to a discovered entity, unless this setting is overridden in the discovery rule.



Finally, we have created a discovery rule that looks as shown below. It has five item prototypes, two trigger prototypes, and one graph prototype.



Host prototypes

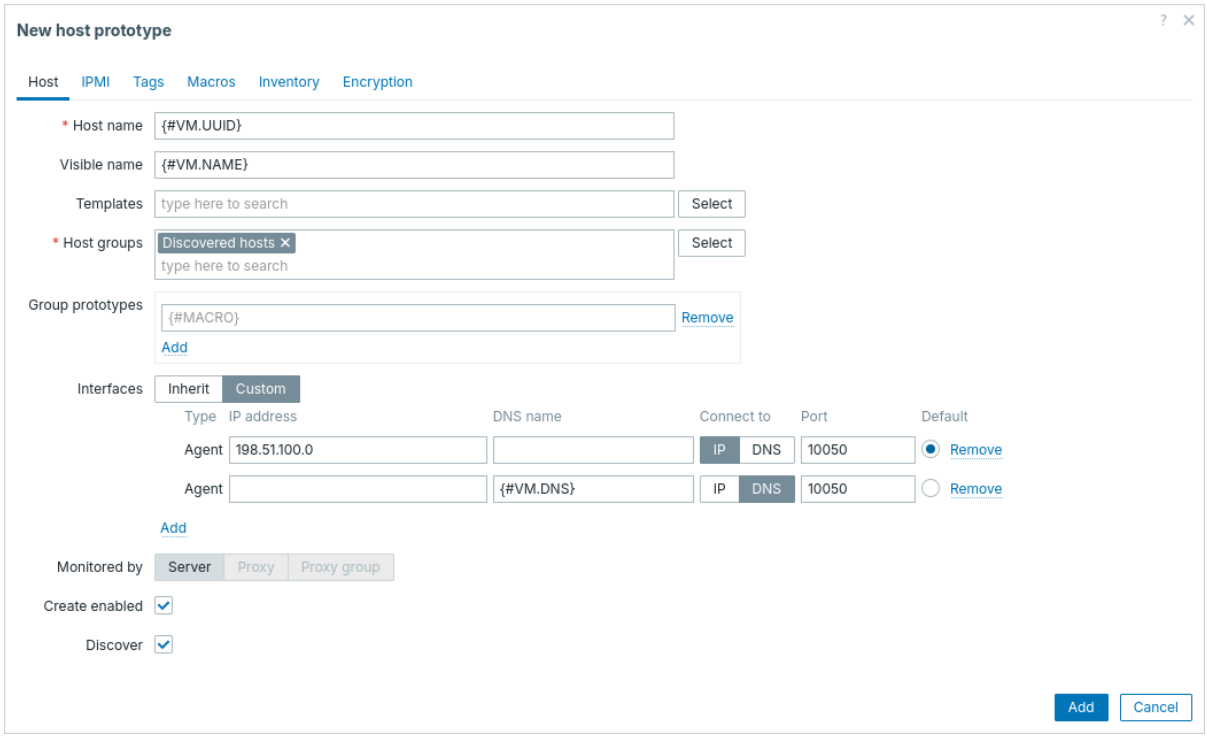
Host prototypes are blueprints for creating hosts through low-level discovery rules. Before being discovered as hosts, these prototypes cannot have items and triggers, except those linked from templates.

Configuration

Host prototypes are configured under low-level discovery rules.

To create a host prototype:

1. Go to *Data collection > Hosts*.
2. Click *Discovery* for the required host to navigate to the list of low-level discovery rules configured for that host.
3. Click *Host prototypes* for the required discovery rule.
4. Click the *Create host prototype* button in the upper-right corner and complete the configuration form.



All mandatory input fields are marked with a red asterisk.

Parameter	Description
Host name	This parameter must contain at least one low-level discovery macro to ensure unique host names for created hosts.
Visible name	Low-level discovery macros are supported.
Group prototypes	Allows specifying host group prototypes by using low-level discovery macros. Based on the specified group prototypes, host groups will be discovered, created and linked to the created hosts; discovered groups that have already been created by other low-level discovery rules will also be linked to the created hosts. However, discovered host groups that match manually created host groups will not be linked to the created hosts.
Interfaces	Set whether discovered hosts inherit the IP from the host that the discovery rule belongs to (default), or get custom interfaces. Low-level discovery macros and user macros are supported.
Monitored by	This field is synced automatically with the respective field in host configuration.
Create enabled	Set the status of discovered hosts; if unchecked, hosts will be created as disabled.
Discover	Set whether hosts will be created from the host prototype; if unchecked, hosts will not be created from the host prototype (unless this setting is overridden in the low-level discovery rule).

Note:
Low-level discovery macros are also supported for tag values and host prototype user macro values.
 Value maps are not supported for host prototypes.

For an example of how to configure a host prototype, see *Virtual machine monitoring*.

Host interfaces

To add custom interfaces, switch the *Interfaces* selector from ”Inherit” to ”Custom”. Click Add and select the interface type - Zabbix agent, SNMP, JMX, IPMI.

Note:
If *Custom* is selected, but no interfaces have been set, the hosts will be created without interfaces.
 If *Inherit* is selected and the host prototype belongs to a template, all discovered hosts will inherit the host interface from the host to which the template is linked.

If multiple custom interfaces are specified, the primary interface can be set in the *Default* column.

For an example of how to configure a custom host interface, see *VMware monitoring setup example*.

Warning:
A host will only be created if a host interface contains correct data.

Discovered hosts

In the host list, discovered hosts are prefixed with the name of the discovery rule that created them.

Discovered hosts inherit most parameters from host prototypes as *read-only*. Only the following discovered host parameters can be configured:

- *Templates* - link new templates or unlink manually added ones. Templates inherited from a host prototype cannot be unlinked.
- *Description* - add or edit the host description.
- *Status* - enable or disable the host.
- *Tags* - add new tags or remove manually added ones. Tags inherited from a host prototype cannot be removed. Note that tags cannot have duplicates (tags with the same name and value). If a new tag is added to the host prototype and it matches a manually added tag on a discovered host, the manually added tag will be replaced with the inherited tag during discovery.
- *Macros* - add new host macros or remove manually added ones, as well as those inherited from a host prototype; change macro values and types.

Note:
Discovered hosts also inherit user macros from the host where the discovery rule is configured. These inherited macros can be removed, and their values and types can be changed.

Discovered hosts can also discover other hosts. For example, if a host prototype (used to discover hypervisors) has a template linked to it that includes a low-level discovery rule with its own host prototype (used to discover virtual machines), Zabbix will discover both hypervisors and their virtual machines. To additionally discover containers on those VMs, you can create low-level discovery rules or link a new template with host prototypes to the discovered VMs or pre-configure the VM host prototype by linking a template that itself includes host prototypes.

Discovered hosts can be deleted manually. Note, however, that they will be discovered again if discovery is enabled for them.

Hosts that are no longer discovered may be:

- automatically disabled (based on the *Disable lost resources* value of the discovery rule)
- automatically deleted (based on the *Delete lost resources* value of the discovery rule).

Discovery prototypes

Overview

Discovery prototypes are **nested low-level discovery** rules within a "parent" discovery rule, allowing multi-level discovery of objects with their own items, triggers, etc. For example, you may want to discover all database instances on a database server, then discover tablespaces for each instance, then discover tables for each tablespace.

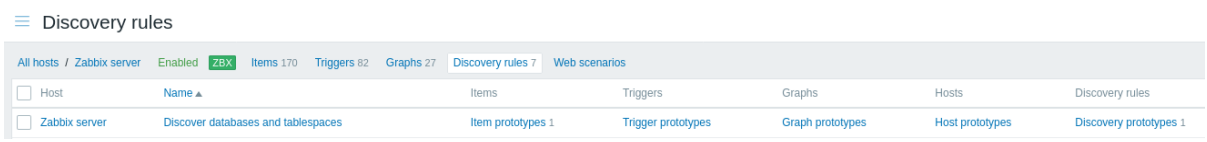
Discovery prototypes have their own item, trigger, graph, host, and discovery prototypes. A nested discovery prototype will use the same JSON value as the parent rule if you specify a *Nested* type.

The levels of nesting for discovery prototypes are unlimited.

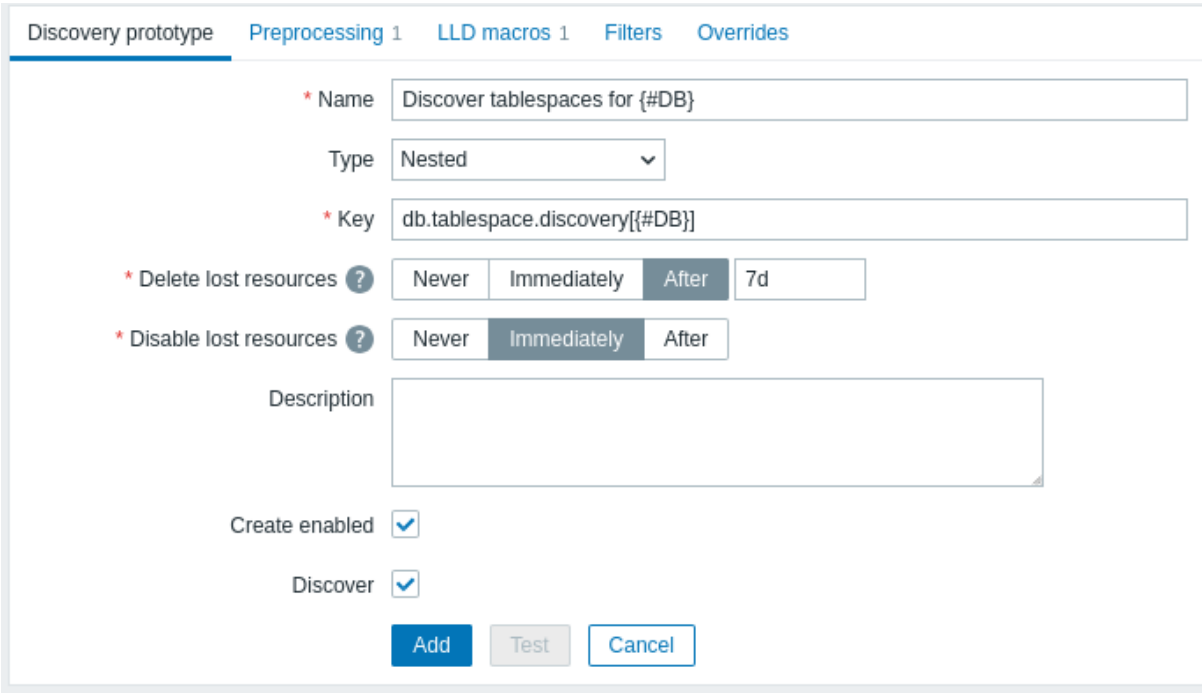
Configuration

To create a discovery prototype:

- Click *Discovery prototypes* in the row of an existing discovery rule



- Click *Create discovery prototype*



Configuration fields of this form are shared with the regular low-level discovery.

If you select "Nested" as *Type* in the opened discovery prototype form, then discovery rules (from the discovery prototype) are generated based on a JSON object from the same JSON value as the parent discovery rule. For example, if the original JSON is [<object A>, <object B>] and there is one nested discovery rule prototype then two discovery rules would be generated based on object A and object B data respectively.

In this case the discovery prototype is activated at the same time as the parent rule. The nested rule thus can use preprocessing, to work on another "slice" of the same data, already acquired by the parent.

LLD macros from the parent LLD rule are available for nested discovery rules.

Nested LLD rules on discovered hosts

A *Nested* low-level discovery rule can be used on a host template assigned to host prototype. If a *Nested* discovery rule exists on a discovered host, then the JSON object used to discover the host is also sent to all LLD rules of nested type on this host. For more details, see example.

LLD macros from the discovery rule that created the host are available for nested discovery rules.

Example

Let's illustrate possible application of discovery prototypes, based on receiving the following example of multi-level JSON.

```
[
  {
    "database": "db1",
    "created_at": "2024-02-01T12:30:00Z",
    "encoding": "UTF8",
    "tablespaces": [
      { "name": "ts1", "max_size": "10GB" },
      { "name": "ts2", "max_size": "20GB" },
      { "name": "ts3", "max_size": "15GB" }
    ]
  },
  {
    "database": "db2",
    "created_at": "2023-11-15T08:45:00Z",
    "encoding": "UTF16",
    "tablespaces": [
      { "name": "ts1", "max_size": "5GB" },
      { "name": "ts2", "max_size": "25GB" },
      { "name": "ts3", "max_size": "30GB" }
    ]
```

```
  },
  {
    "database": "db3",
    "created_at": "2024-01-05T15:10:00Z",
    "encoding": "UTF8",
    "tablespaces": [
      { "name": "ts1", "max_size": "12GB" },
      { "name": "ts2", "max_size": "18GB" },
      { "name": "ts3", "max_size": "22GB" }
    ]
  }
]
```

Case 1

Discovering database instances on a database server, then discovering the tablespaces for each instance.

1. You have at least one host related to database server discovery.
2. Create an LLD rule for this host named *Discover databases and tablespaces*.
3. Switch to the *LLD Macros* tab on this rule, add macro `{#DB}=$.database`.
4. Add an item prototype for this rule named *Active connections to {#DB}* (Type: Agent, Key: `db.connections[{#DB}]`).
5. The items related to each database are discovered:

```
Active connections to db1, Active connections to db2, Active connections to db3.
```

6. Create a discovery prototype for this rule named *Discover tablespaces for {#DB}* (Type: Nested, Key: `db.tablespace.discovery[{#`
7. Switch to the *Preprocessing* tab of this discovery prototype and add the step `JSONPath=$.tablespaces`.
8. Switch to the *LLD Macros* tab of this discovery prototype, add macro `{#TSNAME}=$.name`.
9. Create an item prototype for this discovery prototype named *Size of tablespace {#TSNAME} for {#DB}* (Type: Agent, Key: `db.ts.size[{#DB}, {#TSNAME}]`).
10. The items related to each tablespace of each database are discovered:

```
Size of tablespace ts1 for db1, Size of tablespace ts2 for db1, Size of tablespace ts3 for db1,
Size of tablespace ts1 for db2, Size of tablespace ts2 for db2, Size of tablespace ts3 for db2,
Size of tablespace ts1 for db3, Size of tablespace ts2 for db3, Size of tablespace ts3 for db3.
```

with keys `db.ts.size[db1,ts1]`, `db.ts.size[db1,ts2]`, ... `db.ts.size[db3,ts3]`.

Case 2

Discovering database instances on the database server by representing them as discovered hosts, then discovering the tablespaces for each instance.

1. You have at least one host (root host) related to database server discovery.
2. Create a template to discover the tablespaces for each database.
3. Create an item in this template named *Active connections to {#DB}* (Type: Agent, Key: `db.connections[{#DB}]`).
4. Create an LLD rule for this template named *Discover tablespaces* (Type: Nested).
5. Switch to the *Preprocessing* tab of this rule and add the step `JSONPath=$.tablespaces`.
6. Switch to the *LLD Macros* tab of this rule, add macro `{#TSNAME}=$.name`.
7. Create an item prototype for this rule named *Size of tablespace {#TSNAME} for {#DB}* (Type: Agent, Key: `db.ts.size[{#DB}, {#TSNAME}]`).
8. Back on the root host, create an LLD rule for this host named *Discover databases and tablespaces*.
9. Switch to the *LLD Macros* tab on this rule, add macro `{#DB}=$.database`.
10. Add a host prototype for this rule named *Host for database {#DB}*.
11. Switch to the *Macros* tab on this host prototype, add macro `{$DB}={#DB}` (for the item's name and key from Step 3).
12. Link the template from step 2 to this host prototype.
13. The discovered hosts contain the discovered items related to each database and its tablespaces:

Host	Items
Host for database db1	Active connections to db1
	Size of tablespace ts1 for db1
	Size of tablespace ts2 for db1
	Size of tablespace ts3 for db1
Host for database db2	Active connections to db2
	Size of tablespace ts1 for db2
	Size of tablespace ts2 for db2
	Size of tablespace ts3 for db2
Host for database db3	Active connections to db3
	Size of tablespace ts1 for db3
	Size of tablespace ts2 for db3
	Size of tablespace ts3 for db3

Notes on low-level discovery

Using LLD macros in user macro contexts

LLD macros may be used inside user macro context, for example, in trigger prototypes.

Multiple LLD rules for the same item

It is possible to define several low-level discovery rules with the same discovery item.

To do that you need to define the Alias agent parameter, allowing you to use altered discovery item keys in different discovery rules, for example, `vfs.fs.discovery[foo]`, `vfs.fs.discovery[bar]`, etc.

Data limits for return values

There is no limit for low-level discovery rule JSON data if it is received directly by Zabbix server. This is because the return values are processed without being stored in a database.

There is also no limit for custom low-level discovery rules. However, if custom low-level discovery rule data is retrieved using a user parameter, the user parameter return value limit applies.

If data has to go through Zabbix proxy, it has to store this data in the database. In such a case, database limits apply.

Discovery rule examples

This section provides practical examples of configuring discovery rules for different types of entities, such as filesystems, network interfaces, services, sensors, SNMP objects, Java applications, and databases.

Each example illustrates how a discovery rule can be structured, including the item key or query used, the JSON output with LLD macros, and how those results are used to create item, trigger, and graph prototypes.

- Discovery of block devices
- Discovery of CPUs and CPU cores
- Discovery of host interfaces in Zabbix
- Discovery of IPMI sensors
- Discovery of JMX objects
- Discovery of mounted filesystems
- Discovery of network interfaces
- Discovery of SNMP OIDs
- Discovery of SNMP OIDs (legacy)
- Discovery of systemd services
- Discovery of Windows performance counter instances
- Discovery of Windows services
- Discovery using WMI queries
- Discovery using ODBC SQL queries
- Discovery using Prometheus data

Custom LLD rules

Overview

It is also possible to create a completely custom LLD rule, discovering any type of entities - for example, databases on a database server.

To do so, a custom item should be created that returns a JSON string, specifying the found objects and optionally - some properties of them. The amount of macros per entity is not limited - while the built-in discovery rules return either one or two macros (for example, two for filesystem discovery), it is possible to return more.

Example

The required JSON string format is best illustrated with an example. Suppose you are running an old Zabbix 1.8 agent (one that does not support the `vfs.fs.discovery` key), but you still need to discover file systems. Here is a simple Perl script for Linux that discovers mounted file systems and outputs in JSON, which includes both file system name and type. One way to use it would be as a UserParameter with key `vfs.fs.discovery_perl`:

```
####!/usr/bin/perl

$first = 1;

print "[\n";

for (`cat /proc/mounts`)
{
    ($fsname, $fstype) = m/\S+ (\S+) (\S+)/;

    print "\t,\n" if not $first;
    $first = 0;

    print "\t{\n";
    print "\t\t\"{#FSNAME}\":\"$fsname\",\n";
    print "\t\t\"{#FSTYPE}\":\"$fstype\"\n";
    print "\t}\n";
}

print "]\n";
```

Attention:
Allowed symbols for LLD macro names are **0-9** , **A-Z** , _ , **.**. Lowercase letters are not supported in the names.

An example of its output (reformatted for clarity) is shown below. The JSON for custom discovery checks has to follow the same format.

```
[
    { "{#FSNAME}":"/",                          "{#FSTYPE}":"rootfs"   },
    { "{#FSNAME}":"/sys",                       "{#FSTYPE}":"sysfs"    },
    { "{#FSNAME}":"/proc",                      "{#FSTYPE}":"proc"     },
    { "{#FSNAME}":"/dev",                       "{#FSTYPE}":"devtmpfs" },
    { "{#FSNAME}":"/dev/pts",                   "{#FSTYPE}":"devpts"   },
    { "{#FSNAME}":"/lib/init/rw",               "{#FSTYPE}":"tmpfs"    },
    { "{#FSNAME}":"/dev/shm",                   "{#FSTYPE}":"tmpfs"    },
    { "{#FSNAME}":"/home",                      "{#FSTYPE}":"ext3"     },
    { "{#FSNAME}":"/tmp",                       "{#FSTYPE}":"ext3"     },
    { "{#FSNAME}":"/usr",                       "{#FSTYPE}":"ext3"     },
    { "{#FSNAME}":"/var",                       "{#FSTYPE}":"ext3"     },
    { "{#FSNAME}":"/sys/fs/fuse/connections",   "{#FSTYPE}":"fusectl"  }
]
```

In the previous example it is required that the keys match the LLD macro names used in prototypes, the alternative is to extract LLD macro values using JSONPath {#FSNAME} → `$.fsname` and {#FSTYPE} → `$.fstype`, thus making such script possible:

```
####!/usr/bin/perl

$first = 1;

print "[\n";

for (`cat /proc/mounts`)
```

```
{
    ($fsname, $fstype) = m/\S+ (\S+) (\S+)/;

    print "\t,\n" if not $first;
    $first = 0;

    print "\t{\n";
    print "\t\t\"fsname\":\"$fsname\",\n";
    print "\t\t\"fstype\":\"$fstype\"\n";
    print "\t}\n";
}

print "]\n";
```

An example of its output (reformatted for clarity) is shown below.

```
[
    { "fsname":"/",                          "fstype":"rootfs"   },
    { "fsname":"/sys",                       "fstype":"sysfs"    },
    { "fsname":"/proc",                      "fstype":"proc"     },
    { "fsname":"/dev",                       "fstype":"devtmpfs" },
    { "fsname":"/dev/pts",                   "fstype":"devpts"   },
    { "fsname":"/lib/init/rw",               "fstype":"tmpfs"    },
    { "fsname":"/dev/shm",                   "fstype":"tmpfs"    },
    { "fsname":"/home",                      "fstype":"ext3"     },
    { "fsname":"/tmp",                       "fstype":"ext3"     },
    { "fsname":"/usr",                       "fstype":"ext3"     },
    { "fsname":"/var",                       "fstype":"ext3"     },
    { "fsname":"/sys/fs/fuse/connections",   "fstype":"fusectl"  }
]
```

Then, in the discovery rule's *Filter* field, you may specify `{#FSTYPE}` as a macro and `rootfs|ext3` as a regular expression.

Note:
You don't have to use macro names like FSNAME/FSTYPE with custom LLD rules, you are free to use whatever names you like. In case JSONPath is used then the LLD row will be an array element that can be an object, but it can be also another array or a value.

Note that, if using a user parameter, the return value is limited to 16 MB. For more details, see data limits for LLD return values.

Virtual machine monitoring

Overview Zabbix can use low-level discovery rules to automatically discover VMware hypervisors and virtual machines, and create hosts to monitor them based on pre-defined host prototypes.

Zabbix also includes ready-to-use templates for monitoring VMware vCenter or ESXi hypervisors.

The minimum required VMware vCenter or vSphere version is 5.1.

Data collection Virtual machine monitoring consists of two steps:

1. Zabbix *vmware collector* processes collect virtual machine data - the processes obtain necessary information from VMware web services over the SOAP protocol, pre-process it, and store it in Zabbix server shared memory.
2. Zabbix *poller* processes retrieve data using Zabbix simple check VMware monitoring item keys.

Zabbix divides collected data into VMware configuration data and VMware performance counter data. Both types of data are collected independently by the *vmware collector* processes.

The following statistics are available based on the VMware performance counter information:

- Datastore
- Disk device
- CPU

- Power
- Network interface
- Custom performance counter items

For the complete list of items that obtain data from VMware performance counters, see *VMware monitoring item keys*.

Note that the frequency of VMware event retrieval depends on the polling interval of vmware.eventlog, but cannot be less then 5 seconds.

Configuration If Zabbix server is compiled from sources, it must be compiled with the `--with-libcurl --with-libxml2` configuration options to enable virtual machine monitoring. Zabbix packages are compiled with these options already enabled.

The following Zabbix server configuration file parameters can be modified for virtual machine monitoring:

- StartVMwareCollectors

Note:
It is recommended to enable more collectors than the number of monitored VMware services; otherwise, the retrieval of VMware performance counter statistics might be delayed by the retrieval of VMware configuration data (which takes a while for large installations).

 Generally, the value of StartVMwareCollectors should not dip below 2 and should not exceed twice the amount of monitored VMware services: Amount of services < StartVMwareCollectors < (Amount of services * 2). For example, when monitoring one VMware service, set StartVMwareCollectors to 2; when monitoring three services, set StartVMwareCollectors to 5.

 Note that the required number of collectors also depends on the scope of the VMware environment, and the VMwareFrequency and VMwarePerfFrequency configuration parameters.

- VMwareCacheSize
- VMwareFrequency
- VMwarePerfFrequency
- VMwareTimeout

Attention:
To support datastore capacity metrics, ensure that the value of the VMware `vpxd.stats.maxQueryMetrics` key is set to at least 64. For more information, see the *VMware Knowledge Base* article.

Discovery

Zabbix can use low-level discovery rules (for example, `vmware.hv.discovery[{$VMWARE.URL}]`) to automatically discover VMware hypervisors and virtual machines. Moreover, Zabbix can use host prototypes to automatically generate real hosts for the discovered entities. For more information, see *Host prototypes*.

Configuration examples

For a basic example of how to set up Zabbix for monitoring VMware using the *VMware FQDN* template, see *Monitor VMware with Zabbix*.

For a more detailed example of how to create a host, a low-level discovery rule, and a host prototype for monitoring VMware, see *Setup example*.

Extended logging The data collected by the *vmware collector* processes can be logged for detailed debugging using debug level 5. The debug level can be configured in the server and proxy configuration files or using the runtime control option `-R log_level_increase="vmware collector,N"`, where `N` is the process number.

For example, to increase the debug level from 4 to 5 for all *vmware collector* processes, run the following command:

```
zabbix_server -R log_level_increase="vmware collector"
```

To increase the debug level from 4 to 5 for the second *vmware collector* process, run the following command:

```
zabbix_server -R log_level_increase="vmware collector,2"
```

When extended logging of VMware collector data is no longer required, it is recommended to decrease the debug level to default (3) by running the `-R log_level_decrease` command.

Troubleshooting

- In case of unavailable metrics, please ensure that they are not made unavailable or turned off by default in recent VMware vSphere versions, or if some limits are not placed on performance-metric database queries. For more information, see ZBX-12094.

- If `config.vpxd.stats.maxQueryMetrics` is invalid or exceeds the maximum number of characters permitted error, add a `config.vpxd.stats.maxQueryMetrics` parameter to the vCenter Server settings. The value of this parameter should be the same as the value of `maxQuerysize` in VMware's `web.xml` file. For more information, see *VMware Knowledge Base* article.
- If you suspect your Zabbix installation is using too much memory, see Profiling excessive memory usage with tcmalloc.

VMware monitoring item keys

Overview This page provides details on the simple checks that can be used to monitor VMware environments. The metrics are grouped by the monitoring target.

Supported item keys The item keys are listed without parameters and additional information. Click the item key to see the full details.

Item key	Description	Item group
vmware.eventlog	The VMware event log.	General service
vmware.fullname	The VMware service full name.	
vmware.version	The VMware service version.	
vmware.cl.perfcounter	The VMware cluster performance counter metrics.	Cluster
vmware.cluster.alarms.get	The VMware cluster alarms data.	
vmware.cluster.discovery	The discovery of VMware clusters.	
vmware.cluster.property	The VMware cluster property.	
vmware.cluster.status	The VMware cluster status.	
vmware.cluster.tags.get	The VMware cluster tags array.	
vmware.datastore.alarms.get	The VMware datastore alarms data.	Datastore
vmware.datastore.discovery	The discovery of VMware datastores.	
vmware.datastore.hv.list	The list of datastore hypervisors.	
vmware.datastore.perfcounter	The VMware datastore performance counter value.	
vmware.datastore.property	The VMware datastore property.	
vmware.datastore.read	The amount of time for a read operation from the datastore.	
vmware.datastore.size	The VMware datastore space in bytes or in percentage from total.	
vmware.datastore.tags.get	The VMware datastore tags array.	
vmware.datastore.write	The amount of time for a write operation to the datastore.	
vmware.dc.alarms.get	The VMware datacenter alarms data.	Datacenter
vmware.dc.discovery	The discovery of VMware datacenters.	
vmware.dc.tags.get	The VMware datacenter tags array.	
vmware.dvswitch.discovery	The discovery of VMware vSphere Distributed Switches.	vSphere Distributed Switch
vmware.dvswitch.fetchports.get	The VMware vSphere Distributed Switch ports data.	
vmware.hv.alarms.get	The VMware hypervisor alarms data.	Hypervisor
vmware.hv.cluster.name	The VMware hypervisor cluster name.	
vmware.hv.connectionstate	The VMware hypervisor connection state.	
vmware.hv.cpu.usage	The VMware hypervisor processor usage (Hz).	
vmware.hv.cpu.usage.perf	The VMware hypervisor processor usage as a percentage during the interval.	
vmware.hv.cpu.utilization	The VMware hypervisor processor usage as a percentage during the interval, depends on power management or HT.	
vmware.hv.datacenter.name	The VMware hypervisor datacenter name.	
vmware.hv.datastore.discovery	The discovery of VMware hypervisor datastores.	
vmware.hv.datastore.list	The list of VMware hypervisor datastores.	
vmware.hv.datastore.multipath	The number of available datastore paths.	
vmware.hv.datastore.read	The average amount of time for a read operation from the datastore.	
vmware.hv.datastore.size	The VMware datastore space in bytes or in percentage from total.	
vmware.hv.datastore.write	The average amount of time for a write operation to the datastore.	
vmware.hv.discovery	The discovery of VMware hypervisors.	
vmware.hv.diskinfo.get	The VMware hypervisor disk data.	
vmware.hv.fullname	The complete product name, including version information.	
vmware.hv.hw.cpu.freq	The VMware hypervisor processor frequency.	
vmware.hv.hw.cpu.model	The VMware hypervisor processor model.	

Item key	Description	Item group
vmware.hv.hw.cpu.num	The number of processor cores on VMware hypervisor.	
vmware.hv.hw.cpu.threads	The number of processor threads on VMware hypervisor.	
vmware.hv.hw.memory	The VMware hypervisor total memory size.	
vmware.hv.hw.model	The VMware hypervisor model.	
vmware.hv.hw.sensors.get	The VMware hypervisor hardware sensors value.	
vmware.hv.hw.serialnumber	The VMware hypervisor serial number.	
vmware.hv.hw.uuid	The VMware hypervisor BIOS UUID.	
vmware.hv.hw.vendor	The VMware hypervisor vendor name.	
vmware.hv.maintenance	The VMware hypervisor maintenance status.	
vmware.hv.memory.size.ballooned	The VMware hypervisor ballooned memory size.	
vmware.hv.memory.used	The VMware hypervisor used memory size.	
vmware.hv.net.if.discovery	The discovery of VMware hypervisor network interfaces.	
vmware.hv.network.in	The VMware hypervisor network input statistics.	
vmware.hv.network.linkspeed	The VMware hypervisor network interface speed.	
vmware.hv.network.out	The VMware hypervisor network output statistics.	
vmware.hv.perfcounter	The VMware hypervisor performance counter value.	
vmware.hv.property	The VMware hypervisor property.	
vmware.hv.power	The VMware hypervisor power usage.	
vmware.hv.sensor.health.state	The VMware hypervisor health state rollup sensor.	
vmware.hv.sensors.get	The VMware hypervisor HW vendor state sensors.	
vmware.hv.status	The VMware hypervisor status.	
vmware.hv.tags.get	The VMware hypervisor tags array.	
vmware.hv.uptime	The VMware hypervisor uptime.	
vmware.hv.version	The VMware hypervisor version.	
vmware.hv.vm.num	The number of virtual machines on the VMware hypervisor.	
vmware.rp.cpu.usage	The CPU usage in hertz during the interval on VMware Resource Pool.	Resource pool
vmware.rp.memory	The memory metrics of VMware resource pool.	
vmware.alarms.get	The VMware virtual center alarms data.	Virtual center
vmware.vm.alarms.get	The VMware virtual machine alarms data.	Virtual machine
vmware.vm.attribute	The VMware virtual machine custom attribute value.	
vmware.vm.cluster.name	The VMware virtual machine name.	
vmware.vm.consolidationneeded	The VMware virtual machine disk requires consolidation.	
vmware.vm.cpu.latency	The percentage of time the virtual machine is unable to run because it is contending for access to the physical CPU(s).	
vmware.vm.cpu.num	The number of processors on VMware virtual machine.	
vmware.vm.cpu.readiness	The percentage of time that the virtual machine was ready, but could not get scheduled to run on the physical CPU.	
vmware.vm.cpu.ready	The time that the virtual machine was ready, but could not get scheduled to run on the physical CPU.	
vmware.vm.cpu.swapwait	The percentage of CPU time spent waiting for swap-in.	
vmware.vm.cpu.usage	The VMware virtual machine processor usage (Hz).	
vmware.vm.cpu.usage.perf	The VMware virtual machine processor usage as a percentage during the interval.	
vmware.vm.datacenter.name	The VMware virtual machine datacenter name.	
vmware.vm.discovery	The discovery of VMware virtual machines.	
vmware.vm.guest.memory.swapped	The amount of guest physical memory that is swapped out to the swap space.	
vmware.vm.guest.osuptime	The total time elapsed since the last operating system boot-up.	
vmware.vm.hv.maintenance	The VMware virtual machine hypervisor maintenance status.	
vmware.vm.hv.name	The VMware virtual machine hypervisor name.	
vmware.vm.memory.size	The VMware virtual machine total memory size.	
vmware.vm.memory.size.ballooned	The VMware virtual machine ballooned memory size.	
vmware.vm.memory.size.compressed	The VMware virtual machine compressed memory size.	
vmware.vm.memory.size.consumed	The amount of host physical memory consumed for backing up guest physical memory pages.	
vmware.vm.memory.size.private	The VMware virtual machine private memory size.	
vmware.vm.memory.size.shared	The VMware virtual machine shared memory size.	
vmware.vm.memory.size.swapped	The VMware virtual machine swapped memory size.	
vmware.vm.memory.size.usage.guest	The VMware virtual machine guest memory usage.	
vmware.vm.memory.size.usage.host	The VMware virtual machine host memory usage.	

Item key	Description	Item group
vmware.vm.memory.usage	The percentage of host physical memory that has been consumed.	
vmware.vm.net.if.discovery	The discovery of VMware virtual machine network interfaces.	
vmware.vm.net.if.in	The VMware virtual machine network interface input statistics.	
vmware.vm.net.if.out	The VMware virtual machine network interface output statistics.	
vmware.vm.net.if.usage	The VMware virtual machine network utilization during the interval.	
vmware.vm.perfcounter	The VMware virtual machine performance counter value.	
vmware.vm.powerstate	The VMware virtual machine power state.	
vmware.vm.property	The VMware virtual machine property.	
vmware.vm.snapshot.get	The VMware virtual machine snapshot state.	
vmware.vm.state	The VMware virtual machine state.	
vmware.vm.storage.committed	The VMware virtual machine committed storage space.	
vmware.vm.storage.readoio	The average number of outstanding read requests to the virtual disk during the collection interval.	
vmware.vm.storage.totalreadlatency	The average time a read from the virtual disk takes.	
vmware.vm.storage.totalwritelatency	The average time a write to the virtual disk takes.	
vmware.vm.storage.uncommitted	The VMware virtual machine uncommitted storage space.	
vmware.vm.storage.unshared	The VMware virtual machine unshared storage space.	
vmware.vm.storage.writeoio	The average number of outstanding write requests to the virtual disk during the collection interval.	
vmware.vm.tags.get	The VMware virtual machine tags array.	
vmware.vm.tools	The VMware virtual machine guest tools state or version.	
vmware.vm.uptime	The VMware virtual machine uptime.	
vmware.vm.vfs.dev.discovery	The discovery of VMware virtual machine disk devices.	
vmware.vm.vfs.dev.read	The VMware virtual machine disk device read statistics.	
vmware.vm.vfs.dev.write	The VMware virtual machine disk device write statistics.	
vmware.vm.vfs.fs.discovery	The discovery of VMware virtual machine file systems.	
vmware.vm.vfs.fs.size	The VMware virtual machine file system statistics.	

Item key details Parameters without angle brackets are mandatory. Parameters marked with angle brackets **< >** are optional.

vmware.eventlog[url,<mode>,<severity>]

 The VMware event log.
 Return value: *Log*.

Parameters:

- **url** - the VMware service URL;
- **mode** - *all* (default) or *skip* - skip the processing of older data;
- **severity** - filter by severity: *error*, *warning*, *info* or *user*. This parameter must be quoted if more than one severity is specified in a comma-delimited list (e.g.,`"error,warning,info,user"`). Disabled by default.

Comments:

- There must be only one `vmware.eventlog` item key per URL;
- See also example of filtering VMware event log records.

vmware.fullnameurl

 The VMware service full name.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL.

vmware.versionurl

 The VMware service version.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL.

vmware.cl.perfcounter[url,id,path,<instance>]

 The VMware cluster performance counter metrics.
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **id** - the VMware cluster ID. `id` can be received from `vmware.cluster.discovery[]` as {#CLUSTER.ID}.

- **path** - the performance counter path[1];
- **instance** - the performance counter instance.

vmware.cluster.alarms.get[url,id]

 The VMware cluster alarms data.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **id** - the VMware cluster ID.

vmware.cluster.discoveryurl

 The discovery of VMware clusters.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL.

vmware.cluster.property[url,id,prop]

 The VMware cluster property.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **id** - the VMware cluster ID;
- **prop** - the property path which is the name of a property of the VM object as defined in the VMware SDK.

Examples:

```
vmware.vm.property[{$VMWARE.URL},{$VMWARE.VM.UUID},overallStatus]
vmware.vm.property[{$VMWARE.URL},{$VMWARE.VM.UUID},runtime.powerState]
```

vmware.cluster.status[url,name]

 The VMware cluster status.
 Return value: *0* - gray; *1* - green; *2* - yellow; *3* - red.

Parameters:

- **url** - the VMware service URL;
- **name** - the VMware cluster name.

vmware.cluster.tags.get[url,id]

 The VMware cluster tags array.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **id** - the VMware cluster ID.

This item works with vSphere 6.5 and newer.

vmware.datastore.alarms.get[url,uuid]

 The VMware datastore alarms data.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware datastore global unique identifier.

vmware.datastore.discovery[url,<filter_uuid>]

 The discovery of VMware datastores.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **filter_uuid** - a single HV or VM unique identifier (default: empty). If specified, only linked datastores will be discovered.

vmware.datastore.hv.list[url,datastore]

 The list of datastore hypervisors.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;

- **datastore** - the datastore UUID or name.

Output example:

```
esx7-01-host.zabbix.sandbox
esx7-02-host.zabbix.sandbox
```

vmware.datastore.perfcounter[url,uuid,path,<instance>]

 The VMware datastore performance counter value.
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware datastore global unique identifier;
- **path** - the performance counter path[1];
- **instance** - the performance counter instance. Use empty instance for aggregate values (default). `instance` can be received from `vmware.datastore.discovery[]` as part of the {#DATASTORE.EXTENT} array.

vmware.datastore.property[url,uuid,prop]

 The VMware datastore property.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware datastore global unique identifier;
- **prop** - the property path.

vmware.datastore.read[url,datastore,<mode>]

 The amount of time for a read operation from the datastore (milliseconds).
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **datastore** - the datastore UUID or name;
- **mode** - *latency* (average value, default) or *maxlatency* (maximum value).

vmware.datastore.size[url,datastore,<mode>]

 The VMware datastore space in bytes or in percentage from total.
 Return value: *Integer* - for bytes; *Float* - for percentage.

Parameters:

- **url** - the VMware service URL;
- **datastore** - the datastore UUID or name;
- **mode** - possible values: *total* (default), *free*, *pfree* (free percentage), *uncommitted*.

vmware.datastore.tags.get[url,uuid]

 The VMware datastore tags array.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware datastore global unique identifier.

This item works with vSphere 6.5 and newer.

vmware.datastore.write[url,datastore,<mode>]

 The amount of time for a write operation to the datastore (milliseconds).
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **datastore** - the datastore UUID or name;
- **mode** - *latency* (average value, default) or *maxlatency* (maximum value).

vmware.dc.alarms.get[url,id]

 The VMware datacenter alarms data.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **id** - the VMware datacenter ID.

vmware.dc.discoveryurl

 The discovery of VMware datacenters.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL.

vmware.dc.tags.get[url,id]

 The VMware datacenter tags array.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **id** - the VMware datacenter ID.

This item works with vSphere 6.5 and newer.

vmware.dvswitch.discoveryurl

 The discovery of VMware vSphere Distributed Switches.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL.

vmware.dvswitch.fetchports.get[url,uuid,<filter>,<mode>]

 The VMware vSphere Distributed Switch ports data.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware DVSwitch global unique identifier.
- **filter** - a single string with comma-separated criteria for selecting ports;
- **mode** - *state* (all XML without ”config” XML nodes, default) or *full*.

The **filter** parameter supports the criteria available in the VMware data object DistributedVirtualSwitchPortCriteria.

Example:

```
vmware.dvswitch.fetchports.get[{$VMWARE.URL},{$VMWARE.DVS.UUID},"connected:true,active:true,uplinkPort:fal
```

vmware.hv.alarms.get[url,uuid]

 The VMware hypervisor alarms data.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.cluster.name[url,uuid]

 The VMware hypervisor cluster name.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.connectionstate[url,uuid]

 The VMware hypervisor connection state.
 Return value: *String*: *connected*, *disconnected*, or *notResponding*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.cpu.usage[url,uuid]

 The VMware hypervisor processor usage (Hz).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.cpu.usage.perf[url,uuid]

 The VMware hypervisor processor usage as a percentage during the interval.
 Return value: *Float*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.cpu.utilization[url,uuid]

 The VMware hypervisor processor usage as a percentage during the interval, depends on power management or HT.
 Return value: *Float*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.datacenter.name[url,uuid]

 The VMware hypervisor datacenter name.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.datastore.discovery[url,uuid]

 The discovery of VMware hypervisor datastores.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.datastore.list[url,uuid]

 The list of VMware hypervisor datastores.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

Output example:

```
SSD-RAID1-VAULT1
SSD-RAID1-VAULT2
SSD-RAID10
```

vmware.hv.datastore.multipath[url,uuid,<datastore>,<partitionid>]

 The number of available datastore paths.
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier;
- **datastore** - the datastore UUID or name;
- **partitionid** - the internal ID of physical device from `vmware.hv.datastore.discovery`.

vmware.hv.datastore.read[url,uuid,datastore,<mode>]

 The average amount of time for a read operation from the datastore (milliseconds).
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier;
- **datastore** - the datastore UUID or name;
- **mode** - *latency* (default).

vmware.hv.datastore.size[url,uuid,datastore,<mode>]

 The VMware datastore space in bytes or in percentage from total.
 Return value: *Integer* - for bytes; *Float* - for percentage.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier;
- **datastore** - the datastore UUID or name;
- **mode** - possible values: *total* (default), *free*, *pfree* (free percentage), *uncommitted*.

vmware.hv.datastore.write[url,uuid,datastore,<mode>]

 The average amount of time for a write operation to the datastore (milliseconds).
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier;
- **datastore** - the datastore UUID or name;
- **mode** - *latency* (default).

vmware.hv.discoveryurl

 The discovery of VMware hypervisors.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL.

vmware.hv.diskinfo.get[url,uuid]

 The VMware hypervisor disk data.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.fullname[url,uuid]

 The complete product name, including version information.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.hw.cpu.freq[url,uuid]

 The VMware hypervisor processor frequency (Hz).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.hw.cpu.model[url,uuid]

 The VMware hypervisor processor model.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.hw.cpu.num[url,uuid]

 The number of processor cores on VMware hypervisor.
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.hw.cpu.threads[url,uuid]

 The number of processor threads on VMware hypervisor.
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.hw.memory[url,uuid]

 The VMware hypervisor total memory size (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.hw.model[url,uuid]

 The VMware hypervisor model.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.hw.sensors.get[url,uuid]

 The VMware hypervisor hardware sensors value.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.hw.serialnumber[url,uuid]

 The VMware hypervisor serial number.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

This item works with vSphere API 6.7 and newer.

vmware.hv.hw.uuid[url,uuid]

 The VMware hypervisor BIOS UUID.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

This item works with vSphere API 6.7 and newer.

vmware.hv.hw.vendor[url,uuid]

 The VMware hypervisor vendor name.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

This item works with vSphere API 6.7 and newer.

vmware.hv.maintenance[url,uuid]

 The VMware hypervisor maintenance status.
 Return value: *0* - not in maintenance; *1* - in maintenance.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.memory.size.ballooned[url,uuid]

 The VMware hypervisor ballooned memory size (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.memory.used[url,uuid]

 The VMware hypervisor used memory size (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.net.if.discovery[url,uuid]

 The discovery of VMware hypervisor network interfaces.
 Return value: *JSON object*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.network.in[url,uuid,<mode>]

 The VMware hypervisor network input statistics (bytes per second).
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier;
- **mode** - *bps* (default), *packets*, *dropped*, *errors*, *broadcast*.

vmware.hv.network.linkspeed[url,uuid,ifname]

 The VMware hypervisor network interface speed.
 Return value: *Integer*. Returns *0*, if the network interface is down; otherwise, the speed value of the interface.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier;
- **ifname** - the interface name.

vmware.hv.network.out[url,uuid,<mode>]

 The VMware hypervisor network output statistics (bytes per second).
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier;
- **mode** - *bps* (default), *packets*, *dropped*, *errors*, *broadcast*.

vmware.hv.perfcounter[url,uuid,path,<instance>]

 The VMware hypervisor performance counter value.
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier;
- **path** - the performance counter path [1];
- **instance** - the performance counter instance. Use empty instance for aggregate values (default).

vmware.hv.property[url,uuid,prop]

 The VMware hypervisor property.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier;
- **prop** - the property path.

vmware.hv.power[url,uuid,<max>]

 The VMware hypervisor power usage (W).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier;

- **max** - the maximum allowed power usage.

vmware.hv.sensor.health.state[url,uuid]

 The VMware hypervisor health state rollup sensor.
 Return value: *Integer*: *0* - gray; *1* - green; *2* - yellow; *3* - red.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

Note that the item might not work in VMware vSphere 6.5 and newer, because VMware has deprecated the *VMware Rollup Health State* sensor.

vmware.hv.sensors.get[url,uuid]

 The VMware hypervisor HW vendor state sensors.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.status[url,uuid]

 The VMware hypervisor status.
 Return value: *Integer*: *0* - gray; *1* - green; *2* - yellow; *3* - red.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

This item uses the host system overall status property.

vmware.hv.tags.get[url,uuid]

 The VMware hypervisor tags array.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

This item works with vSphere 6.5 and newer.

vmware.hv.uptime[url,uuid]

 The VMware hypervisor uptime (seconds).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

This item uses the host system overall status property.

vmware.hv.version[url,uuid]

 The VMware hypervisor version.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.hv.vm.num[url,uuid]

 The number of virtual machines on the VMware hypervisor.
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware hypervisor global unique identifier.

vmware.rp.cpu.usage[url,rpid]

 The CPU usage in hertz during the interval on VMware Resource Pool.
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;

- **rpid** - the VMware resource pool ID.

vmware.rp.memory[url,rpid,<mode>]

 The memory metrics of VMware resource pool.
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **rpid** - the VMware resource pool ID;
- **mode** - possible values:
consumed (default) - the amount of host physical memory consumed for backing up guest physical memory pages
 ballooned - the amount of guest physical memory reclaimed from the virtual machine by the balloon driver in the guest
overhead - the host physical memory consumed by ESXi data structures for running the virtual machines

vmware.alarms.geturl

 The VMware virtual center alarms data.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL.

vmware.vm.alarms.get[url,uuid]

 The VMware virtual machine alarms data.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.attribute[url,uuid,name]

 The VMware virtual machine custom attribute value.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **name** - the custom attribute name.

vmware.vm.cluster.name[url,uuid]

 The VMware virtual machine name.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.consolidationneeded[url,uuid]

 The VMware virtual machine disk requires consolidation.
 Return value: *String*: *true* - consolidation is needed; *false* - consolidation is not needed.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.cpu.latency[url,uuid]

 The percentage of time the virtual machine is unable to run because it is contending for access to the physical CPU(s).
 Return value: *Float*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.cpu.num[url,uuid]

 The number of processors on VMware virtual machine.
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.cpu.readiness[url,uuid,<instance>]

 The percentage of time that the virtual machine was ready, but could not get scheduled to run on the physical CPU.
 Return value: *Float*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **instance** - the CPU instance.

vmware.vm.cpu.ready[url,uuid]

 The time (in milliseconds) that the virtual machine was ready, but could not get scheduled to run on the physical CPU. CPU ready time is dependent on the number of virtual machines on the host and their CPU loads (%).
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.cpu.swapwait[url,uuid,<instance>]

 The percentage of CPU time spent waiting for swap-in.
 Return value: *Float*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **instance** - the CPU instance.

vmware.vm.cpu.usage[url,uuid]

 The VMware virtual machine processor usage (Hz).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.cpu.usage.perf[url,uuid]

 The VMware virtual machine processor usage as a percentage during the interval.
 Return value: *Float*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.datacenter.name[url,uuid]

 The VMware virtual machine datacenter name.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.discoveryurl

 The discovery of VMware virtual machines.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL.

vmware.vm.guest.memory.size.swapped[url,uuid]

 The amount of guest physical memory that is swapped out to the swap space (KB).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.guest.osuptime[url,uuid]

 The total time elapsed since the last operating system boot-up (in seconds).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.hv.maintenance[url,uuid]

 The VMware virtual machine hypervisor maintenance status.
 Return value: *0* - not in maintenance; *1* - in maintenance.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.hv.name[url,uuid]

 The VMware virtual machine hypervisor name.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.memory.size[url,uuid]

 The VMware virtual machine total memory size (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.memory.size.ballooned[url,uuid]

 The VMware virtual machine ballooned memory size (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.memory.size.compressed[url,uuid]

 The VMware virtual machine compressed memory size (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.memory.size.consumed[url,uuid]

 The amount of host physical memory consumed for backing up guest physical memory pages (KB).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.memory.size.private[url,uuid]

 The VMware virtual machine private memory size (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.memory.size.shared[url,uuid]

 The VMware virtual machine shared memory size (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.memory.size.swapped[url,uuid]

 The VMware virtual machine swapped memory size (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.memory.size.usage.guest[url,uuid]

 The VMware virtual machine guest memory usage (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.memory.size.usage.host[url,uuid]

 The VMware virtual machine host memory usage (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.memory.usage[url,uuid]

 The percentage of host physical memory that has been consumed.
 Return value: *Float*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.net.if.discovery[url,uuid]

 The discovery of VMware virtual machine network interfaces.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.net.if.in[url,uuid,instance,<mode>]

 The VMware virtual machine network interface input statistics (bytes/packets per second).
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **instance** - the network interface instance;
- **mode** - *bps* (default) or *pps* - bytes or packets per second.

vmware.vm.net.if.out[url,uuid,instance,<mode>]

 The VMware virtual machine network interface output statistics (bytes/packets per second).
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **instance** - the network interface instance;
- **mode** - *bps* (default) or *pps* - bytes or packets per second.

vmware.vm.net.if.usage[url,uuid,<instance>]

 The VMware virtual machine network utilization (combined transmit-rates and receive-rates) during the interval (KBps).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **instance** - the network interface instance.

vmware.vm.perfcounter[url,uuid,path,<instance>]

 The VMware virtual machine performance counter value.
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **path** - the performance counter path [1];
- **instance** - the performance counter instance. Use empty instance for aggregate values (default).

vmware.vm.powerstate[url,uuid]

 The VMware virtual machine power state.
 Return value: *0* - poweredOff; *1* - poweredOn; *2* - suspended.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.property[url,uuid,prop]

 The VMware virtual machine property.
 Return value: *String*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **prop** - the property path.

vmware.vm.snapshot.get[url,uuid]

 The VMware virtual machine snapshot state.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.state[url,uuid]

 The VMware virtual machine state.
 Return value: *String*: *notRunning*, *resetting*, *running*, *shuttingDown*, *standby*, or *unknown*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.storage.committed[url,uuid]

 The VMware virtual machine committed storage space (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.storage.readoio[url,uuid,instance]

 The average number of outstanding read requests to the virtual disk during the collection interval.
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **instance** - the disk device instance.

vmware.vm.storage.totalreadlatency[url,uuid,instance]

 The average time a read from the virtual disk takes (milliseconds).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **instance** - the disk device instance.

vmware.vm.storage.totalwritelatency[url,uuid,instance]

 The average time a write to the virtual disk takes (milliseconds).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;

- **uuid** - the VMware virtual machine global unique identifier;
- **instance** - the disk device instance.

vmware.vm.storage.uncommitted[url,uuid]

 The VMware virtual machine uncommitted storage space (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.storage.unshared[url,uuid]

 The VMware virtual machine unshared storage space (bytes).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.storage.writeoio[url,uuid,instance]

 The average number of outstanding write requests to the virtual disk during the collection interval.
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **instance** - the disk device instance.

vmware.vm.tags.get[url,uuid]

 The VMware virtual machine tags array.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

This item works with vSphere 6.5 and newer.

vmware.vm.tools[url,uuid,<mode>]

 The VMware virtual machine guest tools state.
 Return value: *String*.
 In *status* mode: *guestToolsExecutingScripts* - VMware Tools is starting; *guestToolsNotRunning* - VMware Tools is not running; *guestToolsRunning* - VMware Tools is running.
 In *version* mode: version.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **mode** - *version*, *status* (default).

vmware.vm.uptime[url,uuid]

 The VMware virtual machine uptime (seconds).
 Return value: *Integer*.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.vfs.dev.discovery[url,uuid]

 The discovery of VMware virtual machine disk devices.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

vmware.vm.vfs.dev.read[url,uuid,instance,<mode>]

 The VMware virtual machine disk device read statistics (bytes/operations per second).
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;

- **uuid** - the VMware virtual machine global unique identifier;
- **instance** - the disk device instance;
- **mode** - *bps* (default) or *ops* - bytes or operations per second.

vmware.vm.vfs.dev.write[url,uuid,instance,<mode>]

 The VMware virtual machine disk device write statistics (bytes/operations per second).
 Return value: *Integer* [2].

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **instance** - the disk device instance;
- **mode** - *bps* (default) or *ops* - bytes or operations per second.

vmware.vm.vfs.fs.discovery[url,uuid]

 The discovery of VMware virtual machine file systems.
 Return value: JSON object.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier.

VMware Tools must be installed on the guest virtual machine for this item to work.

vmware.vm.vfs.fs.size[url,uuid,fsname,<mode>]

 The VMware virtual machine file system statistics (bytes/percentages).
 Return value: *Integer* - for bytes; *Float* - for percentage.

Parameters:

- **url** - the VMware service URL;
- **uuid** - the VMware virtual machine global unique identifier;
- **fsname** - the file system name;
- **mode** - *total*, *free*, *used*, *pfree*, or *pused*.

VMware Tools must be installed on the guest virtual machine for this item to work.

Footnotes

[1] See Creating custom performance counter names for VMware.

[2] The value of these items is obtained from VMware performance counters and the VMwarePerfFrequency parameter is used to refresh their data in Zabbix VMware cache:

- vmware.cl.perfcounter
- vmware.hv.datastore.read
- vmware.hv.datastore.write
- vmware.hv.network.in
- vmware.hv.network.out
- vmware.hv.perfcounter
- vmware.vm.cpu.ready
- vmware.vm.net.if.in
- vmware.vm.net.if.out
- vmware.vm.perfcounter
- vmware.vm.vfs.dev.read
- vmware.vm.vfs.dev.write

More info

See Virtual machine monitoring for detailed information how to configure Zabbix to monitor VMware environments.

Virtual machine discovery key fields

The following table lists fields returned by virtual machine related discovery keys.

Item key		
Description	**Field**	**Retrieved content**

Item key		
vmware.cluster.discovery		
Performs cluster discovery.	{#CLUSTER.ID}	Cluster identifier.
	{#CLUSTER.NAME}	Cluster name.

	"resource_pool"	An array containing resource pool data, including resource group ID, tags array, resource pool path, number of virtual machines.
		
		Array structure:
		`[{`
		`"rpid":"resource group id",`
		`"tags":[{}],`
		`"rpath":"resource group path",`
		`"vm_count":0`
		`}]`
		
		For `"tags"` array structure, see the "tags" field.

	"tags"	An array containing tags with tag name, description and category.
		
		Array structure:
		`[{`
		`"tag":"tag name",`
		`"tag_description":"tag description",`
		`"category":"tag category"`
		`}]`

vmware.datastore.discovery		
Performs datastore discovery.	{#DATASTORE}	Datastore name.

	{#DATASTORE.EXTENT}	An array containing datastore extent partition ID and instance name.
		
		Array structure:
		`[{`
		`"partitionid":1,`
		`"instance":"name"`
		`}]`

	{#DATASTORE.TYPE}	Datastore type. Value examples: `VMFS`, `NFS`, `vsan`, etc.
	{#DATASTORE.UUID}	Datastore identifier.

	"tags"	An array containing tags with tag name, description and category.
		
		Array structure:
		`[{`
		`"tag":"tag name",`
		`"tag_description":"tag description",`
		`"category":"tag category"`
		`}]`

vmware.dc.discovery		
Performs datacenter discovery.	{#DATACENTER}	Datacenter name.
	{#DATACENTERID}	Datacenter identifier.

	"tags"	An array containing tags with tag name, description and category.
		
		Array structure:
		`[{`
		`"tag":"tag name",`
		`"tag_description":"tag description",`
		`"category":"tag category"`
		`}]`

vmware.dvswitch.discovery		
Performs vSphere distributed switches discovery.	{#DVS.NAME}	Switch name.

Item key		
	{#DVS.UUID}	Switch identifier.
vmware.hv.discovery		
Performs hypervisor discovery.	{#HV.UUID}	Unique hypervisor identifier.
	{#HV.ID}	Hypervisor identifier (HostSystem managed object name).
	{#HV.NAME}	Hypervisor name.
	{#HV.NETNAME}	Hypervisor network host name.
	{#HV.IP}	Hypervisor IP address, might be empty. In case of an HA configuration with multiple net interfaces, the following selection priority for interface is observed: - prefer the IP which shares the IP-subnet with the vCenter IP; - prefer the IP from IP-subnet with default gateway; - prefer the IP from interface with the lowest ID.
	{#CLUSTER.NAME}	Cluster name, might be empty.
	{#DATACENTER.NAME}	Datacenter name.
	{#PARENT.NAME}	Name of container that stores the hypervisor.
	{#PARENT.TYPE}	Type of container in which the hypervisor is stored. The values could be `Datacenter`, `Folder`, `ClusterComputeResource`, `VMware`, where ”VMware” stands for unknown container type.

	”resource_pool”	An array containing resource pool data, including resource group ID, tags array, resource pool path, number of virtual machines.
		
		Array structure:
		`[{`
		`"rpid":"resource group id",`
		`"tags":[{}],`
		`"rpath":"resource group path",`
		`"vm_count":0`
		`}]`
		
		For `"tags"` array structure, see the ”tags” field.

	”tags”	An array containing tags with tag name, description and category.
		
		Array structure:
		`[{`
		`"tag":"tag name",`
		`"tag_description":"tag description",`
		`"category":"tag category"`
		`}]`

vmware.hv.datastore.discovery		
Performs hypervisor datastore discovery. Note that multiple hypervisors can use the same datastore.	{#DATASTORE}	Datastore name.
	{#DATASTORE.TYPE}	Datastore type. Value examples: `VMFS`, `NFS`, `vsan`, etc.
	{#DATASTORE.UUID}	Datastore identifier.
	{#MULTIPATH.COUNT}	Registered number of datastore paths.
	{#MULTIPATH.PARTITION.COUNT}	Number of available disk partitions.

	”datastore_extent”	An array containing datastore extent instance name and partition ID.
		
		Array structure:
		`[{`
		`"partitionid":1,`
		`"instance":"name"`
		`}]`

Item key		

	”tags”	An array containing tags with tag name, description and category.
		
		Array structure:
		`[{`
		`"tag":"tag name",`
		`"tag_description":"tag description",`
		`"category":"tag category"`
		`}]`

vmware.hv.net.if.discovery		
Performs hypervisor network interfaces discovery.	{#IFNAME}	Interface name.
	{#IFDRIVER}	Interface driver.
	{#IFDUPLEX}	Interface duplex settings.
	{#IFSPEED}	Interface speed.
	{#IFMAC}	Interface mac address.
vmware.vm.discovery		
Performs virtual machine discovery.	{#VM.UUID}	Unique virtual machine identifier.
	{#VM.ID}	Virtual machine identifier (VirtualMachine managed object name).
	{#VM.NAME}	Virtual machine name.
	{#HV.NAME}	Hypervisor name.
	{#HV.UUID}	Unique hypervisor identifier.
	{#HV.ID}	Hypervisor identifier (HostSystem managed object name).
	{#CLUSTER.NAME}	Cluster name, might be empty.
	{#DATACENTER.NAME}	Datacenter name.
	{#DATASTORE.NAME}	Datastore name.
	{#DATASTORE.UUID}	Datastore identifier.
	{#VM.IP}	Virtual machine IP address, might be empty.
	{#VM.DNS}	Virtual machine DNS name, might be empty.
	{#VM.GUESTFAMILY}	Guest virtual machine OS family, might be empty.
	{#VM.GUESTFULLNAME}	Full guest virtual machine OS name, might be empty.
	{#VM.FOLDER}	The chain of virtual machine parent folders, can be used as value for nested groups; folder names are combined with ”/”. Might be empty.
	{#VM.TOOLS.STATUS}	VMware virtual machine tools state.
	{#VM.POWERSTATE}	VMware virtual machine power state (`poweredOff`, `poweredOn`, or `suspended`).
	{#VM.RPOOL.ID}	Resource pool identifier.
	{#VM.RPOOL.PATH}	Full resource pool path excluding the ”root” name ”Resources”. Folder names are combined with ”/”.
	{#VM.SNAPSHOT.COUNT}	Number of VM snapshots.

	”tags”	An array containing tags with tag name, description and category.
		
		Array structure:
		`[{`
		`"tag":"tag name",`
		`"tag_description":"tag description",`
		`"category":"tag category"`
		`}]`

	”vm_customattribute”	An array of virtual machine custom attributes (if defined).
		
		Array structure:
		`[{`
		`"name":"custom field name",`
		`"value":"custom field value"`
		`}]`

Item key		
	”net_if”	An array of virtual machine network interfaces. Array structure: [{ "ifname": "interface name", "ifdesc": "interface description", "ifmac": "00:00:00:00:00:00", "ifconnected": true, "iftype": "interface type", "ifbackingdevice": "interface backing device", "ifdvswitch_uuid": "interface switch uuid", "ifdvswitch_portgroup": "interface switch port group", "ifdvswitch_port": "interface switch port", "ifip": ["interface ip addresses"] }] For the description of returned data, see the ”vmware.vm.net.if.discovery” item key.
vmware.vm.net.if.discovery		
Performs virtual machine network interface discovery.	{#IFNAME}	Network interface name.
	{#IFDESC}	Interface description.
	{#IFMAC}	Interface mac address.
	{#IFCONNECTED}	Interface connection status (`false` - disconnected; `true` - connected).
	{#IFTYPE}	Interface type.
	{#IFBACKINGDEVICE}	Name of the backing device.
	{#IFDVSWITCH.UUID}	Unique vSphere Distributed Switch identifier.
	{#IFDVSWITCH.PORTGROUP}	Distributed port group.
	{#IFDVSWITCH.PORT}	vSphere Distributed Switch port.
	”ifip”	An array of interface addresses.
vmware.vm.vfs.dev.discovery		
Performs virtual machine disk device discovery.	{#DISKNAME}	Disk device name.
vmware.vm.vfs.fs.discovery		
Performs virtual machine file system discovery.	{#FSNAME}	File system name.

JSON examples for VMware items

Overview This section provides additional information about JSON objects returned by various VMware items.

vmware.*.alarms.get The items **vmware.alarms.get[]**, **vmware.cluster.alarms.get[]**, **vmware.datastore.alarms.get[]**, **vmware.dc.alarms.get[]**, **vmware.hv.alarms.get[]**, **vmware.vm.alarms.get[]** return JSON objects with the following structure (values are provided as an example):

```
{
    "alarms": [
        {
            "name": "Host connection and power state",
            "system_name": "alarm.HostConnectionStateAlarm",
            "description": "Default alarm to monitor host connection and power state",
            "enabled": true,
            "key": "alarm-1.host-2013",
            "time": "2022-06-27T05:27:38.759976Z",
```

```
            "overall_status": "red",
            "acknowledged": false,
            "entity": {
                "id": "host-2013",
                "uuid": "4201b707-f737-7cf8-881f-49215d1aad2a",
                "type": "HostSystem",
                "name": "esx-riga-01.prod.local"
            }
        },
        {
            "name": "Host memory usage",
            "system_name": "alarm.HostMemoryUsageAlarm",
            "description": "Default alarm to monitor host memory usage",
            "enabled": true,
            "key": "alarm-4.host-1004",
            "time": "2022-05-16T13:32:42.47863Z",
            "overall_status": "yellow",
            "acknowledged": false,
            "entity": {
                "id": "host-1004",
                "uuid": "5315c818-g848-8dg9-992g-50326e2bbe3a",
                "type": "HostSystem",
                "name": "esx-riga-02.prod.local",
            }
        },
        {
            // other alarms
        }
    ]
}
```

vmware.*.tags.get The items **vmware.cluster.tags.get[]**, **vmware.datastore.tags.get[]**, **vmware.dc.tags.get[]**, **vmware.hv.tags.get[]**, **vmware.vm.tags.get[]** return JSON objects with the following structure (values are provided as an example):

```
{
    "tags": [
        {
            "name": "Windows",
        "description": "tag for cat OS type",
        "category": "OS type"
    },
    {
            "name": "SQL Server",
        "description": "tag for cat application name",
        "category": "application name"
        },
    {
            // other tags
        }
    ]
}
```

vmware.hv.diskinfo.get The item **vmware.hv.diskinfo.get[]** returns JSON objects with the following structure (values are provided as an example):

```
[
  {
    "instance": "mpx.vmhba32:C0:T0:L0",
    "hv_uuid": "8002299e-d7b9-8728-d224-76004bbb6100",
    "datastore_uuid": "",
    "operational_state": [
      "ok"
```

```
    ],
    "lun_type": "disk",
    "queue_depth": 1,
    "model": "USB DISK",
    "vendor": "SMI Corp",
    "revision": "1100",
    "serial_number": "CCYYMMDDHHmmSS9S62CK",
    "vsan": {}
  },
  {
    // other instances
  }
]
```

vmware.dvswitch.fetchports.get The item **vmware.dvswitch.fetchports.get[]** returns JSON objects with the following structure (values are provided as an example):

```
{
    "FetchDVPortsResponse":
    {
        "returnval": [
            {
                "key": "0",
                "dvsUuid": "50 36 6a 24 25 c0 10 9e-05 4a f6 ea 4e 3d 09 88",
                "portgroupKey": "dvportgroup-2023",
                "proxyHost":
                {
                    "@type": "HostSystem",
                    "#text": "host-2021"
                },
                "connectee":
                {
                    "connectedEntity":
                    {
                        "@type": "HostSystem",
                        "#text": "host-2021"
                    },
                    "nicKey": "vmnic0",
                    "type": "pnic"
                },
                "conflict": "false",
                "state":
                {
                    "runtimeInfo":
                    {
                        "linkUp": "true",
                        "blocked": "false",
                        "vlanIds":
                        {
                            "start": "0",
                            "end": "4094"
                        },
                        "trunkingMode": "true",
                        "linkPeer": "vmnic0",
                        "macAddress": "00:00:00:00:00:00",
                        "statusDetail": null,
                        "vmDirectPathGen2Active": "false",
                        "vmDirectPathGen2InactiveReasonOther": "portNptIncompatibleConnectee"
                    },
                    "stats":
                    {
                        "packetsInMulticast": "2385470",
                        "packetsOutMulticast": "45",
```

```
                        "bytesInMulticast": "309250248",
                        "bytesOutMulticast": "5890",
                        "packetsInUnicast": "155601537",
                        "packetsOutUnicast": "113008658",
                        "bytesInUnicast": "121609489384",
                        "bytesOutUnicast": "47240279759",
                        "packetsInBroadcast": "1040420",
                        "packetsOutBroadcast": "7051",
                        "bytesInBroadcast": "77339771",
                        "bytesOutBroadcast": "430392",
                        "packetsInDropped": "0",
                        "packetsOutDropped": "0",
                        "packetsInException": "0",
                        "packetsOutException": "0"
                    }
                },
                "connectionCookie": "1702765133",
                "lastStatusChange": "2022-03-25T14:01:11Z",
                "hostLocalPort": "false"
            },
            {
                //other keys
            }
        ]
    }
}
```


vmware.hv.hw.sensors.get The item **vmware.hv.hw.sensors.get[]** returns JSON objects with the following structure (values are provided as an example):


```
{
    "val":
        {
            "@type": "HostHardwareStatusInfo",
            "storageStatusInfo": [
                {
                    "name": "Intel Corporation HD Graphics 630 #2",
                "status":
                        {
                            "label": "Unknown",
                    "summary": "Cannot report on the current status of the physical element",
                    "key": "Unknown"
                }
            },
            {
                    "name": "Intel Corporation 200 Series/Z370 Chipset Family USB 3.0 xHCI Controller #20"
                "status":
                        {
                     "label": "Unknown",
                     "summary": "Cannot report on the current status of the physical element",
                     "key": "Unknown"
                 }
            },
            {
                    // other hv hw sensors
                }
        ]
        }
}
```


vmware.hv.sensors.get The item **vmware.hv.sensors.get[]** returns JSON objects with the following structure (values are provided as an example):

```
{
    "val":
        {
            "@type": "ArrayOfHostNumericSensorInfo", "HostNumericSensorInfo": [
                {
                    "@type": "HostNumericSensorInfo",
                    "name": "System Board 1 PwrMeter Output --- Normal",
                    "healthState":
                        {
                            "label": "Green",
                            "summary": "Sensor is operating under normal conditions",
                    "key": "green"
                },
                    "currentReading": "10500",
                "unitModifier": "-2",
                "baseUnits": "Watts",
                "sensorType": "other"
            },
                {
                    "@type": "HostNumericSensorInfo",
                    "name": "Power Supply 1 PS 1 Output --- Normal",
                    "healthState":
                        {
                            "label": "Green",
                    "summary": "Sensor is operating under normal conditions",
                    "key": "green"
                },
                    "currentReading": "10000",
            "unitModifier": "-2",
            "baseUnits": "Watts",
            "sensorType": "power"
        },
                {
                        // other hv sensors
                }
        ]
        }
}
```


vmware.vm.snapshot.get If any snapshots exist, the item **vmware.snapshot.get[]** returns a JSON object with the following structure (values are provided as an example):


```
{
  "snapshot": [
    {
      "name": "VM Snapshot 4%2f1%2f2022, 9:16:39 AM",
      "description": "Descr 1",
      "createtime": "2022-04-01T06:16:51.761Z",
      "size": 5755795171,
      "uniquesize": 5755795171
    },
    {
      "name": "VM Snapshot 4%2f1%2f2022, 9:18:21 AM",
      "description": "Descr 2",
      "createtime": "2022-04-01T06:18:29.164999Z",
      "size": 118650595,
      "uniquesize": 118650595
    },
    {
      "name": "VM Snapshot 4%2f1%2f2022, 9:37:29 AM",
      "description": "Descr 3",
      "createtime": "2022-04-01T06:37:53.534999Z",
      "size": 62935016,
```

```
      "uniquesize": 62935016
    }
  ],
  "count": 3,
  "latestdate": "2022-04-01T06:37:53.534999Z",
  "latestage": 22729203,
  "oldestdate": "2022-04-01T06:16:51.761Z",
  "oldestage": 22730465,
  "size": 5937380782,
  "uniquesize": 5937380782
}
```


If no snapshot exists, the item **vmware.snapshot.get[]** returns a JSON object with empty values:


```
{
  "snapshot": [],
  "count": 0,
  "latestdate": null,
  "latestage": 0,
  "oldestdate": null,
  "oldestage": 0,
  "size": 0,
  "uniquesize": 0
}
```


VMware monitoring setup example

Overview

The following example describes how to set up Zabbix for monitoring VMware virtual machines. This involves:

- creating a host that represents your VMware environment;
- creating a low-level discovery rule that discovers virtual machines in your VMware environment;
- creating a host prototype, based on which Zabbix will generate real hosts for virtual machines discovered by the low-level discovery rule.

Prerequisites

> **Note:**
> This example does not cover the configuration of VMware. It is assumed that VMware is already configured.

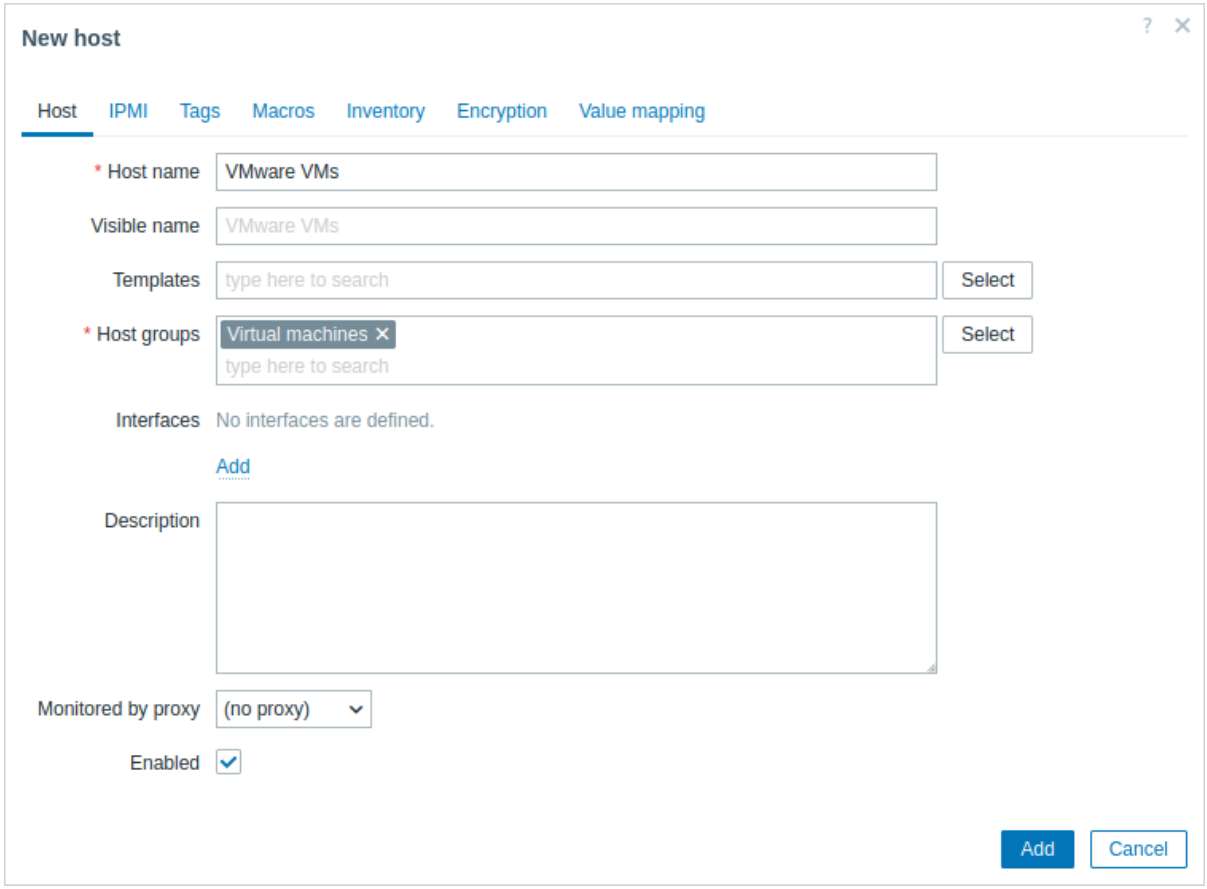
Before proceeding, set the `StartVMwareCollectors` parameter in Zabbix server configuration file to 2 or more (the default value is 0).

Create a host

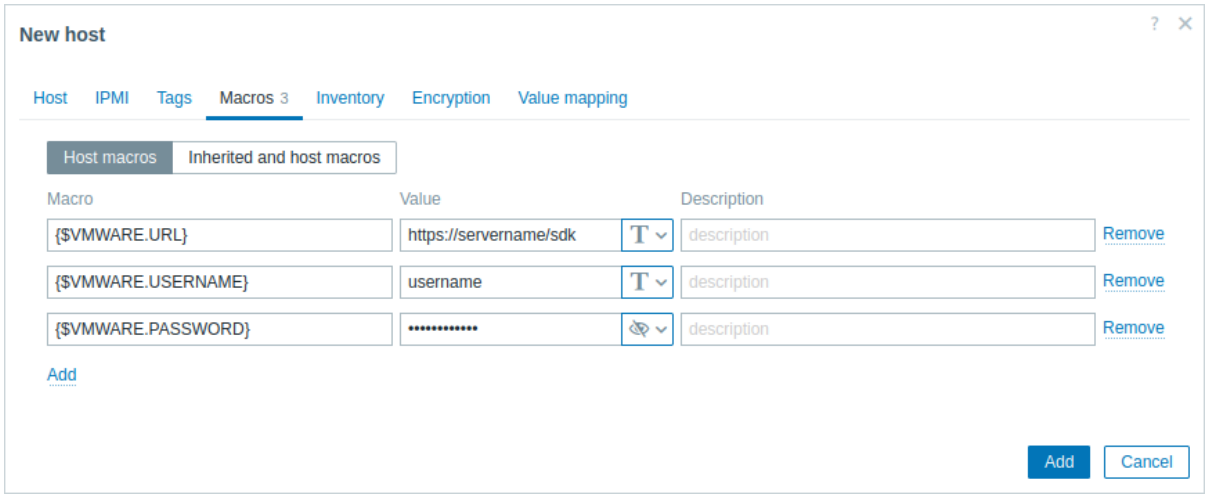
1. Go to *Data collection* > *Hosts*.

2. Create a host:

- In the *Host name* field, enter a host name (for example, "VMware VMs").
- In the *Host groups* field, type or select a host group (for example, "Virtual machines").



- In the *Macros* tab, set the following host macros:
  - {$VMWARE.URL} - VMware service (ESXi hypervisor) SDK URL (https://servername/sdk)
  - {$VMWARE.USERNAME} - VMware service user name
  - {$VMWARE.PASSWORD} - VMware service {$VMWARE.USERNAME} user password



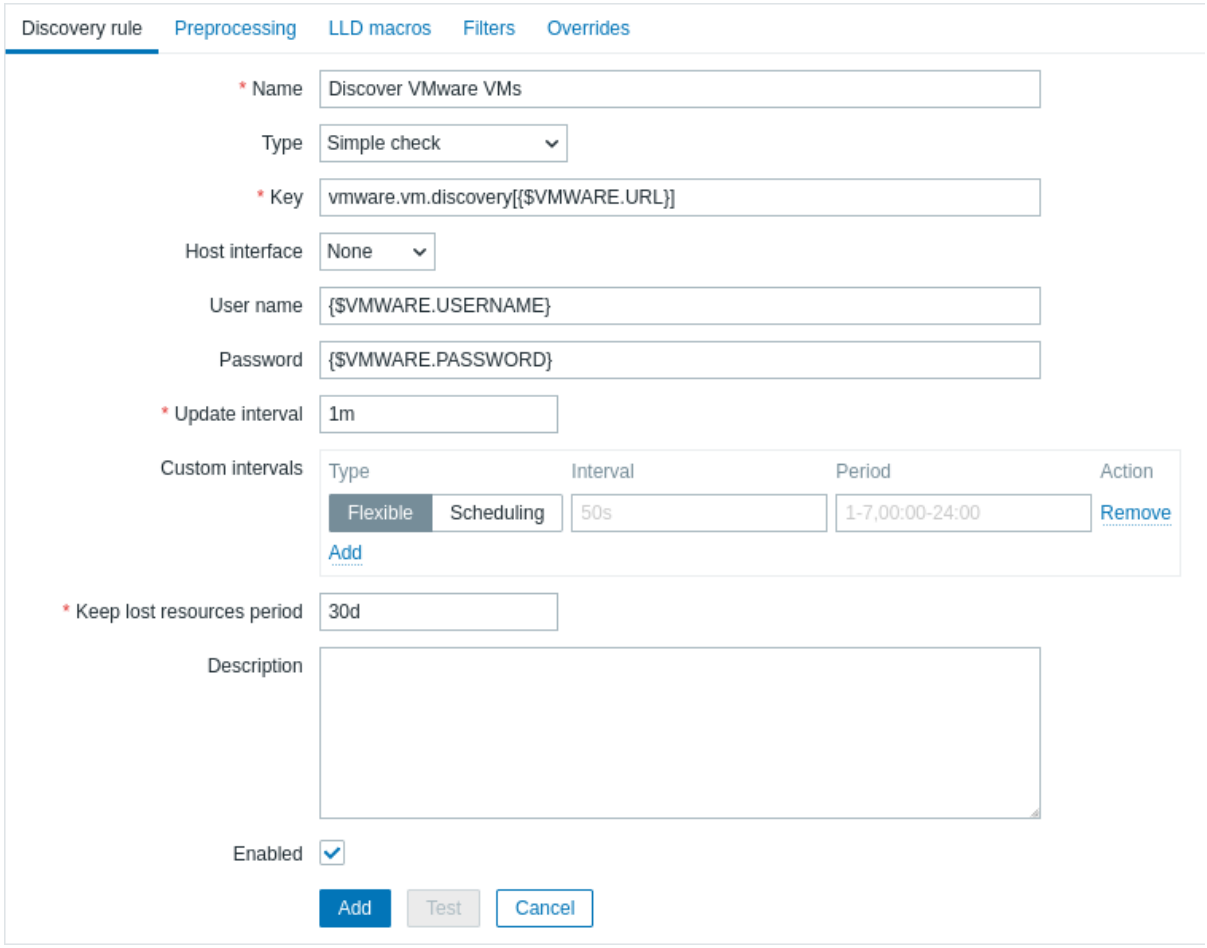
3. Click the *Add* button to create the host. This host will represent your VMware environment.

Create a low-level discovery rule

1. Click *Discovery* for the created host to go to the list of low-level discovery rules for that host.

2. Create a low-level discovery rule:

- In the *Name* field, enter a low-level discovery rule name (for example, "Discover VMware VMs").
- In the *Type* field, select "Simple check".
- In the *Key* field, enter the built-in item key for discovering VMware virtual machines: `vmware.vm.discovery[{$VMWARE.URL}]`
- In the *User name* and *Password* fields, enter the corresponding macros previously configured on the host.



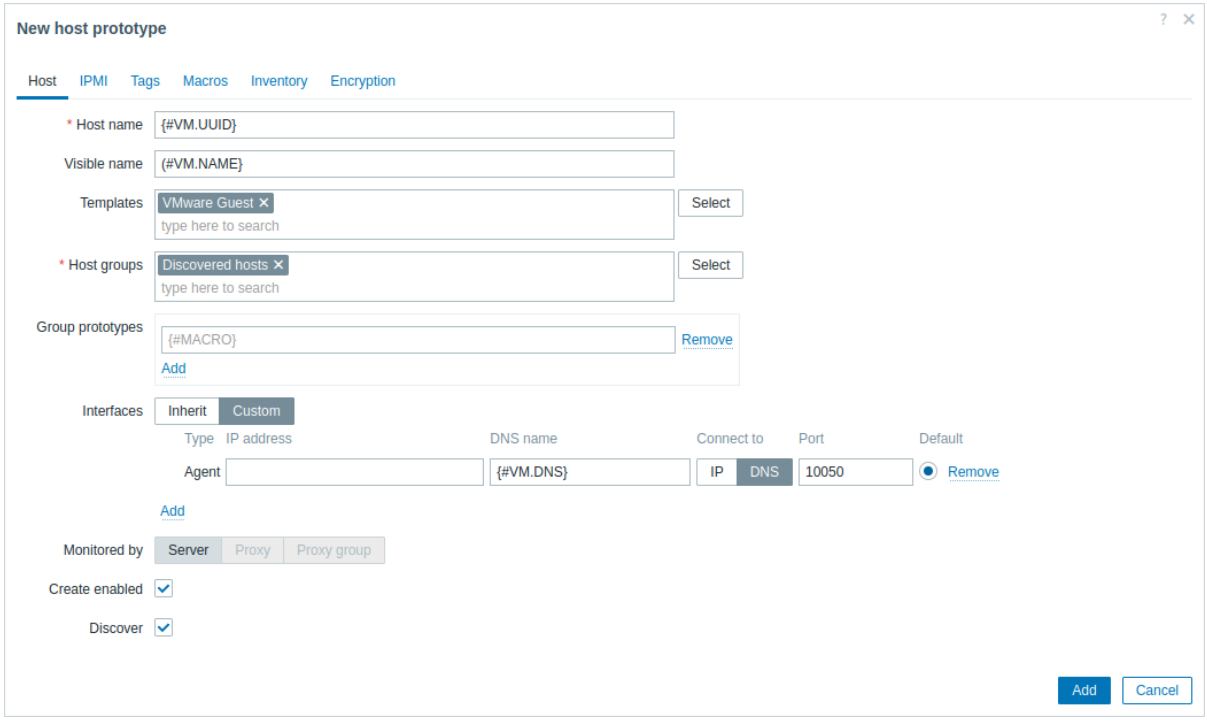
3. Click the *Add* button to create the low-level discovery rule. This discovery rule will discover virtual machines in your VMware environment.

Create a host prototype

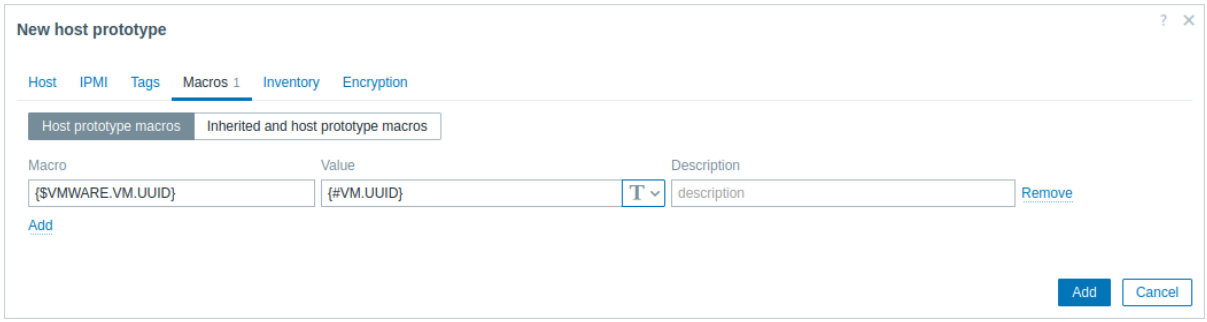
1. In the list of low-level discovery rules, click *Host prototypes* for the previously created low-level discovery rule.

2. Create a host prototype. Since host prototypes are blueprints for creating hosts through low-level discovery rules, most fields will contain low-level discovery macros. This ensures that the hosts are created with properties based on the content retrieved by the previously created low-level discovery rule.

- In the *Host name* field, enter the macro {#VM.UUID}.
- In the *Visible name* field, enter the macro {#VM.NAME}.
- In the *Templates* field, enter or select the ”VMware Guest” template. This template contains VMware items and discovery rules for monitoring the power state of a virtual machine, CPU usage, memory usage, network devices, etc.
- In the *Host groups* field, enter or select a host group (for example, ”Discovered hosts”).
- In the *Interfaces* field, add a custom host interface. Then, enter the macro {#VM.DNS} in the *DNS name* field, or enter the macro {#VM.IP} in the *IP address* field. Alternatively, if your VMware environment virtual machines have multiple interfaces, proceed to the *Advanced host interface configuration* section. Configuring a custom host interface is necessary for the correct functioning of the *VMware Guest* template.



- In the *Macros* tab, set the {$VMWARE.VM.UUID} macro with the value {#VM.UUID}. This is necessary for the correct functioning of the *VMware Guest* template that uses this macro as a host-level user macro in item parameters (for example, *vmware.vm.net.if.discovery[{$VMWARE.URL}, {$VMWARE.VM.UUID}]*).

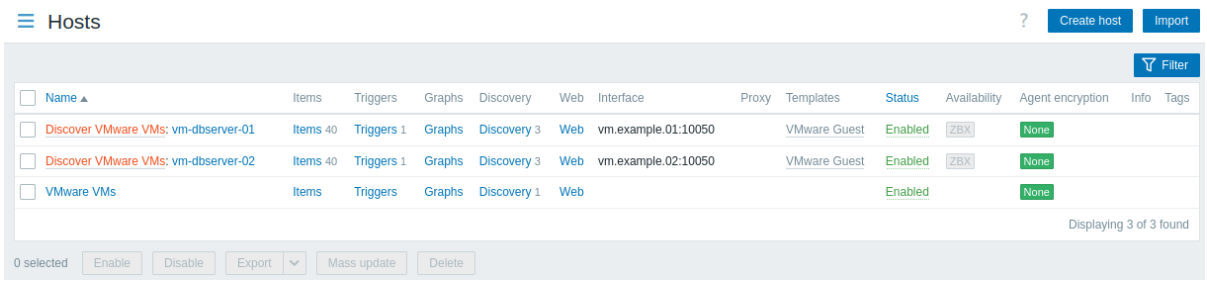


3. Click the *Add* button to create the host prototype. This host prototype will be used to create hosts for virtual machines discovered by the previously created low-level discovery rule.

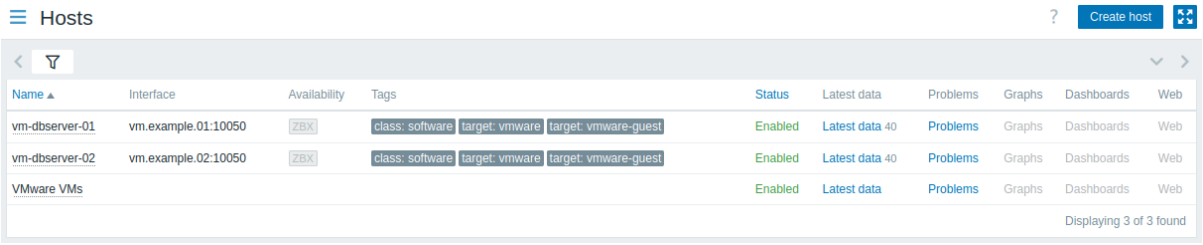
View hosts and metrics

After the host prototype has been created, the low-level discovery rule will create hosts for discovered VMware virtual machines, and Zabbix will start to monitor them. Note that the discovery and creation of hosts can also be executed manually, if necessary.

To view the created hosts, navigate to the *Data collection > Hosts* menu section.



To view collected metrics, navigate to the *Monitoring > Hosts* menu section and click *Latest data* for one of the hosts.



Advanced host interface configuration

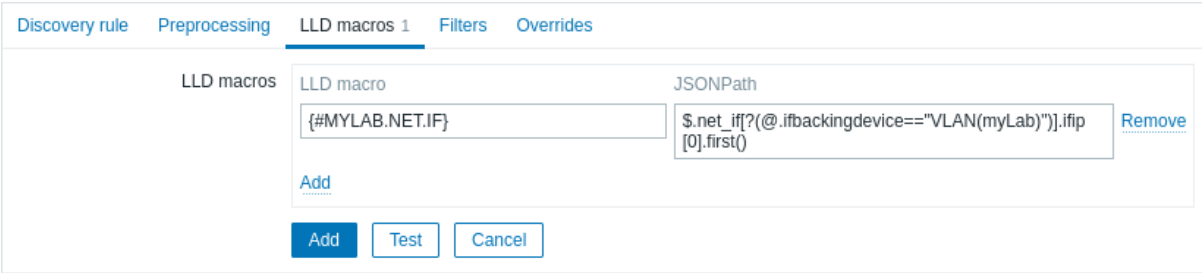
The `vmware.vm.discovery[{$VMWARE.URL}]` item key, configured in the *Create a low-level discovery rule* section, returns network interfaces data in the "net_if" field:

```
"net_if": [
      {
        "ifname": "5000",
        "ifdesc": "Network adapter 1",
        "ifmac": "00:11:22:33:44:55",
        "ifconnected": true,
        "iftype": "VirtualVmxnet3",
        "ifbackingdevice": "VLAN(myLab)",
        "ifdvswitch_uuid": "",
        "ifdvswitch_portgroup": "",
        "ifdvswitch_port": "",
        "ifip": [
          "127.0.0.1",
          "::1"
        ]
      },
      {
        "ifname": "5001",
        "ifdesc": "Network adapter 2",
        "ifmac": "00:11:22:33:44:55",
        "ifconnected": false,
        "iftype": "VirtualVmxnet3",
        "ifbackingdevice": "VLAN(myLab2)",
        "ifdvswitch_uuid": "",
        "ifdvswitch_portgroup": "",
        "ifdvswitch_port": "",
        "ifip": []
      }
    ]
```

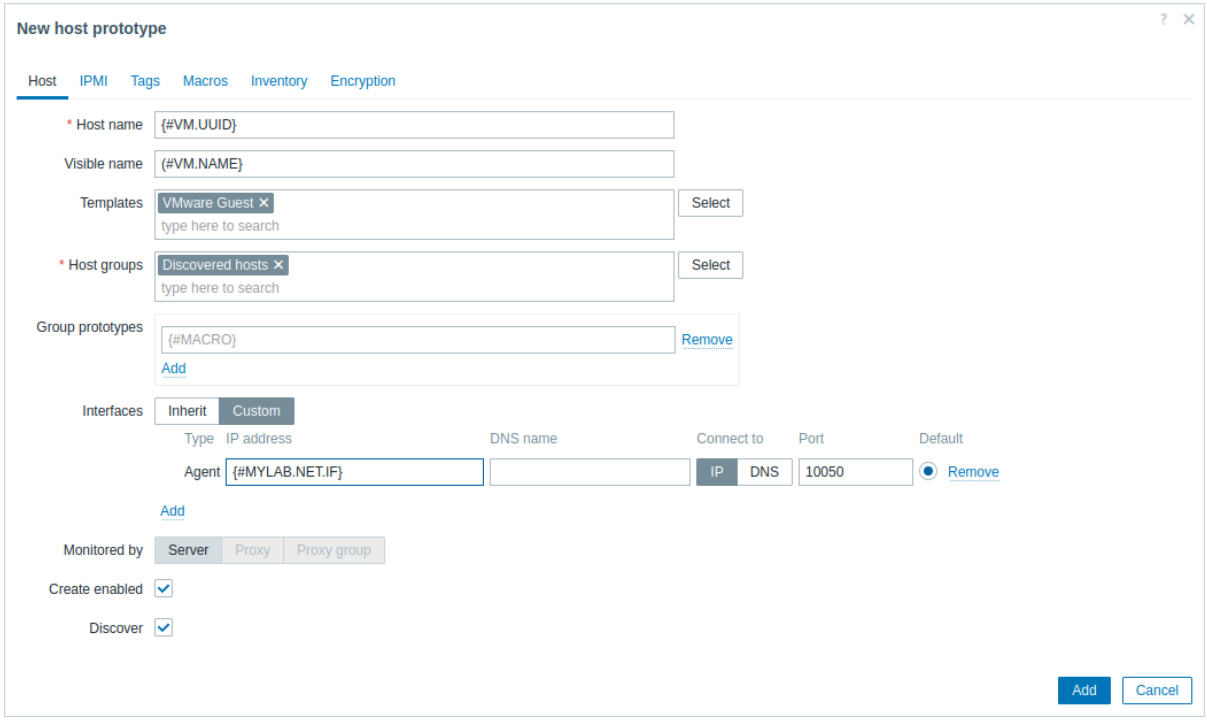
This data can be used to configure a custom host interface.

1. When creating a low-level discovery rule, additionally configure a low-level discovery macro. In the *LLD macros* tab, create a custom LLD macro with a JSONPath value. For example:

- `{#MYLAB.NET.IF}` - `$.net_if[?(@.ifbackingdevice=="VLAN(myLab)")].ifip[0].first()`



2. When creating a host prototype, add a custom host interface and enter the LLD macro in the *DNS name* or *IP address* field.



Web monitoring

Overview With Zabbix you can check several availability aspects of web sites.

Attention:
To perform web monitoring Zabbix server must be initially configured with cURL (libcurl) support.

To activate web monitoring you need to define web scenarios. A web scenario consists of one or several HTTP requests or "steps". The steps are periodically executed by Zabbix server in a pre-defined order. If a host is monitored by proxy, the steps are executed by the proxy.

Web scenarios are attached to hosts/templates in the same way as items, triggers, etc. That means that web scenarios can also be created on a template level and then applied to multiple hosts in one move.

The following information is collected in any web scenario:

- average download speed per second for all steps of whole scenario
- number of the step that failed
- last error message

The following information is collected in any web scenario step:

- download speed per second
- response time
- response code

For more details, see web monitoring items.

Data collected from executing web scenarios is kept in the database. The data is automatically used for graphs, triggers and notifications.

Zabbix can also check if a retrieved HTML page contains a pre-defined string. It can execute a simulated login and follow a path of simulated mouse clicks on the page.

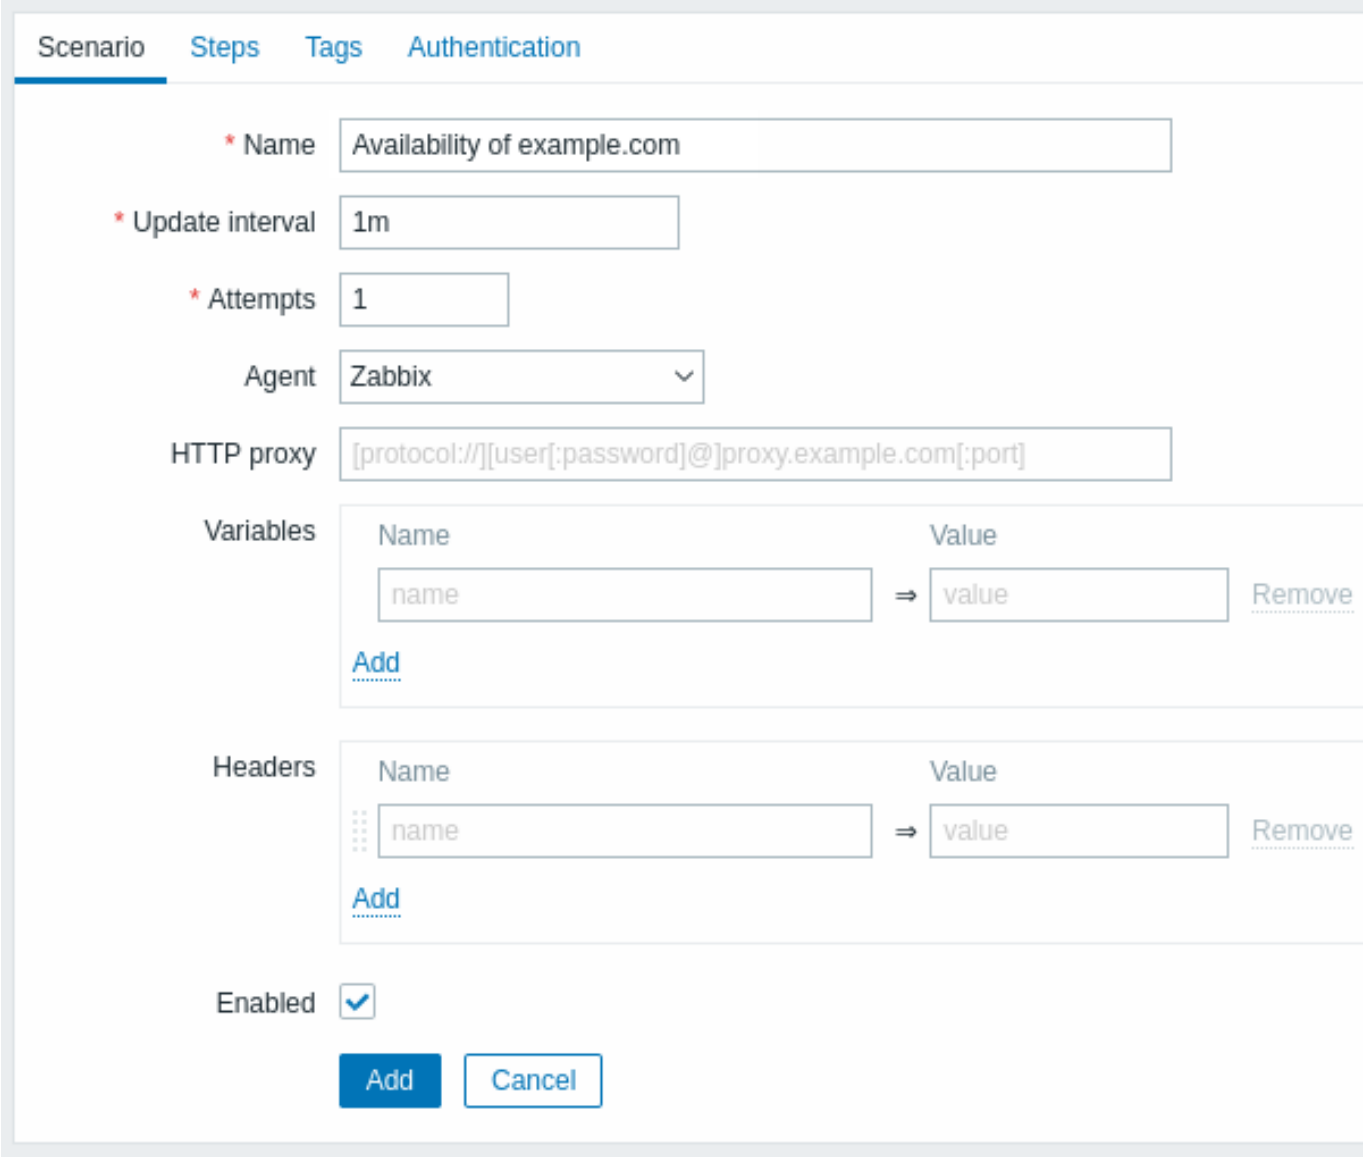
Zabbix web monitoring supports both HTTP and HTTPS. When running a web scenario, Zabbix will optionally follow redirects (see option *Follow redirects* below). Maximum number of redirects is hard-coded to 10 (using cURL option CURLOPT_MAXREDIRS). All cookies are preserved during the execution of a single scenario.

Configuring a web scenario To configure a web scenario:

1. Go to: *Data collection* > *Hosts* (or *Templates*).

2. Click *Web* in the row of the host/template.
3. Click *Create web scenario* to the right (or the scenario name to edit an existing scenario).
4. Enter parameters of the scenario in the form.

The **Scenario** tab allows you to configure the general parameters of a web scenario.



All mandatory input fields are marked with a red asterisk.

Scenario parameters:

Parameter	Description
Name	Unique scenario name. User macros are supported. *Note* that if user macros are used, these macros will be left unresolved in web monitoring item names.
Update interval	How often the scenario will be executed. Time suffixes are supported, e.g., 30s, 1m, 2h, 1d. User macros are supported. *Note* that if a user macro is used and its value is changed (e.g., 5m → 30s), the next check will be executed according to the previous value (farther in the future with the example values). New web scenarios will be checked within 60 seconds of their creation.
Attempts	The number of attempts for executing web scenario steps. In case of network problems (timeout, no connectivity, etc) Zabbix can repeat executing a step several times. The figure set will equally affect each step of the scenario. Up to 10 attempts can be specified, default value is 1. *Note*: Zabbix will not repeat a step because of a wrong response code or the mismatch of a required string.

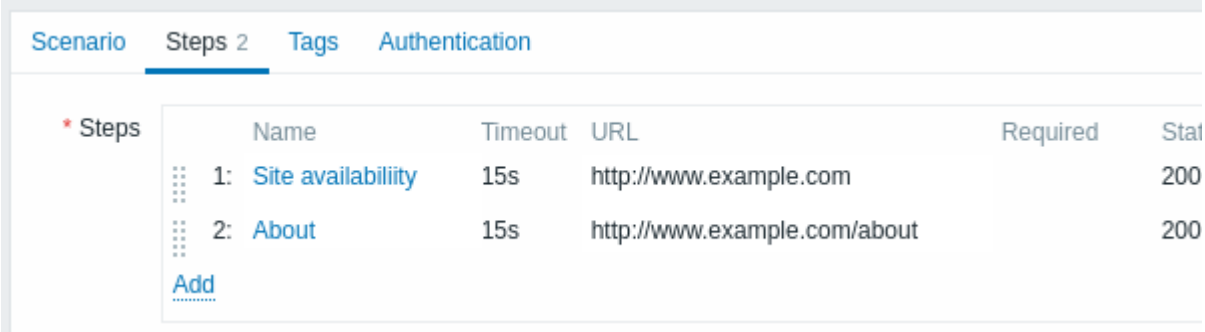
Parameter	Description
Agent	Select a client agent. Zabbix will pretend to be the selected browser. This is useful when a website returns different content for different browsers. User macros can be used in this field.
HTTP proxy	You can specify an HTTP proxy to use, using the format `[protocol://][username[:password]@]proxy.example.com[:port]`. This sets the CURLOPT_PROXY cURL option. The optional `protocol://` prefix may be used to specify alternative proxy protocols (the protocol prefix support was added in cURL 7.21.7). With no protocol specified, the proxy will be treated as an HTTP proxy. By default, 1080 port will be used. If specified, the proxy will overwrite proxy related environment variables like http_proxy, HTTPS_PROXY. If not specified, the proxy will not overwrite proxy-related environment variables. The entered value is passed on "as is", no sanity checking takes place. You may also enter a SOCKS proxy address. If you specify the wrong protocol, the connection will fail and the item will become unsupported. *Note* that only simple authentication is supported with HTTP proxy. User macros can be used in this field.
Variables	Variables that may be used in scenario steps (URL, post variables). They have the following format: **{macro1}**=value1 **{macro2}**=value2 **{macro3}**=regex:<regular expression> **{macro4}**=jsonpath:<jsonpath> **{macro5}**=xmlxpath:<xmlxpath> **{macro6}**={{macro}.function()} (see macro functions) For example: {username}=Alexei {password}=kj3h5kJ34bd {hostid}=regex:hostid is ([0-9]+) {url}=jsonpath:$.host_url {status}=xmlxpath://host/response/status {newvar}={{myvar}.btoa()} The macros can then be referenced in the steps as {username}, {password}, {hostid}, etc. Zabbix will automatically replace them with actual values. Note that variables with `regex:` need one step to get the value of the regular expression so the extracted value can only be applied to the step after. If the value part starts with `regex:` then the part after it is treated as a regular expression that searches the web page and, if found, stores the match in the variable. At least one subgroup must be present so that the matched value can be extracted. User macros and `{HOST.*}` macros are supported. Variables are automatically URL-encoded when used in query fields or form data for post variables, but must be URL-encoded manually when used in raw post or directly in URL.
Headers	HTTP Headers are used when performing a request. Default and custom headers can be used. Headers will be assigned using default settings depending on the Agent type selected from a drop-down list on a scenario level, and will be applied to all the steps, unless they are custom defined on a step level. **It should be noted that defining the header on a step level automatically discards all the previously defined headers, except for a default header that is assigned by selecting the 'User-Agent' from a drop-down list on a scenario level.** However, even the 'User-Agent' default header can be overridden by specifying it on a step level. To unset the header on a scenario level, the header should be named and attributed with no value on a step level. Headers should be listed using the same syntax as they would appear in the HTTP protocol, optionally using some additional features supported by the CURLOPT_HTTPHEADER cURL option. For example: Accept-Charset=utf-8 Accept-Language=en-US Content-Type=application/xml; charset=utf-8 User macros and `{HOST.*}` macros are supported.
Enabled	The scenario is active if this box is checked; otherwise, it is disabled.

Note that when editing an existing scenario, two extra buttons are available in the form:

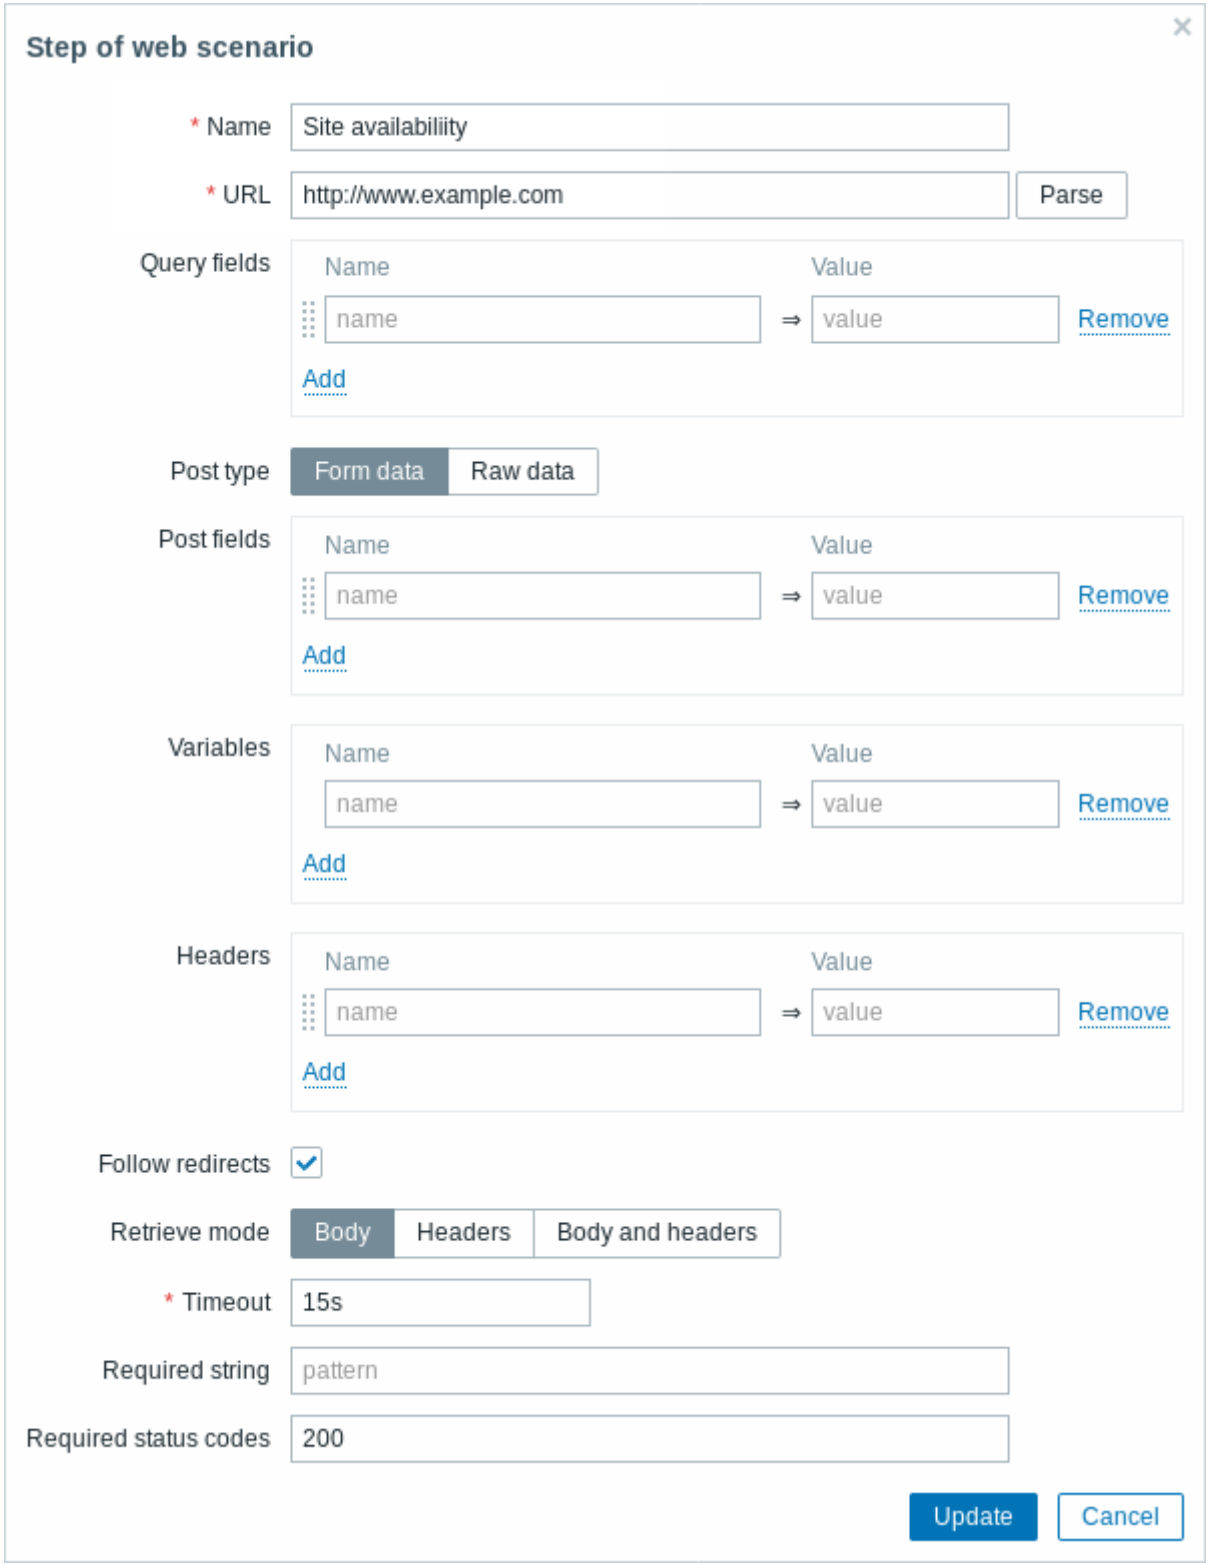
Clone	Create another scenario based on the properties of the existing one.
Clear history and trends	Delete history and trend data for the scenario. This will make the server perform the scenario immediately after deleting the data.

Note:
If *HTTP proxy* field is left empty, another way for using an HTTP proxy is to set proxy related environment variables.
For HTTP checks - set the `http_proxy` environment variable for the Zabbix server user. For example, `http_proxy=http://proxy_ip:proxy_port`.
For HTTPS checks - set the `HTTPS_PROXY` environment variable. For example, `HTTPS_PROXY=http://proxy_ip:proxy_port`.
More details are available by running a shell command: `# man curl`.

The **Steps** tab allows you to configure the web scenario steps. To add a web scenario step, click *Add* in the *Steps* block.



Note:
Secret user macros must not be used in URLs as they will resolve to "******".



Configuring steps

Step parameters:

Parameter	Description
Name	Unique step name. User macros are supported. *Note* that if user macros are used, these macros will be left unresolved in web monitoring item names.

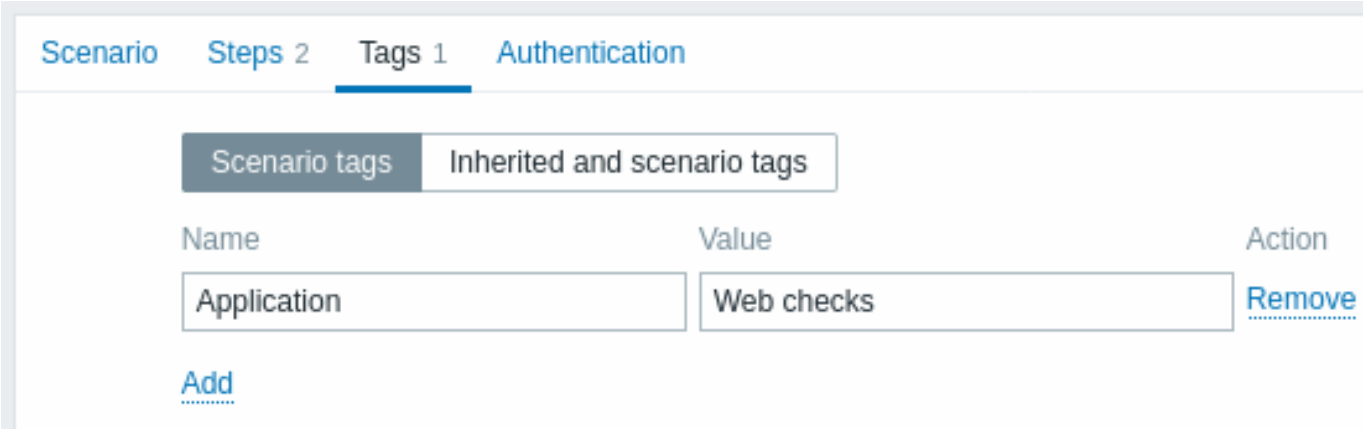
Parameter	Description
URL	URL to connect to and retrieve data. For example: https://www.example.com http://www.example.com/download Domain names can be specified in Unicode characters. They are automatically punycode-converted to ASCII when executing the web scenario step. The *Parse* button can be used to separate optional query fields (like `?key=value&limit=10`) from the URL, moving the attributes and values into *Query fields* for automatic URL-encoding. Variables can be used in the URL, using the {macro} syntax. Variables can be URL-encoded manually using a {{macro}.urlencode()} syntax. User macros and `{HOST.*}` macros are supported. For security reasons, do not include actual usernames or passwords in URLs. Use user macros instead to avoid exposing credentials. Limited to 2048 characters.
Query fields	HTTP GET variables for the URL. Specified as attribute and value pairs. Values are URL-encoded automatically. Values from scenario variables, user macros or `{HOST.*}` macros are resolved and then URL-encoded automatically. Using a {{macro}.urlencode()} syntax will double URL-encode them. User macros and `{HOST.*}` macros are supported.
Post	HTTP POST variables. In **Form data** mode, specified as attribute and value pairs. Values are URL-encoded automatically. Values from scenario variables, user macros or `{HOST.*}` macros are resolved and then URL-encoded automatically. In **Raw data** mode, attributes/values are displayed on a single line and concatenated with a **&** symbol. Raw values can be URL-encoded/decoded manually using a {{macro}.urlencode()} or {{macro}.urldecode()} syntax. For example: id=2345&userid={user} If {user} is defined as a variable of the web scenario, it will be replaced by its value when the step is executed. If you wish to URL-encode the variable, substitute {user} with {{user}.urlencode()}. User macros and `{HOST.*}` macros are supported.
Variables	Step-level variables that may be used for GET and POST functions. Specified as attribute and value pairs. Step-level variables override scenario-level variables or variables from the previous step. However, the value of a step-level variable only affects the step after (and not the current step). They have the following format: **{macro}**=value **{macro}**=regex:<regular expression> For more information see variable description on the scenario level. Variables are automatically URL-encoded when used in query fields or form data for post variables, but must be URL-encoded manually when used in raw post or directly in URL.
Headers	Custom HTTP headers that will be sent when performing a request. Specified as attribute and value pairs. A header defined on a step level will be used for that particular step. **It should be noted that defining the header on a step level automatically discards all the previously defined headers, except for a default header that is assigned by selecting the 'User-Agent' from a drop-down list on a scenario level.** However, even the 'User-Agent' default header can be overridden by specifying it on a step level. For example, assigning the name to a header, but setting no value will unset the default header on a scenario level. User macros and `{HOST.*}` macros are supported. This sets the `CURLOPT_HTTPHEADER` cURL option.
Follow redirects	Mark the checkbox to follow HTTP redirects. This sets the CURLOPT_FOLLOWLOCATION cURL option.
Retrieve mode	Select the retrieve mode: **Body** - retrieve only body from the HTTP response **Headers** - retrieve only headers from the HTTP response **Body and headers** - retrieve body and headers from the HTTP response

Parameter	Description
Timeout	Zabbix will not spend more than the set amount of time on processing the URL (from one second to maximum of 1 hour). Actually this parameter defines the maximum time for making connection to the URL and maximum time for performing an HTTP request. Therefore, Zabbix will not spend more than **2 x Timeout** seconds on the step. Time suffixes are supported, e.g., 30s, 1m, 1h. User macros are supported.
Required string	Required regular expression pattern. Unless retrieved content (HTML) matches the required pattern the step will fail. If empty, no check on required string is performed. For example: Homepage of Zabbix Welcome.*admin *Note*: Referencing regular expressions created in the Zabbix frontend is not supported in this field. User macros and `{HOST.*}` macros are supported.
Required status codes	List of expected HTTP status codes. If Zabbix gets a code which is not in the list, the step will fail. If empty, no check on status codes is performed. For example: 200,201,210-299 User macros are supported.

Note:
Any changes in web scenario steps will only be saved when the whole scenario is saved.

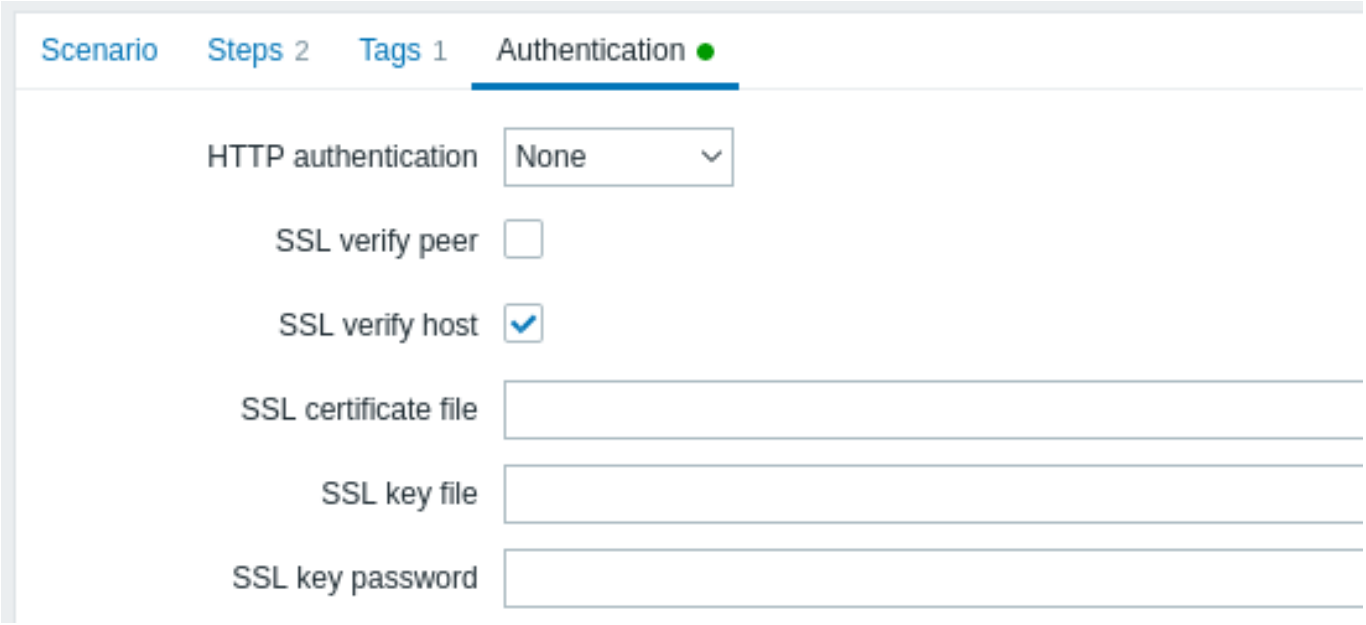
See also a real-life example of how web monitoring steps can be configured.

Configuring tags The **Tags** tab allows you to define scenario-level tags.



Tagging makes it possible to filter web scenarios and web monitoring items.

Configuring authentication The **Authentication** tab allows you to configure scenario authentication options. A green dot next to the tab name indicates that some type of HTTP authentication is enabled.



Authentication parameters:

Parameter	Description
HTTP authentication	Select the authentication option: **None** - no authentication used; **Basic** - basic authentication is used; **NTLM** - NTLM (Windows NT LAN Manager) authentication is used; **Kerberos** - Kerberos authentication is used (see also: Configuring Kerberos with Zabbix); **Digest** - Digest authentication is used.
User	Enter the user name (up to 255 characters). This field is available if *HTTP authentication* is set to Basic, NTLM, Kerberos, or Digest. User macros are supported.
Password	Enter the user password (up to 255 characters). This field is available if *HTTP authentication* is set to Basic, NTLM, Kerberos, or Digest. User macros are supported.
SSL verify peer	Mark the checkbox to verify the SSL certificate of the web server. The server certificate is verified against the trusted certificate authority (CA) certificates from the system-wide CA location. You can override the location of CA files using Zabbix server or proxy configuration parameter SSLCALocation. This sets the CURLOPT_SSL_VERIFYPEER cURL option.
SSL verify host	Mark the checkbox to verify that the *Common Name* field or the *Subject Alternate Name* field of the web server certificate matches. This sets the CURLOPT_SSL_VERIFYHOST cURL option.
SSL certificate file	Name of the SSL certificate file used for client authentication. The certificate file must be in PEM[1] format. If the certificate file contains also the private key, leave the *SSL key file* field empty. If the key is encrypted, specify the password in *SSL key password* field. The directory containing this file is specified by Zabbix server or proxy configuration parameter SSLCertLocation. `HOST.*` macros and user macros can be used in this field. This sets the CURLOPT_SSLCERT cURL option.
SSL key file	Name of the SSL private key file used for client authentication. The private key file must be in PEM[1] format. The directory containing this file is specified by Zabbix server or proxy configuration parameter SSLKeyLocation. `HOST.*` macros and user macros can be used in this field. This sets the CURLOPT_SSLKEY cURL option.
SSL key password	SSL private key file password. User macros can be used in this field. This sets the CURLOPT_KEYPASSWD cURL option.

Attention:
[1] Zabbix supports certificate and private key files in PEM format only. In case you have your certificate and private key data in PKCS #12 format file (usually with extension *.p12 or *.pfx) you may generate the PEM file from it using the following commands:

```
openssl pkcs12 -in ssl-cert.p12 -clcerts -nokeys -out ssl-cert.pem
openssl pkcs12 -in ssl-cert.p12 -nocerts -nodes  -out ssl-cert.key
```

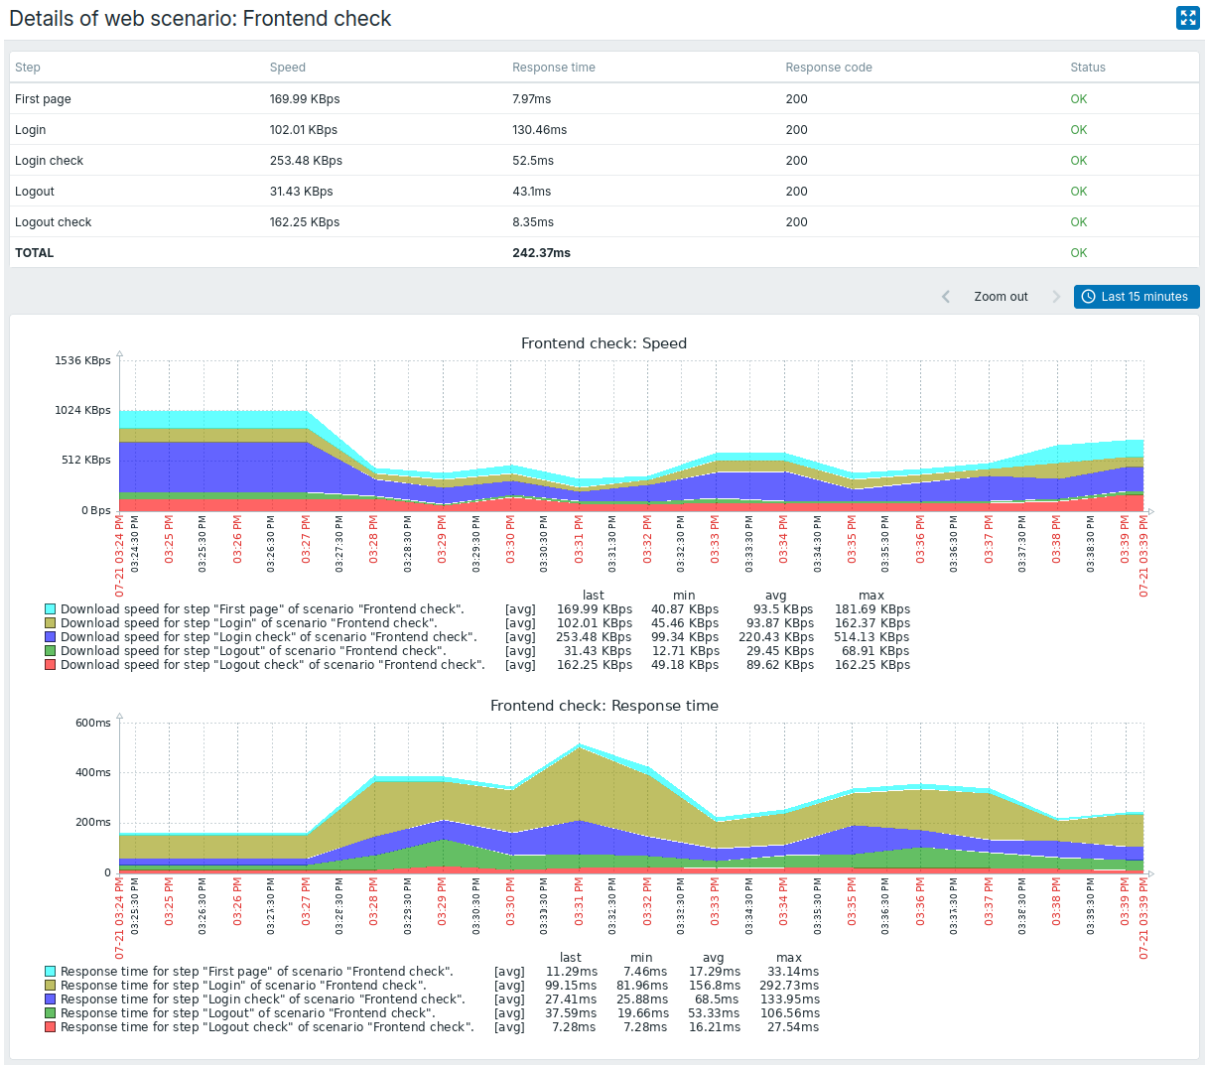
Note:
Zabbix server picks up changes in certificates without a restart.

Note:
If you have client certificate and private key in a single file just specify it in a "SSL certificate file" field and leave "SSL key file" field empty. The certificate and key must still be in PEM format. Combining certificate and key is easy:

```
cat client.crt client.key > client.pem
```

Display To view web scenarios configured for a host:

1. Go to *Monitoring* > *Hosts*.
2. Locate the host in the list and click the *Web* hyperlink in the last column.
3. Click the scenario name to get detailed information.



An overview of web scenarios can also be displayed in *Dashboards* by the Web monitoring widget.

Recent results of the web scenario execution are available in the *Monitoring* > *Latest data* section.

Extended monitoring Sometimes it is necessary to log received HTML page content. This is especially useful if some web scenario step fails. Debug level 5 (trace) serves that purpose. This level can be set in server and proxy configuration files or

using a runtime control option (`-R log_level_increase="http poller,N"`, where N is the process number). The following examples demonstrate how extended monitoring can be started provided debug level 4 is already set:

```
# Increase log level of all http pollers:
zabbix_server -R log_level_increase="http poller"

# Increase log level of second http poller:
zabbix_server -R log_level_increase="http poller,2"
```

If extended web monitoring is not required it can be stopped using the `-R log_level_decrease` option.

Web monitoring items

Overview

Some new items are automatically added for monitoring when web scenarios are created.

All items inherit tags from the web scenario.

Scenario items

As soon as a scenario is created, Zabbix automatically adds the following items for monitoring.

Item	Description
Download speed for scenario <Scenario>	This item will collect information about the download speed (bytes per second) of the whole scenario, i.e. average for all steps. Item key: web.test.in[Scenario„bps] Type: *Numeric(float)*
Failed step of scenario <Scenario>	This item will display the number of the step that failed on the scenario. If all steps are executed successfully, 0 is returned. Item key: web.test.fail[Scenario] Type: *Numeric(unsigned)*
Last error message of scenario <Scenario>	This item returns the last error message text of the scenario. A new value is stored only if the scenario has a failed step. If all steps are ok, no new value is collected. Item key: web.test.error[Scenario] Type: *Character*

The actual scenario name will be used instead of ”Scenario”.

Note:
If the scenario name contains user macros, these macros will be left unresolved in web monitoring item names.

 If the scenario name starts with a doublequote or contains a comma or a square bracket, it will be properly quoted in item keys. In other cases no additional quoting will be performed.

Note:
Web monitoring items are added with a 30 day history and a 90 day trend retention period.

These items can be used to create triggers and define notification conditions.

Example 1

To create a ”Web scenario failed” trigger, you can define a trigger expression:

```
last(/host/web.test.fail[Scenario])<>0
```

Make sure to replace ’Scenario’ with the real name of your scenario.

Example 2

To create a ”Web scenario failed” trigger with a useful problem description in the trigger name, you can define a trigger with name:

```
Web scenario "Scenario" failed: {ITEM.VALUE}
```

and trigger expression:

```
length(last(/host/web.test.error[Scenario]))>0 and last(/host/web.test.fail[Scenario])>0
```

Make sure to replace 'Scenario' with the real name of your scenario.

Example 3

To create a "Web application is slow" trigger, you can define a trigger expression:

```
last(/host/web.test.in[Scenario,,bps])<10000
```

Make sure to replace 'Scenario' with the real name of your scenario.

Scenario step items

As soon as a step is created, Zabbix automatically adds the following items for monitoring.

Item	Description
Download speed for step <Step> of scenario <Scenario>	This item will collect information about the download speed (bytes per second) of the step. Item key: web.test.in[Scenario,Step,bps] Type: *Numeric(float)*
Response time for step <Step> of scenario <Scenario>	This item will collect information about the response time of the step in seconds. Response time is counted from the beginning of the request until all information has been transferred. Item key: web.test.time[Scenario,Step,resp] Type: *Numeric(float)*
Response code for step <Step> of scenario <Scenario>	This item will collect response codes of the step. Item key: web.test.rspcode[Scenario,Step] Type: *Numeric(unsigned)*

Actual scenario and step names will be used instead of "Scenario" and "Step" respectively.

Note:
Web monitoring items are added with a 30 day history and a 90 day trend retention period.

Note:
If scenario name starts with a doublequote or contains comma or square bracket, it will be properly quoted in item keys. In other cases no additional quoting will be performed.

These items can be used to create triggers and define notification conditions. For example, to create a "Zabbix GUI login is too slow" trigger, you can define a trigger expression:

```
last(/zabbix/web.test.time[ZABBIX GUI,Login,resp])>3
```

Real-life scenario

Overview

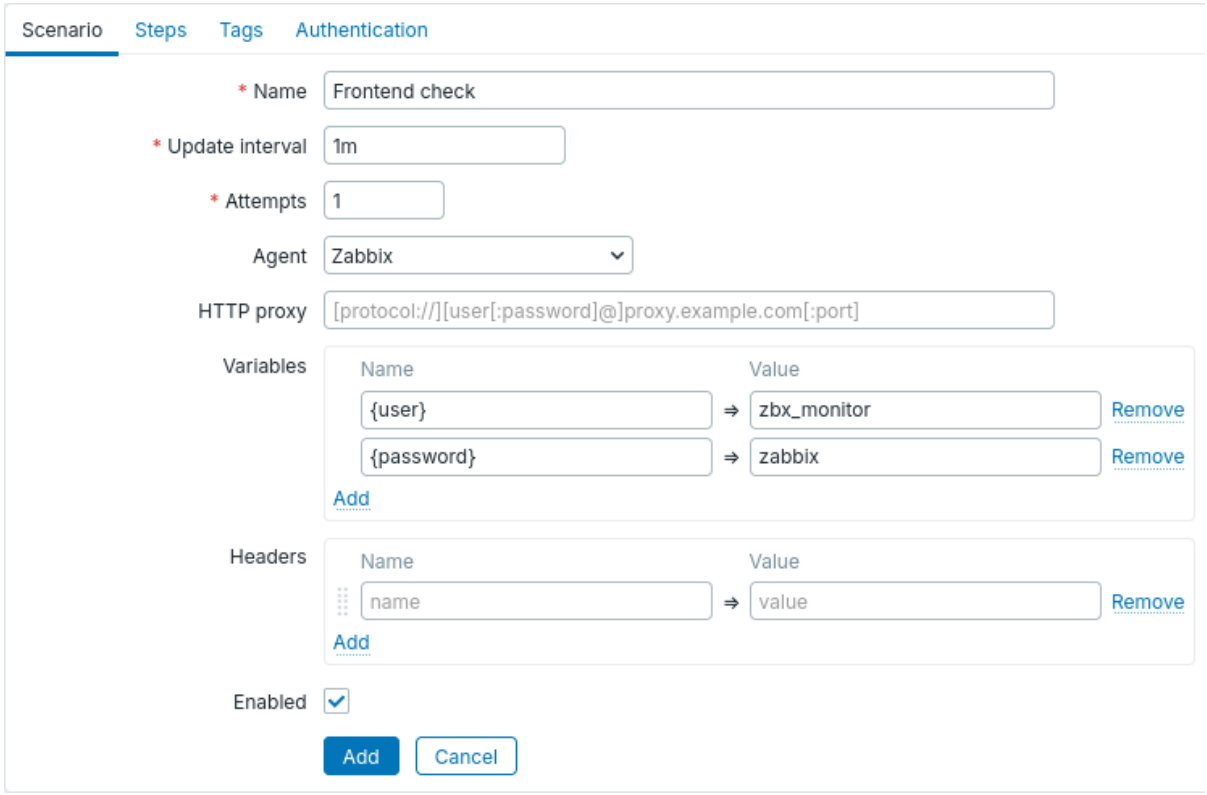
This section presents a step-by-step real-life example of how web monitoring can be used.

Zabbix web monitoring will be used to monitor Zabbix frontend. The goal is to determine if it is available, provides the right content, and how quickly it works. To achieve this, several steps are required, including checking the availability of the first page, logging in with a username and password, verifying the login success, logging out, and confirming the logout.

Scenario

Add a new web scenario

1. Go to *Data collection* > *Hosts*.
2. Choose a host and click *Web* in the row of that host.
3. Click *Create web scenario*.



4. In the new scenario form, on the *Scenario* tab, fill out the following fields:

    - *Name*: Frontend check.
    - *Update interval*: 1m.
    - *Attempts*: 1.
    - *Agent*: Zabbix.
 In the *Variables* section, add the following name-value pairs:
    - `{user}`: your existing Zabbix user credentials;
    - `{password}`: your existing Zabbix user credentials.

Note:
For safety reasons, it is recommended to create a separate user with minimal permissions to use for monitoring purposes.

Optionally, switch to the *Tags* tab and add web scenario tags. You can use web scenario tags to quickly identify related items and triggers or search through collected data. For example, suitable tags for this tutorial are `component: web-scenario` and/or `target: frontend`.

Configure web scenario steps

Switch to the *Steps* tab and define steps for the scenario.
 After configuring each step, click the *Add* button.

Attention:
Note that Zabbix frontend uses JavaScript redirect when logging in, so login must occur first, and logged-in features can be checked only in further steps. Additionally, the login step must use the full URL to **index.php** file.

Common fields for all steps

In addition to the step-specific fields, fill out the following common fields for each step described below:

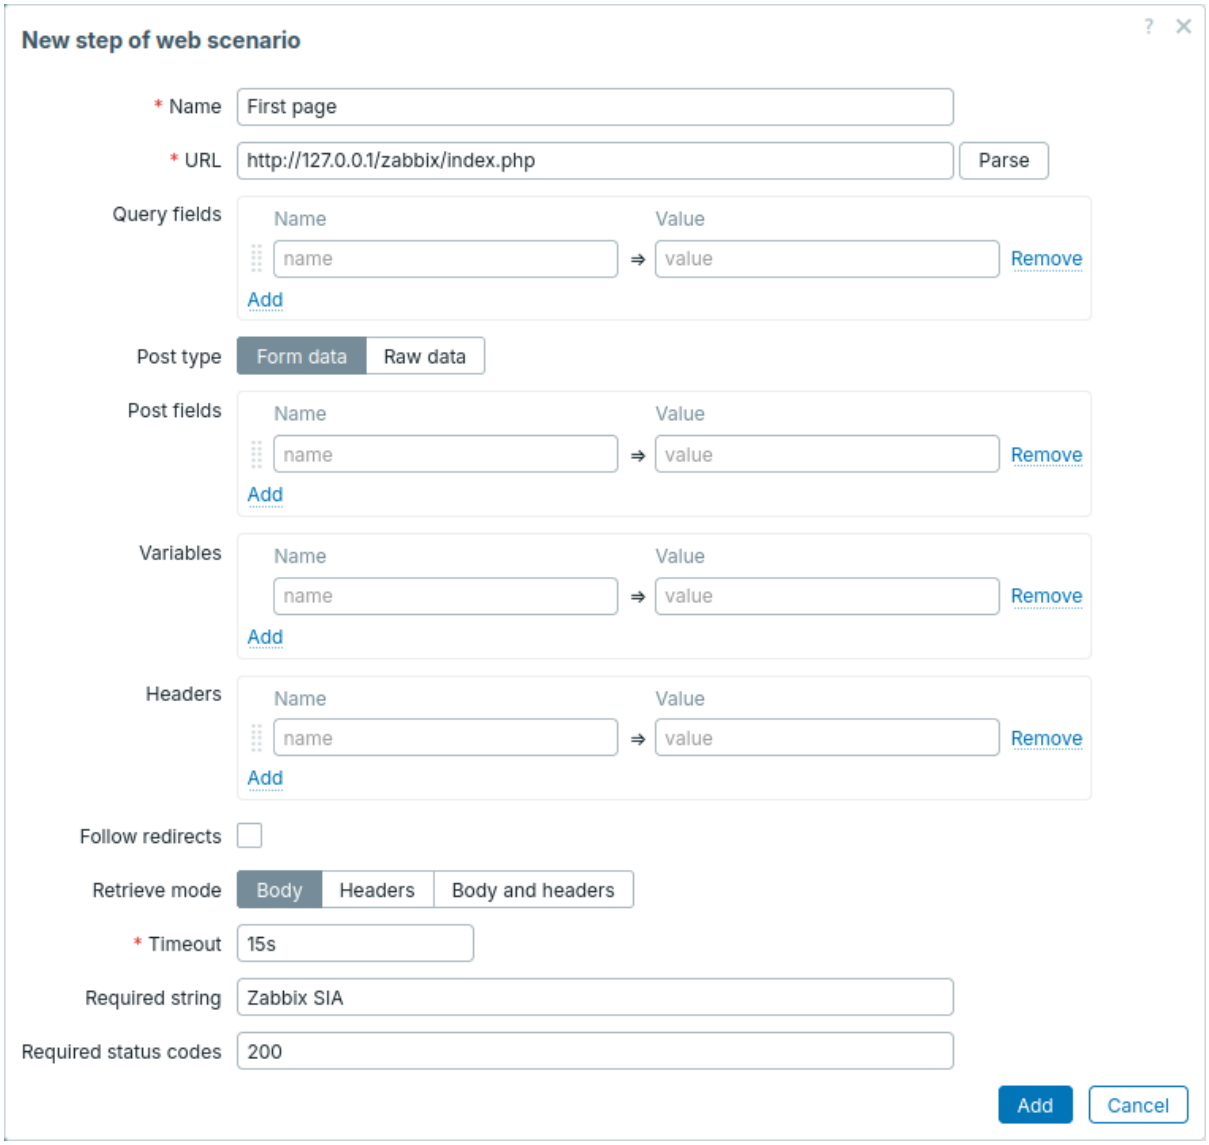
- *URL*: the URL of the Zabbix frontend (`http://<server_ip_or_name>/zabbix/index.php` (for Apache) or `http://<server_ip_o` (for NGINX)). For Web scenario step 3, use a different URL as described below.
- *Timeout*: 15s.
- *Required status codes*: 200.

Web scenario step 1

Check that the first page responds correctly, returns HTTP status code 200, and contains the text *Zabbix SIA*.
 Fill out the following fields:

- *Name*: First page.
- *Timeout*: 15s.
- *Required string*: Zabbix SIA.

- The common fields for all steps.

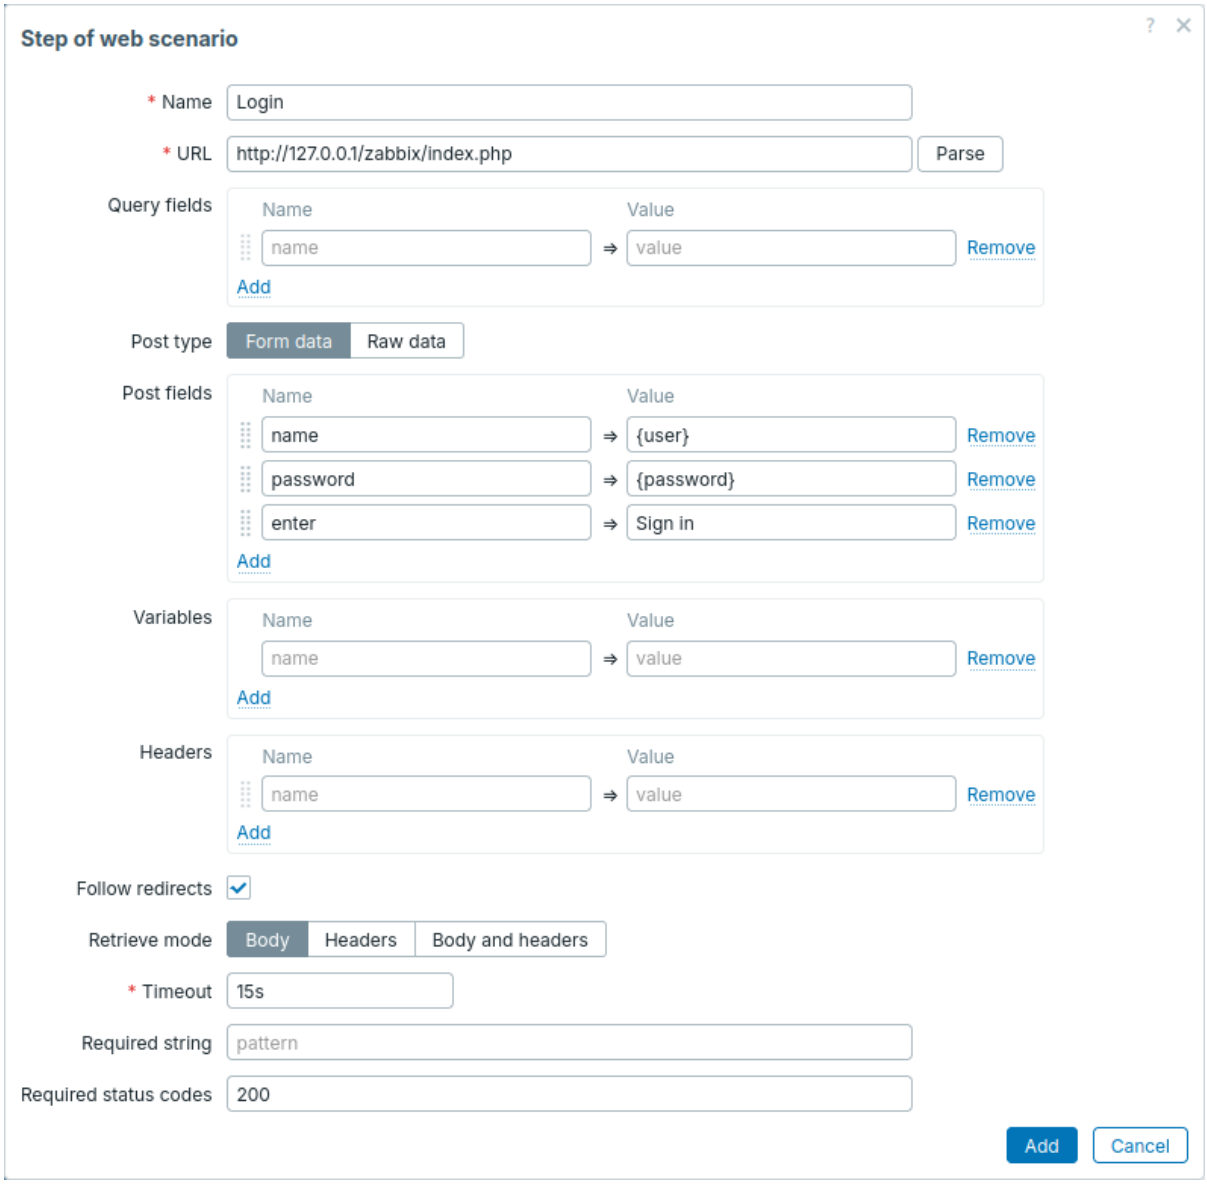


Web scenario step 2

Log in to the Zabbix frontend using the macros (variables) defined in the *Scenario* tab: *{user}* and *{password}*.
 Fill out the following fields:

- *Name*: Login.
- In the *Post fields* section, add the following name-value pairs:
  - name: `{user}`;
  - password: `{password}`;
  - enter: Sign in.
- The common fields for all steps.

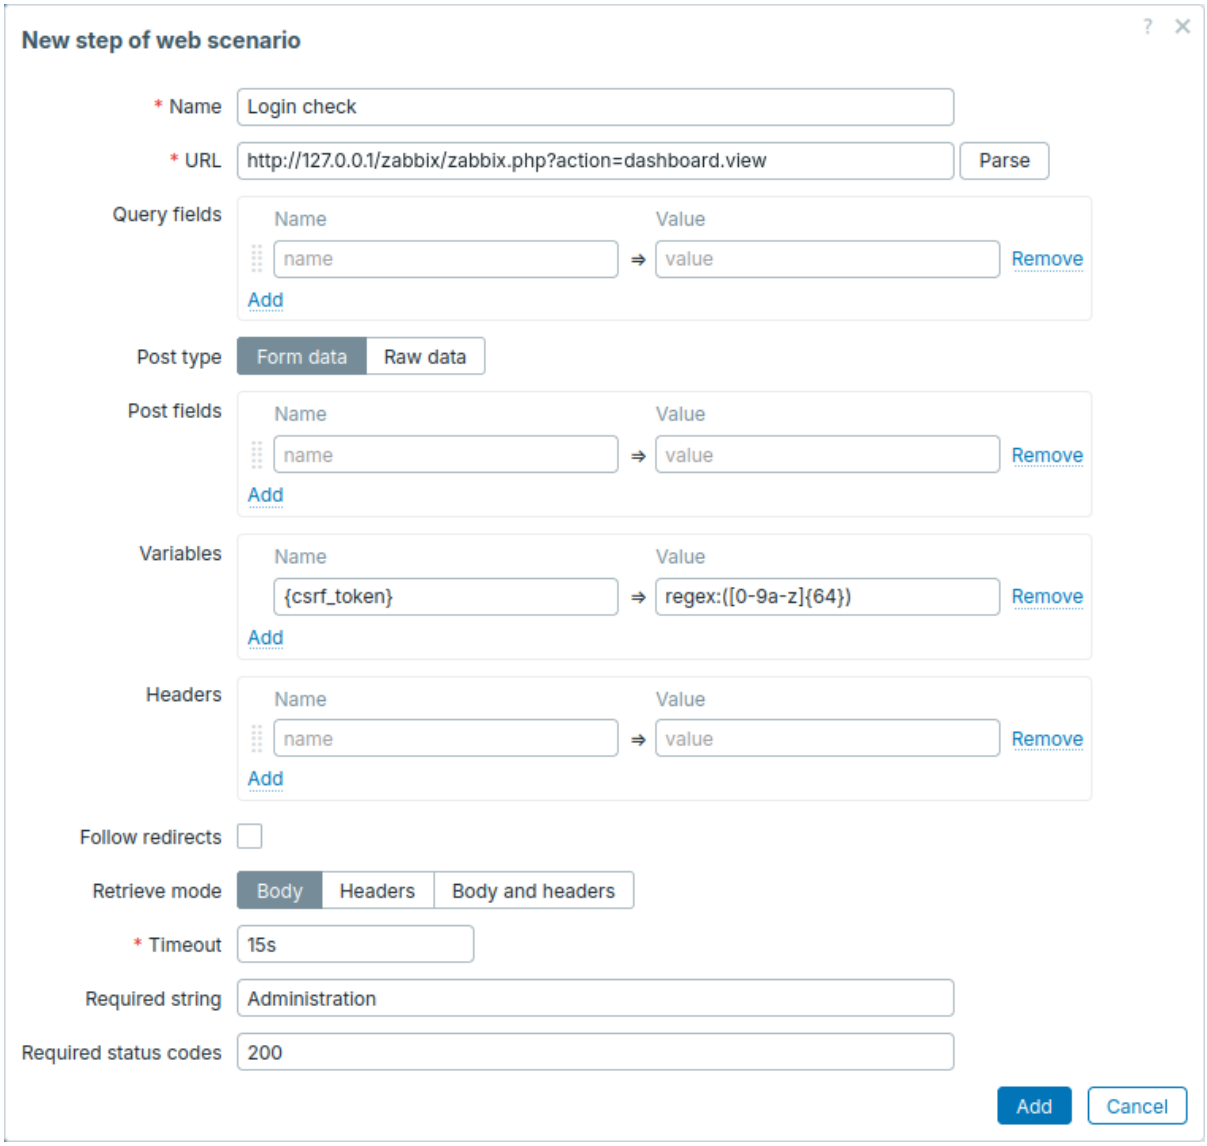
Select the *Follow redirects* checkbox.



Web scenario step 3

After logging in, verify success by checking for a string visible only when logged in, for example, *Administration*.
 Fill out the following fields:

- *Name*: Login check.
- *URL*: the URL to which the user is redirected after login (for example, http://127.0.0.1/zabbix/zabbix.php?action=dashboard.view).
- In the *Variables* section, add the name-value pair `{csrf_token}: regex:([0-9a-z]{64})`. This variable will catch the value of the assigned CSRF token to reuse in the web scenario step 4.
- *Required string*: Administration.
- The common fields for all steps.



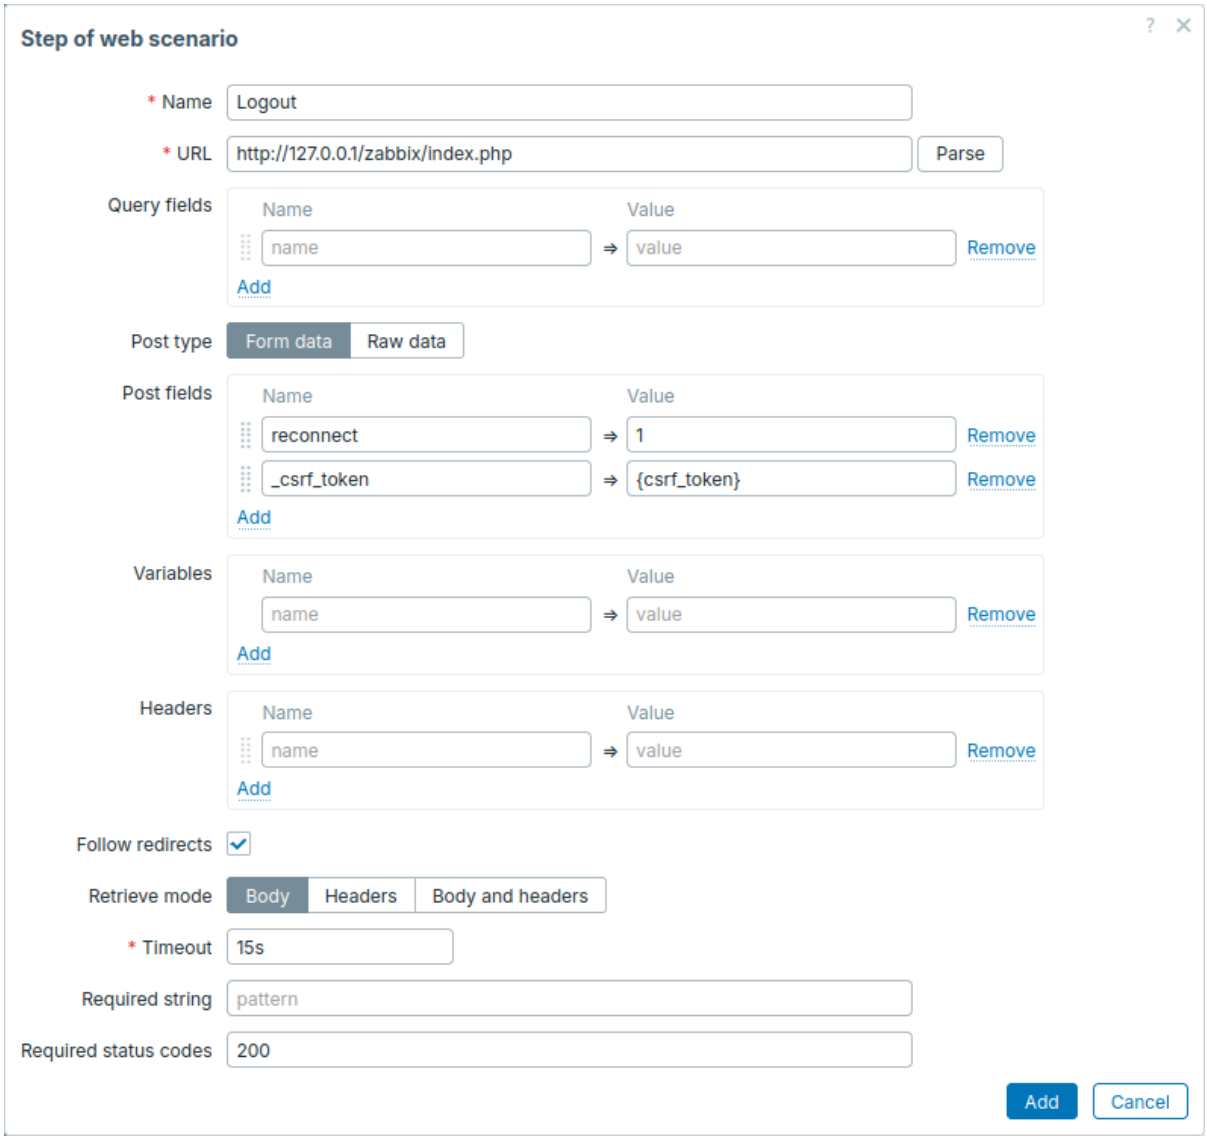
Web scenario step 4

Once the frontend's accessibility and login have been verified, add a logout step; otherwise, the Zabbix database will become cluttered with many open session records.
 Fill out the following fields:

- *Name*: Logout.
- In the *Post fields* section, add the following name-value pairs:
  - reconnect: 1;
  - _csrf_token: `{csrf_token}`.
- The common fields for all steps.

Select the *Follow redirects* checkbox.

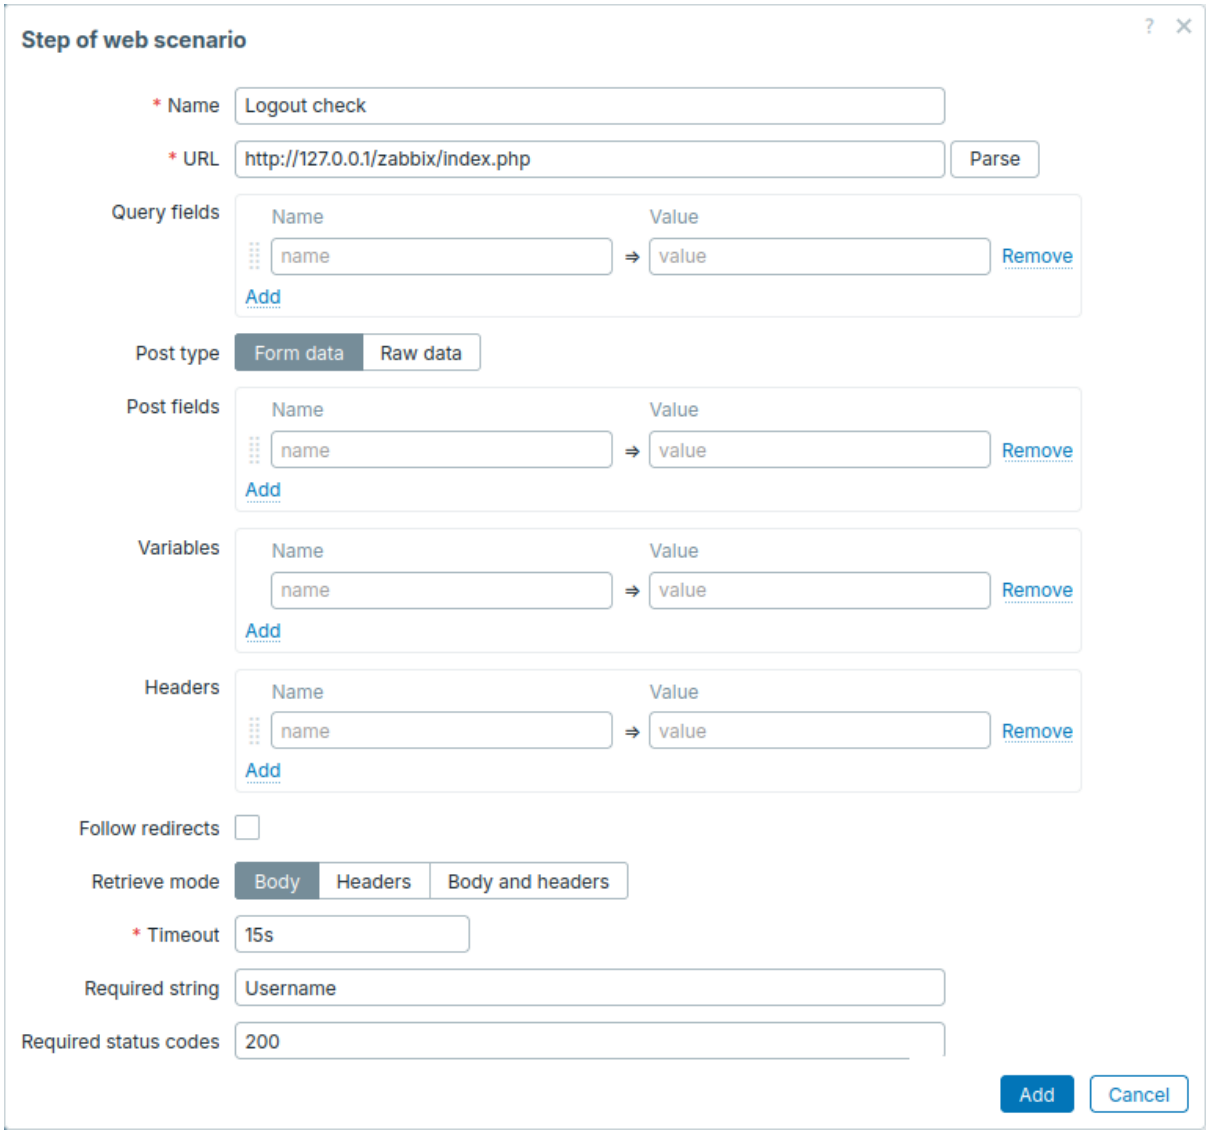
This step uses the variable `{csrf_token}` obtained in the web scenario step 3.



Web scenario step 5

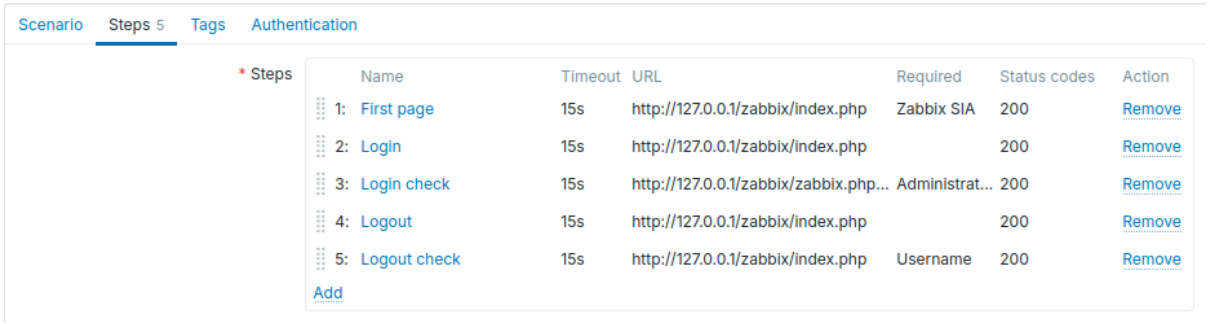
To confirm the logout, check for the *Username* string.
 Fill out the following fields:

- *Name*: Logout check.
- *Required string*: Username.
- The common fields for all steps.



Full configuration of steps

A complete configuration of web scenario steps should look like this:

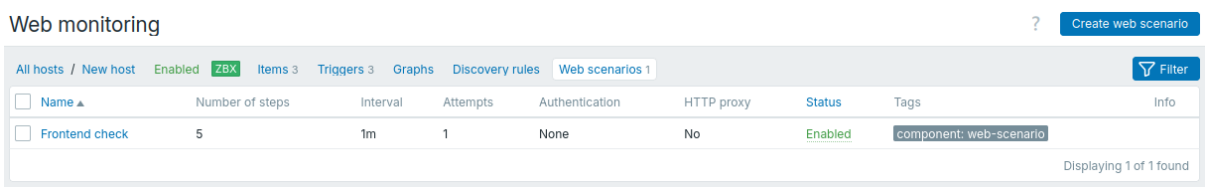


Check the results

Save the finished web monitoring scenario.

When the web scenario is created, new web monitoring items are automatically added for monitoring.

The scenario will be added to the host. To view web scenario information go to *Monitoring* > *Hosts*, locate the host in the list and click the *Web* hyperlink in the last column.



Click the scenario name to see detailed statistics:



Service monitoring

Overview Service monitoring is a business-level monitoring that can be used to get an overview of the entire IT infrastructure service tree, identify weak places of the infrastructure, calculate SLA of various IT services, and check out other information at a higher level. Service monitoring focuses on the overall availability of a service instead of low-level details, such as the lack of disk space, high processor load, etc. Service monitoring also provides functionality to find the root cause of a problem if a service is not performing as expected.

Service monitoring provides a way to create a hierarchical representation of monitored data.

A very simple service structure may look like:

```
Service
|
|-Workstations
| |
| |-Workstation1
| |
| |-Workstation2
|
|-Servers
```

Each node of the structure has attribute status. The status is calculated and propagated to upper levels according to the selected algorithm. The status of individual nodes is affected by the status of the mapped problems. Problem mapping is accomplished with tagging.

Zabbix can send notifications or automatically execute a script on the Zabbix server in case a change of service state is detected. It is possible to define flexible rules whether a parent service should go into a 'Problem state' based on the statuses of child services. Services problem data can then be used to calculate SLA and send SLA reports based on the flexible set of conditions.

Service monitoring is configured in the Services menu, which consists of the following sections:

- Services

Services section allows you to build a hierarchy of your monitored infrastructure by adding parent services, and then - child services to the parent services.

In addition to configuring service tree, this section provides an overview of the whole infrastructure and allows you to quickly identify the problems that led to a change of service state.

- SLA

In this section you can define service level agreements and set service level objectives for specific services.

- SLA report

In this section you can view SLA reports.

Service actions

You may also configure service actions.

Service actions are optional and allow you to:

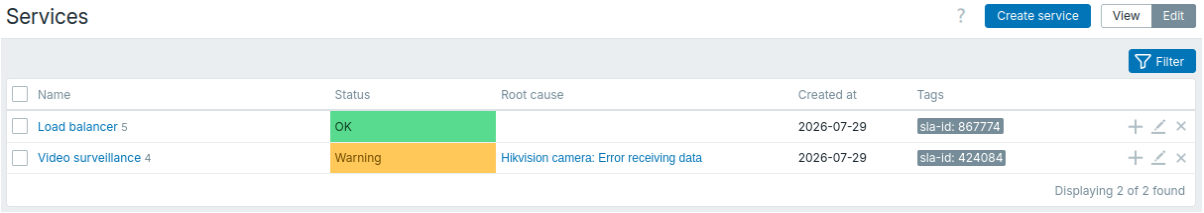
- send a notification that a service is down
- execute a remote command on Zabbix server upon a change of service state
- send a recovery notification when a service is up again.

See also:

- SLA monitoring configuration example
- Notes about upgrading services from Zabbix versions below 6.0

Service tree

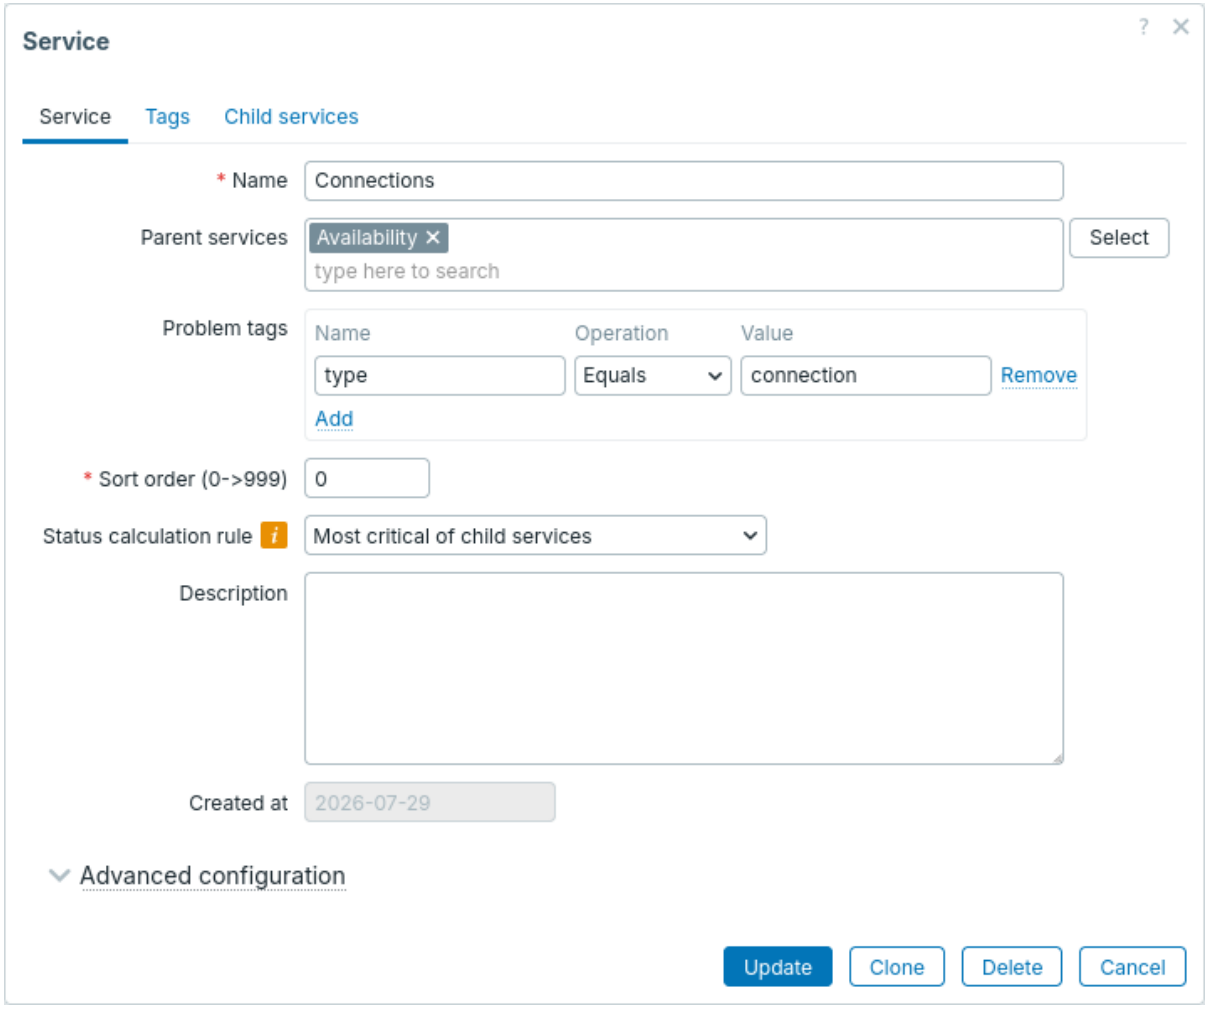
The service tree is configured in the *Services* > *Services* menu section. In the upper-right corner, switch from view to edit mode.



To configure a new service, click the *Create service* button in the upper-right corner.

To quickly add a child service, you can alternatively click the plus icon next to the parent service. This will open the same service configuration form, but the *Parent services* parameter will be pre-filled.

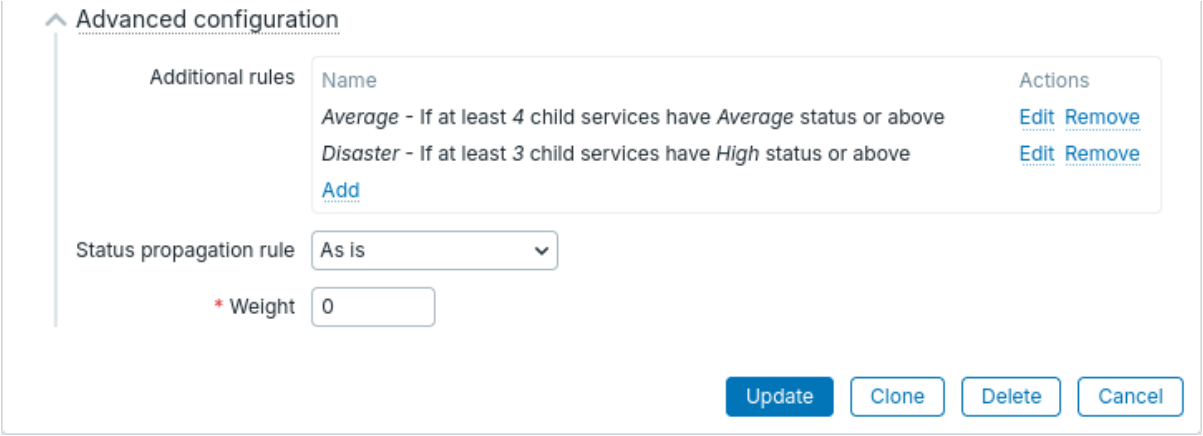
Service configuration In the **Service** tab, specify required service parameters:



All mandatory input fields are marked with a red asterisk.

Parameter	Description
Name	Service name.
Parent services	Parent services the service belongs to. Leave this field empty if you are adding the service of highest level. One service may have multiple parent services. In this case, it will be displayed in the service tree under each of the parent services.
Problem tags	Specify tags to map problem data to the service. Several conditions can be set. Tag name matching is always case-sensitive. There are two operators available for each condition: **Equals** - include the specified tag names and values (case-sensitive) **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-sensitive)
Sort order	Sort order for display, lowest comes first.
Status calculation rule	Rule for calculating service status: **Most critical if all children have problems** - the most critical problem in the child services is used to color the service status, if all children have problems **Most critical of child services** - the most critical problem in the child services is used to color the service status **Set status to OK** - do not calculate service status Additional status calculation rules can be configured in the advanced configuration options.
Description	Service description.
Created at	The time when the service was created; displayed when editing an existing service.
Advanced configuration	Click the *Advanced configuration* header to display advanced configuration options.

Advanced configuration



Parameter	Description
Additional rules	Click *Add* to configure additional status calculation rules.
Set status to	Set service status to either *OK* (default), *Not classified*, *Information*, *Warning*, *Average*, *High* or *Disaster* in case of a condition match.
Condition	Select the condition for direct child services: **if at least (N) child services have (Status) status or above** **if at least (N%) of child services have (Status) status or above** **if less than (N) child services have (Status) status or below** **if less than (N%) of child services have (Status) status or below** **if weight of child services with (Status) status or above is at least (W)** **if weight of child services with (Status) status or above is at least (N%)** **if weight of child services with (Status) status or below is less than (W)** **if weight of child services with (Status) status or below is less than (N%)** If several conditions are specified and the situation matches more than one condition, the highest severity will be set.
N (W)	Set the value of N or W (1-100000), or N% (1-100) in the condition.
Status	Select the value of *Status* in the condition: *OK* (default), *Not classified*, *Information*, *Warning*, *Average*, *High* or *Disaster*.
Status propagation rule	Rule for propagating the service status to the parent service: **As is** - the status is propagated without change **Increase by** - you may increase the propagated status by 1 to 5 severities **Decrease by** - you may decrease the propagated status by 1 to 5 severities **Ignore this service** - the status is not propagated to the parent service at all **Fixed status** - the status is propagated statically, i.e. as always the same
Weight	Weight of the service (integer in the range from 0 (default) to 1000000).

Note:
Additional status calculation rules can only be used to increase severity level over the level calculated according to the main *Status calculation rule* parameter. If according to additional rules the status should be Warning, but according to the *Status calculation rule* the status is Disaster - the service will have status Disaster.

The *Tags* tab contains service-level tags. Service-level tags are used to identify a service. Tags of this type are not used to map problems to the service (for that, use *Problem tags* from the first tab).

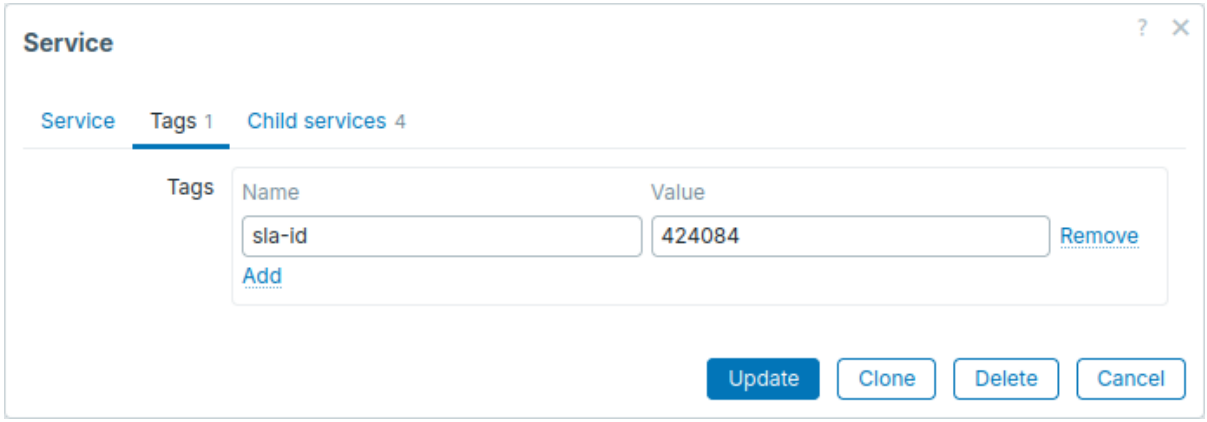
The *Child services* tab allows you to specify dependant services. Click *Add* to add a service from the list of existing services. If you want to add a new child service, save this service first, then click a plus icon next to the service that you have just created.

Tags There are two different types of tags in services:

- Service tags
- Problem tags

Service tags

Service tags are used to match services with service actions and SLAs. These tags are specified at the *Tags* service configuration tab. For mapping SLAs, *OR* logic is used: a service will be mapped to an SLA if it has at least one matching tag. In service actions, mapping rules are configurable and can use either *AND*, *OR*, or *AND/OR* logic.

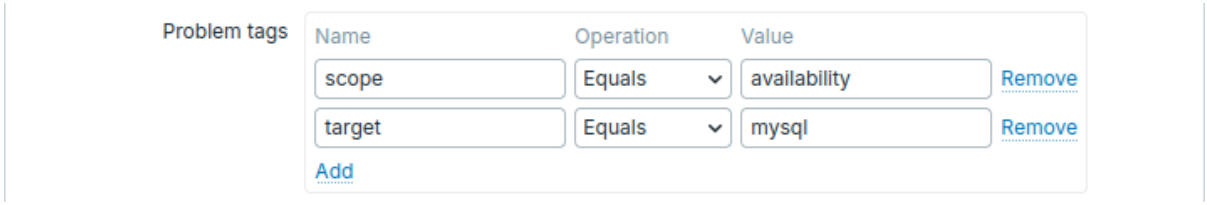


Problem tags

Problem tags are used to match problems and services. These tags are specified at the primary service configuration tab.

Only child services of the lowest hierarchy level may have problem tags defined and be directly correlated to problems. If problem tags match, the service status will change to the same status as the problem has. In case of several problems, a service will have the status of the most severe one. Status of a parent service is then calculated based on child services statuses according to Status calculation rules.

If several tags are specified, *AND* logic is used: a problem must have all tags specified in the service configuration to be mapped to the service.

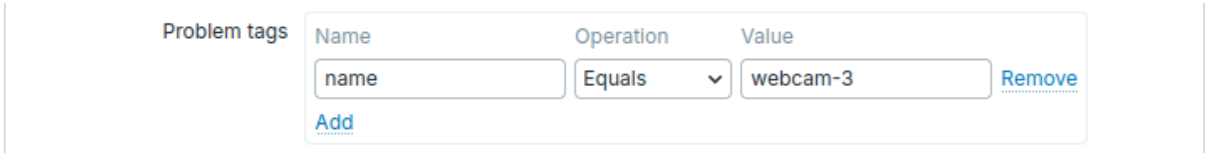


Note:
A problem in Zabbix inherits tags from the whole chain of templates, hosts, items, web scenarios, and triggers. Any of these tags can be used for matching problems to services.

Example:

Problem *Web camera 3 is down* has tags `type:video-surveillance`, `floor:1` and `name:webcam-3` and status *Warning*

The service **Web camera 3** has the only problem tag specified: `name:webcam-3`



Service status will change from *OK* to *Warning* when this problem is detected.

If the service **Web camera 3** had problem tags `name:webcam-3` and `floor:2`, its status would not be changed, when the problem is detected, because the conditions are only partially met.

Modifying existing services

Note:
The buttons described below are visible only when the *Services* section is in edit mode.

To edit an existing service, click the pencil icon next to the service.

To clone an existing service, click the pencil icon to open its configuration and then click the *Clone* button. When a service is cloned, its parent links are preserved, while the child links are not.

To delete a service, click the x icon next to it. When you delete a parent service, its child services will not be deleted and will move one level higher in the service tree (1st level children will get the same level as the deleted parent service).

Two buttons below the list of services offer some mass-editing options:

- *Mass update* - mass update service properties

- *Delete* - delete the services

To use these options, mark the checkboxes before the respective services, then click the required button.

SLA

Overview Once the services are created, you can start monitoring whether their performance is on track with service level agreement (SLA).

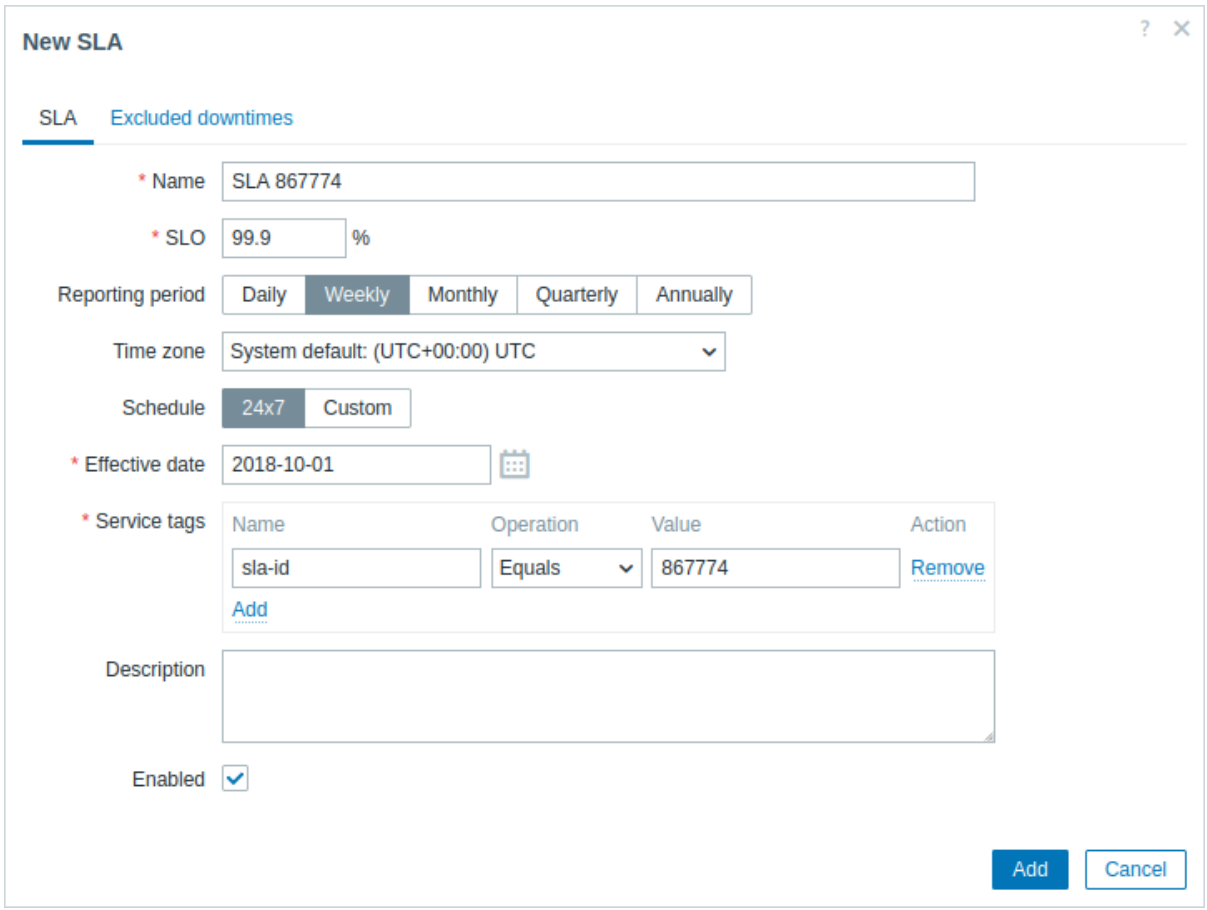
Services > *SLA* menu section allows SLAs to be configured for various services. An SLA in Zabbix defines service level objective (SLO), expected uptime schedule and planned downtimes.

SLAs and services are matched by service tags. The same SLA may be applied to multiple services - performance will be measured for each matching service separately. A single service may have multiple SLAs assigned - data for each of the SLAs will be displayed separately.

In SLA reports Zabbix provides Service level indicator (SLI) data, which measures real service availability. Whether a service meets the SLA targets is determined by comparing SLO (expected availability in %) with SLI (real-life availability in %).

Configuration To create a new SLA, click the *Create SLA* button.

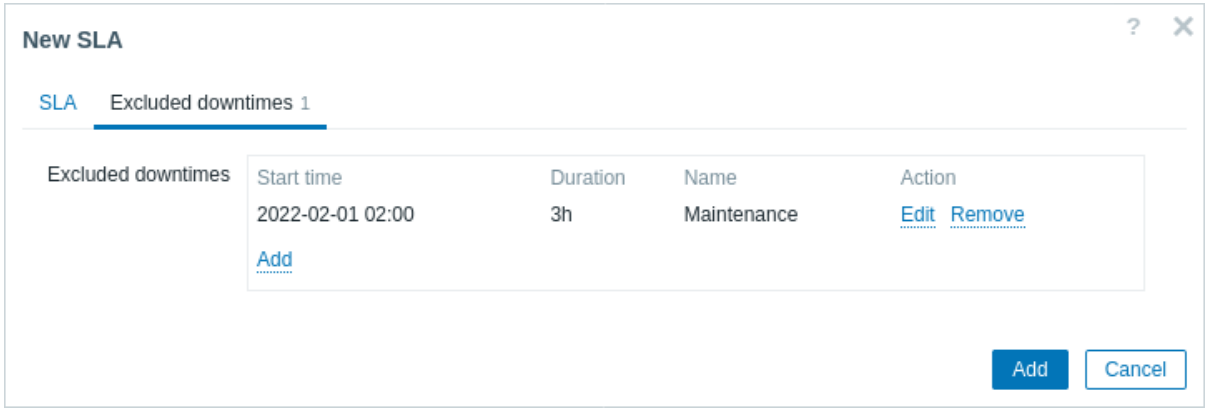
The *SLA* tab allows you to specify general SLA parameters.



Parameter	Description
Name	Enter the SLA name.
SLO	Enter the service level objective (SLO) as percentage.
Reporting period	Select the SLI reporting period used for calculating the SLI data in the SLA report: **Daily** - each day, from 00:00:00 to 23:59:59. **Weekly** - each week, from Sunday 00:00:00 to Saturday 23:59:59. **Monthly** - each month, from the first day 00:00:00 to the last day 23:59:59. **Quarterly** - each calendar quarter: Q1 (Jan-Mar), Q2 (Apr-Jun), Q3 (Jul-Sep), Q4 (Oct-Dec). **Annually** - each calendar year, from January 1 00:00:00 to December 31 23:59:59.
Time zone	Select the SLA time zone.

Parameter	Description
Schedule	Select the SLA schedule - 24x7 or custom.
Effective date	Select the date of starting SLA calculation.
Service tags	Add service tags to identify the services towards which this SLA should be applied. **Name** - service tag name, must be exact match, case-sensitive. **Operation** - select *Equals* if the tag value must match exactly (case-sensitive) or *Contains* if part of the tag value must match (case-insensitive). **Value** - service tag value to search for according to selected operation. The SLA is applied to a service, if at least one service tag matches.
Description	Add a description for the SLA.
Enabled	Mark the checkbox to enable the SLA calculation.

The *Excluded downtimes* tab allows you to specify downtimes that are excluded from the SLA calculation.



Click *Add* to configure excluded downtimes, then enter the period name, start date and duration.

SLA reports How a service performs compared to an SLA is visible in the SLA report. SLA reports can be viewed:

- from the *SLA* section by clicking the SLA report hyperlink;
- from the *Services* section by clicking the SLA name in the info tab;
- in the Dashboard widget *SLA report*.

Once an SLA is configured, the *Info* tab in the services section will also display some information about service performance.

Setup example

Overview This section describes a simple setup for monitoring Zabbix high availability cluster as a service.

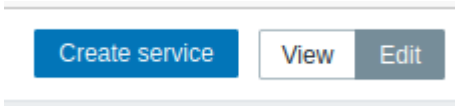
Pre-requisites Prior to configuring service monitoring, you need to have the hosts configured:

- *HA node 1* with at least one trigger and a tag (preferably set on a trigger level) `component:ha-node-1`
- *HA node 2* with at least one trigger and a tag (preferably set on a trigger level) `component:ha-node-2`

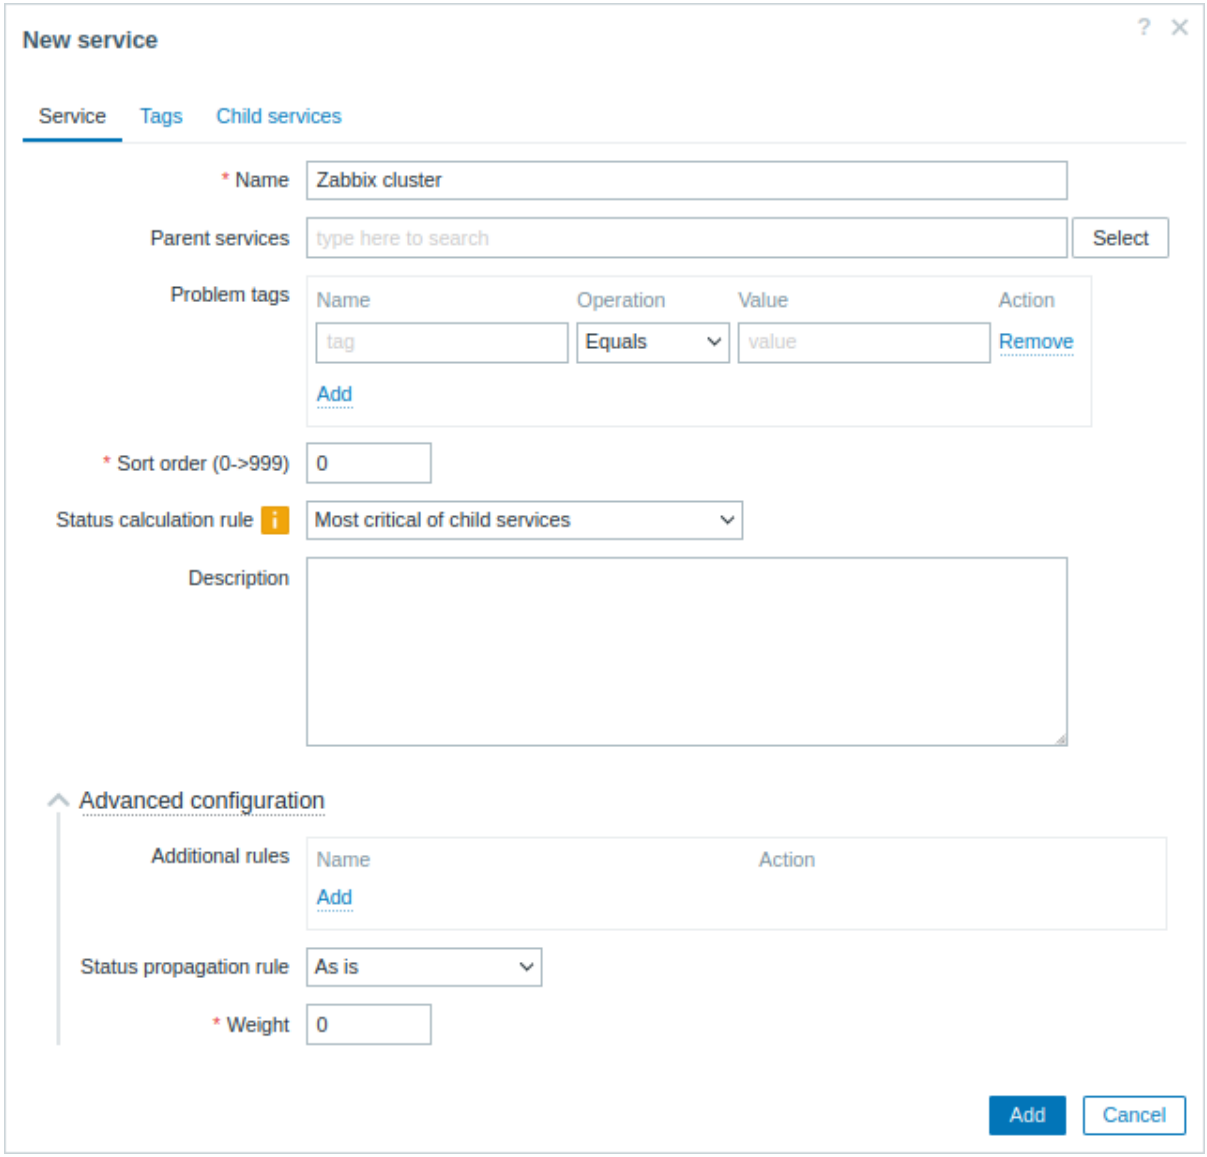
Service tree The next step is to build the service tree. In this example, the infrastructure is very basic and consists of three services: *Zabbix cluster* (parent) and two child services *Zabbix server node 1* and *Zabbix server node 2*.

```
Zabbix cluster
|
|- Zabbix server node 1
|- Zabbix server node 2
```

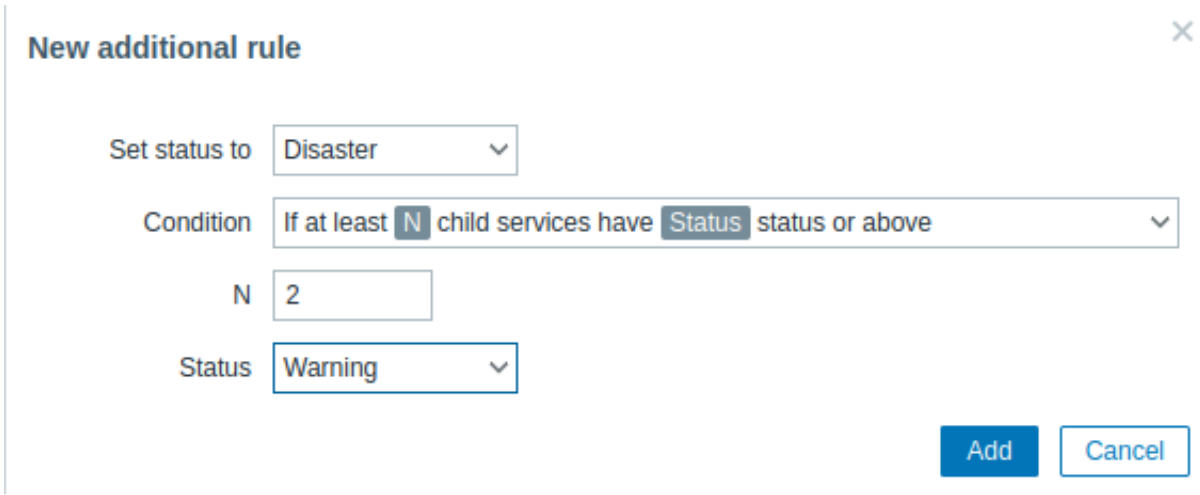
At the Services page, turn on *Edit* mode and click *Create service*:



In the service configuration window, enter name *Zabbix cluster* and click the *Advanced configuration* header to display advanced configuration options.



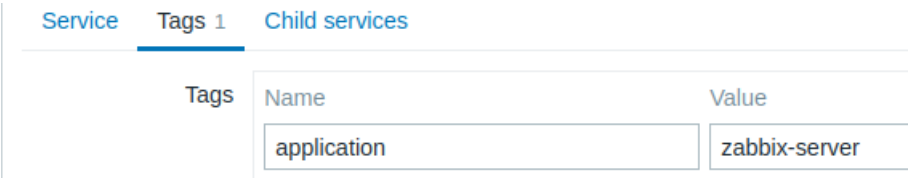
Configure additional rule:



Zabbix cluster will have two child services - one for each of the HA nodes. If both HA nodes have problems of at least *Warning* status, parent service status should be set to *Disaster*. To achieve this, additional rule should be configured as:

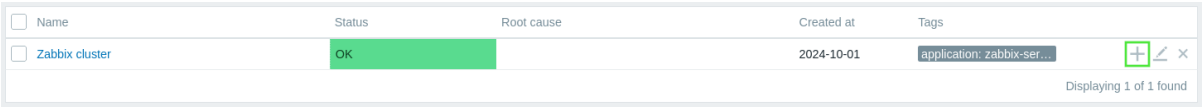
- Set status to: Disaster
- Condition: If at least `N` child services have `Status` status or above
- N: 2
- Status: Warning

Switch to the *Tags* tab and add a tag `application:zabbix-server`. This tag will be used later for service actions and SLA reports.



Save the new service.

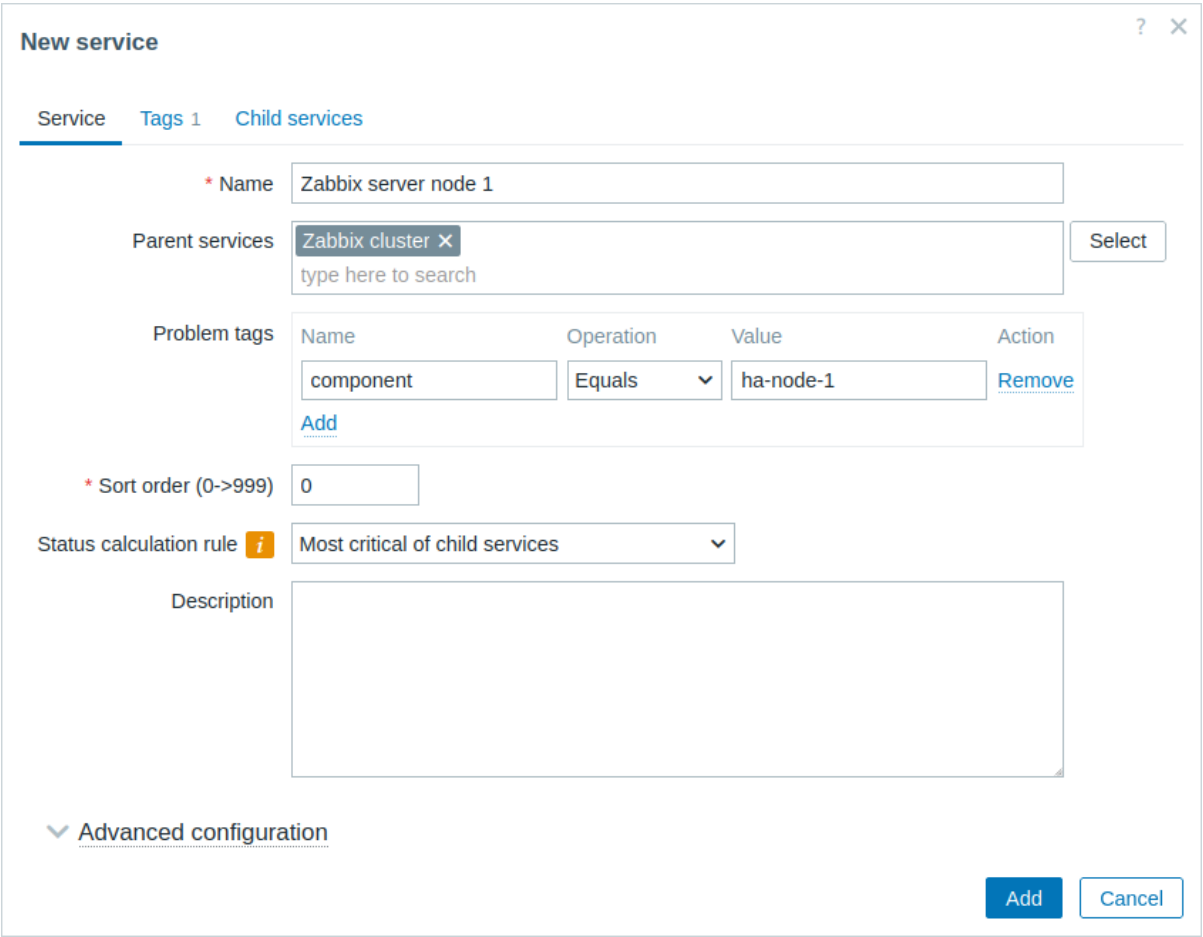
To add a child service, click the plus icon next to the Zabbix cluster service (the icon is visible only in *Edit* mode).



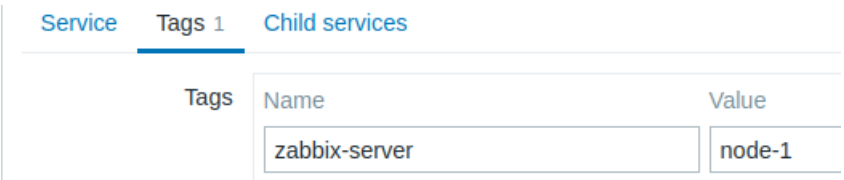
In the service configuration window, enter name *Zabbix server node 1*. Note that the Parent services parameter is already pre-filled with *Zabbix cluster*.

Availability of this service is affected by problems on the host *HA node 1*, marked with `component:ha-node-1` problem tag. In the Problem tags parameter, enter:

- Name: component
- Operation: Equals
- Value: ha-node-1



Switch to the *Tags* tab and add a service tag: `zabbix-server:node-1`. This tag will be used later for service actions and SLA reports.



Save the new service.

Create another child service of Zabbix cluster with name ”Zabbix server node 2”.

Set the Problem tags as:

- Name: component
- Operation: Equals
- Value: ha-node-2

Switch to the *Tags* tab and add a service tag: `zabbix-server:node-2`.

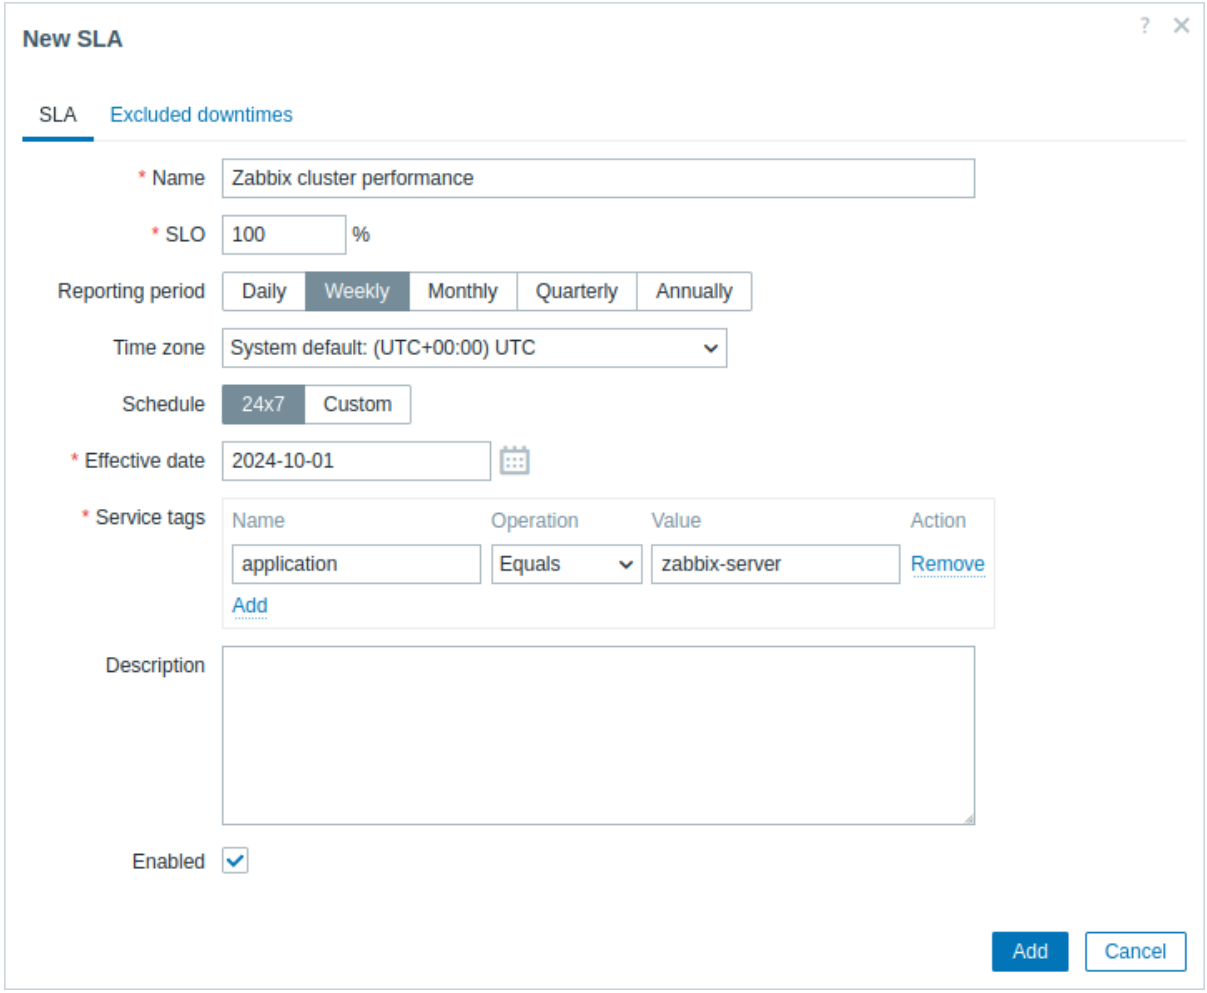
Save the new service.

SLA In this example, expected Zabbix cluster performance is 100% excluding semi-annual one hour maintenance period.

First, you need to add a new service level agreement:

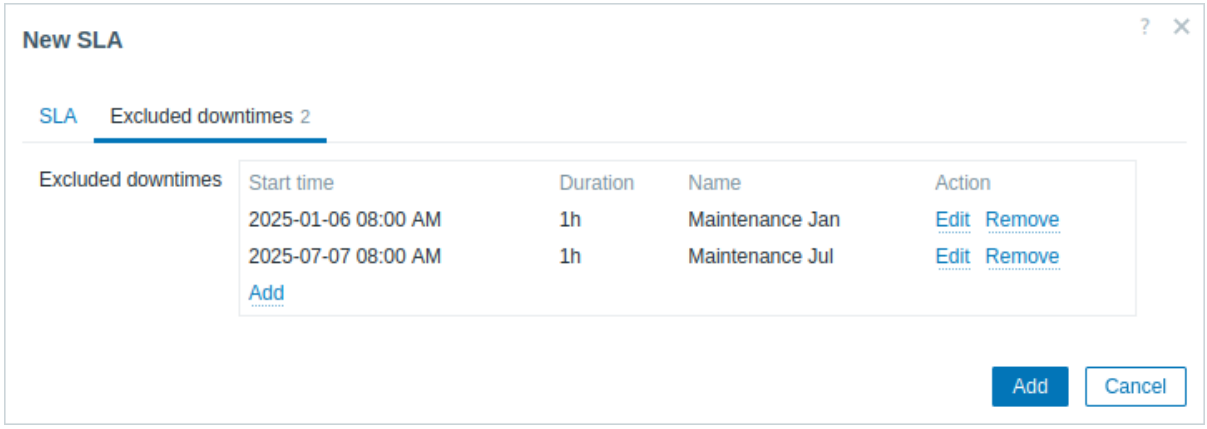
1. Go to *Services* > *SLA*.
2. Click *Create SLA*.
3. Enter the *Zabbix cluster performance* name and set the SLO to 100%. The service Zabbix cluster has a service tag `application:zabbix-server`.
4. To use this SLA for measuring performance of Zabbix cluster, in the *Service tags* parameter, specify:

- Name: application
- Operation: Equals
- Value: zabbix-server



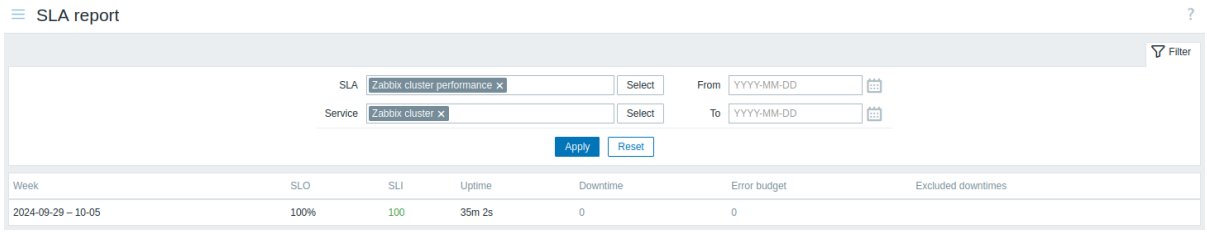
In a real-life setup, you can also update desired reporting period, time zone and start date or change the schedule from 24/7 to custom. For this example, the default settings are sufficient.

5. Switch to the *Excluded downtimes* tab and add downtimes for scheduled maintenance periods to exclude these periods from SLA calculation.
6. In the *Excluded downtimes* section, click the *Add* link, enter the downtime name, planned start time, and duration.

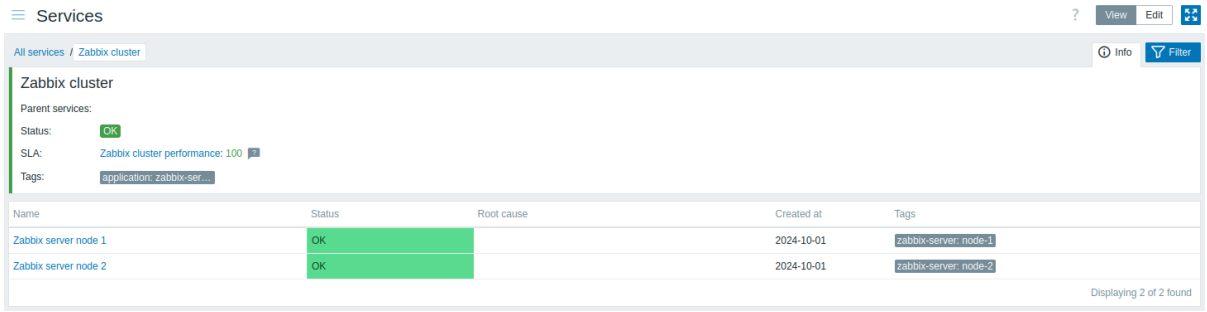


7. Click *Add* to save the new SLA.

Switch to the SLA reports section to view the SLA report for Zabbix cluster.



The SLA info can also be checked in the *Services* section.



Service monitoring upgrade

Overview In Zabbix 6.0, service monitoring functionality has been reworked significantly (see What's new in Zabbix 6.0.0 for the list of changes).

This page describes how services and SLAs, defined in earlier Zabbix versions, are changed during an upgrade to Zabbix 6.0 or newer.

Services In older Zabbix versions, services had two types of dependencies: soft and hard. After an upgrade, all dependencies will become equal.

If a service "Child service" has been previously linked to "Parent service 1" via hard dependency and additionally "Parent service 2" via soft dependency, after an upgrade the "Child service" will have two parent services "Parent service 1" and "Parent service 2".

Trigger-based mapping between problems and services has been replaced by tag-based mapping. In Zabbix 6.0 and newer, service configuration form has a new parameter *Problem tags*, which allows specifying one or multiple tag name and value pairs for problem matching. Triggers that have been linked to a service will get a new tag `ServiceLink : <trigger ID>:<trigger name>` (tag value will be truncated to 32 characters). Linked services will get `ServiceLink` problem tag with the same value.

Status calculation rules

The 'Status calculation algorithm' will be upgraded using the following rules:

- Do not calculate → Set status to OK
- Problem, if at least one child has a problem → Most critical of child services

- Problem, if all children have problems → Most critical if all children have problems

SLAs Previously, SLA targets had to be defined for each service separately. Since Zabbix 6.0, SLA has become a separate entity, which contains information about service schedule, expected service level objective (SLO) and downtime periods to exclude from the calculation. Once configured, an SLA can be assigned to multiple services through service tags.

During an upgrade:

- Identical SLAs defined for each service will be grouped and one SLA per each group will be created.
- Each affected service will get a special tag `SLA:<ID>` and the same tag will be specified in the *Service tags* parameter of the corresponding SLA.
- Service creation time, a new metric in SLA reports, will be set to 01/01/2000 00:00 for existing services.

Encryption and security

Encryption Zabbix supports encrypted communications between Zabbix components using Transport Layer Security (TLS) protocol v.1.2 and 1.3 (depending on the crypto library). Certificate-based and pre-shared key-based encryption is supported.

Encryption can be configured for connections:

- Between Zabbix server, Zabbix proxy, Zabbix agent, Zabbix web service, zabbix_sender and zabbix_get utilities
- To Zabbix database from Zabbix frontend and server/proxy
- Between Zabbix frontend and Zabbix server

Encryption is optional and configurable for individual components:

- Some proxies and agents can be configured to use certificate-based encryption with the server, while others can use pre-shared key-based encryption, and yet others continue with unencrypted communications (as before).
- Server (proxy) can use different encryption configurations for different hosts.

Zabbix daemon programs use one listening port for encrypted and unencrypted incoming connections. Adding an encryption does not require opening new ports on firewalls.

Limitations

- Private keys are stored in plain text in files readable by Zabbix components during startup.
- Pre-shared keys are entered in Zabbix frontend and stored in Zabbix database in plain text.
- Built-in encryption does not protect communications between the web server running Zabbix frontend and user web browser.
- Currently each encrypted connection opens with a full TLS handshake, no session caching and tickets are implemented.
- Adding encryption increases the time for item checks and actions, depending on network latency:
  - For example, if packet delay is 100ms then opening a TCP connection and sending unencrypted request takes around 200ms. With encryption about 1000 ms are added for establishing the TLS connection.
  - Timeouts may need to be increased; otherwise, some items and actions running remote scripts on agents may work with unencrypted connections, but fail with timeout with encrypted.
- Encryption is not supported by network discovery. Zabbix agent checks performed by network discovery will be unencrypted and if Zabbix agent is configured to reject unencrypted connections such checks will not succeed.

Compiling Zabbix with encryption support To support encryption Zabbix must be compiled and linked with one of the supported crypto libraries:

- GnuTLS - from version 3.1.18
- OpenSSL - versions 1.0.1, 1.0.2, 1.1.0, 1.1.1, 3.0.x - 3.5.x
- LibreSSL - tested with versions 2.7.4, 2.8.2:
  - LibreSSL 2.6.x is not supported
  - LibreSSL is supported as a compatible replacement of OpenSSL; the new `tls_*()` LibreSSL-specific API functions are not used. Zabbix components compiled with LibreSSL will not be able to use PSK, only certificates can be used.

> **Note:**
> You can find out more about setting up SSL for Zabbix frontend by referring to these best practices.

The library is selected by specifying the respective option to "configure" script:

- `--with-gnutls[=DIR]`
- `--with-openssl[=DIR]` (also used for LibreSSL)

For example, to configure the sources for server and agent with *OpenSSL* you may use something like:

```
./configure --enable-server --enable-agent --with-mysql --enable-ipv6 --with-net-snmp --with-libcurl --wit
```

Different Zabbix components may be compiled with different crypto libraries (e.g., a server with *OpenSSL*, an agent with *GnuTLS*).

Attention:
If you plan to use pre-shared keys (PSK), consider using *GnuTLS* or *OpenSSL 1.1.0* (or newer) libraries in Zabbix components using PSKs. *GnuTLS* and *OpenSSL 1.1.0* libraries support PSK ciphersuites with Perfect Forward Secrecy. Older versions of the *OpenSSL* library (1.0.1, 1.0.2c) also support PSKs, but available PSK ciphersuites do not provide Perfect Forward Secrecy.

Connection encryption management Connections in Zabbix can use:

- no encryption (default)
- RSA certificate-based encryption
- PSK-based encryption

There are two important parameters used to specify encryption between Zabbix components:

- TLSConnect - specifies what encryption to use for outgoing connections (unencrypted, PSK or certificate)
- TLSAccept - specifies what types of connections are allowed for incoming connections (unencrypted, PSK or certificate). One or more values can be specified.

`TLSConnect` is used in the configuration files for Zabbix proxy (in active mode, specifies only connections to server) and Zabbix agent (for active checks). In Zabbix frontend the `TLSConnect` equivalent is the *Connections to host* field in *Data collection* > *Hosts* > <some host> > *Encryption* tab and the *Connections to proxy* field in *Administration* > *Proxies* > <some proxy> > *Encryption* tab. If the configured encryption type for connection fails, no other encryption types will be tried.

`TLSAccept` is used in the configuration files for Zabbix proxy (in passive mode, specifies only connections from server) and Zabbix agent (for passive checks). In Zabbix frontend the TLSAccept equivalent is the *Connections from host* field in *Data collection* > *Hosts* > <some host> > *Encryption* tab and the *Connections from proxy* field in *Administration* > *Proxies* > <some proxy> > *Encryption* tab.

Normally you configure only one type of encryption for incoming connections. But you may want to switch the encryption type, e.g., from unencrypted to certificate-based with minimum downtime and rollback possibility. To achieve this:

- Set `TLSAccept=unencrypted,cert` in the agent configuration file and restart Zabbix agent.
- Test connection with `zabbix_get` to the agent using certificate. If it works, you can reconfigure encryption for that agent in Zabbix frontend in the *Data collection* > *Hosts* > <some host> > *Encryption* tab by setting *Connections to host* to *Certificate*.
- When server configuration cache gets updated (and proxy configuration is updated if the host is monitored by proxy) then connections to that agent will be encrypted.
- If everything works as expected you can set `TLSAccept=cert` in the agent configuration file and restart Zabbix agent. Now the agent will be accepting only encrypted certificate-based connections. Unencrypted and PSK-based connections will be rejected.

In a similar way it works on server and proxy. If in Zabbix frontend in host configuration *Connections from host* is set to "Certificate" then only certificate-based encrypted connections will be accepted from the agent (active checks) and zabbix_sender (trapper items).

Most likely you will configure incoming and outgoing connections to use the same encryption type or no encryption at all. But technically it is possible to configure it asymmetrically, e.g., certificate-based encryption for incoming and PSK-based for outgoing connections.

Encryption configuration for each host is displayed in the Zabbix frontend, in *Data collection* > *Hosts* in the *Agent encryption* column. For example:

Example	Connections to host	Allowed connections from host	Rejected connections from host
NONE	Unencrypted	Unencrypted	Encrypted, certificate and PSK-based encrypted
CERT NONE PSK CERT	Encrypted, certificate-based	Encrypted, certificate-based	Unencrypted and PSK-based encrypted
PSK NONE PSK CERT	Encrypted, PSK-based	Encrypted, PSK-based	Unencrypted and certificate-based encrypted
PSK NONE PSK CERT	Encrypted, PSK-based	Unencrypted and PSK-based encrypted	Certificate-based encrypted

Example	Connections to host	Allowed connections from host	Rejected connections from host
CERT NONE PSK CERT	Encrypted, certificate-based	Unencrypted, PSK or certificate-based encrypted	-

Attention:
Connections are unencrypted by default. Encryption must be configured for each host and proxy individually.

zabbix_get and zabbix_sender with encryption See zabbix_get and zabbix_sender manpages for using them with encryption.

Ciphersuites Ciphersuites by default are configured internally during Zabbix startup.

Also user-configured ciphersuites are supported for GnuTLS and OpenSSL. Users may configure ciphersuites according to their security policies. Using this feature is optional (built-in default ciphersuites still work).

For crypto libraries compiled with default settings Zabbix built-in rules typically result in the following ciphersuites (in order from higher to lower priority):

Library	Certificate ciphersuites	PSK ciphersuites
GnuTLS 3.1.18	TLS_ECDHE_RSA_AES_128_GCM_SHA256 TLS_ECDHE_RSA_AES_128_CBC_SHA256 TLS_ECDHE_RSA_AES_128_CBC_SHA1 TLS_RSA_AES_128_GCM_SHA256 TLS_RSA_AES_128_CBC_SHA256 TLS_RSA_AES_128_CBC_SHA1	TLS_ECDHE_PSK_AES_128_CBC_SHA256 TLS_ECDHE_PSK_AES_128_CBC_SHA1 TLS_PSK_AES_128_GCM_SHA256 TLS_PSK_AES_128_CBC_SHA256 TLS_PSK_AES_128_CBC_SHA1
OpenSSL 1.0.2c	ECDHE-RSA-AES128-GCM-SHA256 ECDHE-RSA-AES128-SHA256 ECDHE-RSA-AES128-SHA AES128-GCM-SHA256 AES128-SHA256 AES128-SHA	PSK-AES128-CBC-SHA
OpenSSL 1.1.0	ECDHE-RSA-AES128-GCM-SHA256 ECDHE-RSA-AES128-SHA256 ECDHE-RSA-AES128-SHA AES128-GCM-SHA256 AES128-CCM8 AES128-CCM AES128-SHA256 AES128-SHA	ECDHE-PSK-AES128-CBC-SHA256 ECDHE-PSK-AES128-CBC-SHA PSK-AES128-GCM-SHA256 PSK-AES128-CCM8 PSK-AES128-CCM PSK-AES128-CBC-SHA256 PSK-AES128-CBC-SHA
OpenSSL 1.1.1d	TLS_AES_256_GCM_SHA384 TLS_CHACHA20_POLY1305_SHA256 TLS_AES_128_GCM_SHA256 ECDHE-RSA-AES128-GCM-SHA256 ECDHE-RSA-AES128-SHA256 ECDHE-RSA-AES128-SHA AES128-GCM-SHA256 AES128-CCM8 AES128-CCM AES128-SHA256 AES128-SHA	TLS_CHACHA20_POLY1305_SHA256 TLS_AES_128_GCM_SHA256 ECDHE-PSK-AES128-CBC-SHA256 ECDHE-PSK-AES128-CBC-SHA PSK-AES128-GCM-SHA256 PSK-AES128-CCM8 PSK-AES128-CCM PSK-AES128-CBC-SHA256 PSK-AES128-CBC-SHA

User-configured ciphersuites The built-in ciphersuite selection criteria can be overridden with user-configured ciphersuites.

Attention:
User-configured ciphersuites is a feature intended for advanced users who understand TLS ciphersuites, their security and consequences of mistakes, and who are comfortable with TLS troubleshooting.

The built-in ciphersuite selection criteria can be overridden using the following parameters:

Override scope	Parameter	Value	Description
Ciphersuite selection for certificates	TLSCipherCert13	Valid OpenSSL 1.1.1 cipher strings for TLS 1.3 protocol (their values are passed to the OpenSSL function `SSL_CTX_set_ciphersuites()`).	Certificate-based ciphersuite selection criteria for TLS 1.3 Only OpenSSL 1.1.1 or newer.
	TLSCipherCert	Valid OpenSSL cipher strings for TLS 1.2 or valid GnuTLS priority strings. Their values are passed to the `SSL_CTX_set_cipher_list()` or `gnutls_priority_init()` functions, respectively.	Certificate-based ciphersuite selection criteria for TLS 1.2/1.3 (GnuTLS), TLS 1.2 (OpenSSL)
Ciphersuite selection for PSK	TLSCipherPSK13	Valid OpenSSL 1.1.1 cipher strings for TLS 1.3 protocol (their values are passed to the OpenSSL function `SSL_CTX_set_ciphersuites()`).	PSK-based ciphersuite selection criteria for TLS 1.3 Only OpenSSL 1.1.1 or newer.
	TLSCipherPSK	Valid OpenSSL cipher strings for TLS 1.2 or valid GnuTLS priority strings. Their values are passed to the `SSL_CTX_set_cipher_list()` or `gnutls_priority_init()` functions, respectively.	PSK-based ciphersuite selection criteria for TLS 1.2/1.3 (GnuTLS), TLS 1.2 (OpenSSL)
Combined ciphersuite list for certificate and PSK	TLSCipherAll13	Valid OpenSSL 1.1.1 cipher strings for TLS 1.3 protocol (their values are passed to the OpenSSL function `SSL_CTX_set_ciphersuites()`).	Ciphersuite selection criteria for TLS 1.3 Only OpenSSL 1.1.1 or newer.
	TLSCipherAll	Valid OpenSSL cipher strings for TLS 1.2 or valid GnuTLS priority strings. Their values are passed to the `SSL_CTX_set_cipher_list()` or `gnutls_priority_init()` functions, respectively.	Ciphersuite selection criteria for TLS 1.2/1.3 (GnuTLS), TLS 1.2 (OpenSSL)

To override the ciphersuite selection in zabbix_get and zabbix_sender utilities - use the command-line parameters:

- `--tls-cipher13`
- `--tls-cipher`

The new parameters are optional. If a parameter is not specified, the internal default value is used. If a parameter is defined it cannot be empty.

If the setting of a TLSCipher* value in the crypto library fails then the server, proxy or agent will not start and an error is logged.

It is important to understand when each parameter is applicable.

Outgoing connections

The simplest case is outgoing connections:

- For outgoing connections with certificate - use TLSCipherCert13 or TLSCipherCert
- For outgoing connections with PSK - use TLSCipherPSK13 or TLSCipherPSK
- In case of zabbix_get and zabbix_sender utilities the command-line parameters `--tls-cipher13` or `--tls-cipher` can be used (encryption is unambiguously specified with a `--tls-connect` parameter)

Incoming connections

It is a bit more complicated with incoming connections because rules are specific for components and configuration.

For Zabbix **agent**:

Agent connection setup	Cipher configuration
TLSConnect=cert	TLSCipherCert, TLSCipherCert13
TLSConnect=psk	TLSCipherPSK, TLSCipherPSK13
TLSAccept=cert	TLSCipherCert, TLSCipherCert13
TLSAccept=psk	TLSCipherPSK, TLSCipherPSK13

Agent connection setup	Cipher configuration
TLSAccept=cert,psk	TLSCipherAll, TLSCipherAll13

For Zabbix **server** and **proxy**:

Connection setup	Cipher configuration
Outgoing connections using PSK	TLSCipherPSK, TLSCipherPSK13
Incoming connections using certificates	TLSCipherAll, TLSCipherAll13
Incoming connections using PSK if server has no certificate	TLSCipherPSK, TLSCipherPSK13
Incoming connections using PSK if server has certificate	TLSCipherAll, TLSCipherAll13

Some pattern can be seen in the two tables above:

- TLSCipherAll and TLSCipherAll13 can be specified only if a combined list of certificate- **and** PSK-based ciphersuites is used. There are two cases when it takes place: server (proxy) with a configured certificate (PSK ciphersuites are always configured on server, proxy if crypto library supports PSK), agent configured to accept both certificate- and PSK-based incoming connections
- in other cases TLSCipherCert* and/or TLSCipherPSK* are sufficient

The following tables show the `TLSCipher*` built-in default values. They could be a good starting point for your own custom values.

Parameter	GnuTLS 3.6.12
TLSCipherCert	NONE:+VERS-TLS1.2:+ECDHE-RSA:+RSA:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+SHA1:+CURVE-ALL:+COMP-NULL:+SIGN-ALL:+CTYPE-X.509
TLSCipherPSK	NONE:+VERS-TLS1.2:+ECDHE-PSK:+PSK:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+SHA1:+CURVE-ALL:+COMP-NULL:+SIGN-ALL
TLSCipherAll	NONE:+VERS-TLS1.2:+ECDHE-RSA:+RSA:+ECDHE-PSK:+PSK:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+SHA1:+CURVE-ALL:+COMP-NULL:+SIGN-ALL:+CTYPE-X.509

Parameter	OpenSSL 1.1.1d [1]
TLSCipherCert13	
TLSCipherCert	EECDH+aRSA+AES128:RSA+aRSA+AES128
TLSCipherPSK13	TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256
TLSCipherPSK	kECDHEPSK+AES128:kPSK+AES128
TLSCipherAll13	
TLSCipherAll	EECDH+aRSA+AES128:RSA+aRSA+AES128:kECDHEPSK+AES128:kPSK+AES128

[1] Default values are different for older OpenSSL versions (1.0.1, 1.0.2, 1.1.0), for LibreSSL and if OpenSSL is compiled without PSK support.

Examples of user-configured ciphersuites

See below the following examples of user-configured ciphersuites:

- Testing cipher strings and allowing only PFS ciphersuites
- Switching from AES128 to AES256

Testing cipher strings and allowing only PFS ciphersuites

To see which ciphersuites have been selected you need to set 'DebugLevel=4' in the configuration file, or use the `-vv` option for zabbix_sender.

Some experimenting with `TLSCipher*` parameters might be necessary before you get the desired ciphersuites. It is inconvenient to restart Zabbix server, proxy or agent multiple times just to tweak `TLSCipher*` parameters. More convenient options are using zabbix_sender or the `openssl` command. Let's show both.

1. Using zabbix_sender.

Let's make a test configuration file, for example, `/home/zabbix/test.conf`, with the syntax of a `zabbix_agentd.conf` file:

```
Hostname=nonexisting
ServerActive=nonexisting

TLSConnect=cert
TLSCAFile=/home/zabbix/ca.crt
TLSCertFile=/home/zabbix/agent.crt
TLSKeyFile=/home/zabbix/agent.key
TLSPSKIdentity=nonexisting
TLSPSKFile=/home/zabbix/agent.psk
```

You need valid CA and agent certificates and PSK for this example. Adjust certificate and PSK file paths and names for your environment.

If you are not using certificates, but only PSK, you can make a simpler test file:

```
Hostname=nonexisting
ServerActive=nonexisting

TLSConnect=psk
TLSPSKIdentity=nonexisting
TLSPSKFile=/home/zabbix/agentd.psk
```

The selected ciphersuites can be seen by running zabbix_sender (example compiled with OpenSSL 1.1.d):

```
$ zabbix_sender -vv -c /home/zabbix/test.conf -k nonexisting_item -o 1 2>&1 | grep ciphersuites
zabbix_sender [41271]: DEBUG: zbx_tls_init_child() certificate ciphersuites: TLS_AES_256_GCM_SHA384 TLS_
zabbix_sender [41271]: DEBUG: zbx_tls_init_child() PSK ciphersuites: TLS_CHACHA20_POLY1305_SHA256 TLS_AE
zabbix_sender [41271]: DEBUG: zbx_tls_init_child() certificate and PSK ciphersuites: TLS_AES_256_GCM_SHA
```

Here you see the ciphersuites selected by default. These default values are chosen to ensure interoperability with Zabbix agents running on systems with older OpenSSL versions (from 1.0.1).

With newer systems, you can choose to tighten security by allowing only a few ciphersuites, e.g., only ciphersuites with PFS (Perfect Forward Secrecy). Let's configure the `TLSCipher*` parameters to allow only cipher suites that support PFS.

Attention:
The result will not be interoperable with systems using OpenSSL 1.0.1 and 1.0.2, if PSK is used. Certificate-based encryption should work.

Add two lines to the `test.conf` configuration file:

```
TLSCipherCert=EECDH+aRSA+AES128
TLSCipherPSK=kECDHEPSK+AES128
```

and test again:

```
$ zabbix_sender -vv -c /home/zabbix/test.conf -k nonexisting_item -o 1 2>&1 | grep ciphersuites
zabbix_sender [42892]: DEBUG: zbx_tls_init_child() certificate ciphersuites: TLS_AES_256_GCM_SHA384 TLS_
zabbix_sender [42892]: DEBUG: zbx_tls_init_child() PSK ciphersuites: TLS_CHACHA20_POLY1305_SHA256 TLS_AE
zabbix_sender [42892]: DEBUG: zbx_tls_init_child() certificate and PSK ciphersuites: TLS_AES_256_GCM_SHA
```

The "certificate ciphersuites" and "PSK ciphersuites" lists have changed - they are shorter than before, only containing TLS 1.3 ciphersuites and TLS 1.2 ECDHE-* ciphersuites as expected.

2. TLSCipherAll and TLSCipherAll13 cannot be tested with zabbix_sender; they do not affect "certificate and PSK ciphersuites" value shown in the example above. To tweak TLSCipherAll and TLSCipherAll13 you need to experiment with the agent, proxy or server.

So, to allow only PFS ciphersuites you may need to add up to three parameters

```
TLSCipherCert=EECDH+aRSA+AES128
TLSCipherPSK=kECDHEPSK+AES128
TLSCipherAll=EECDH+aRSA+AES128:kECDHEPSK+AES128
```

to `zabbix_agentd.conf`, `zabbix_proxy.conf` and `zabbix_server.conf` if each of them has a configured certificate and agent has also PSK.

If your Zabbix environment uses only PSK-based encryption and no certificates, then only one:

```
TLSCipherPSK=kECDHEPSK+AES128
```

Now that you understand how it works you can test the ciphersuite selection even outside of Zabbix, with the `openssl` command. Let's test all three `TLSCipher*` parameter values:

```
$ openssl ciphers EECDH+aRSA+AES128 | sed 's/:/ /g'
TLS_AES_256_GCM_SHA384 TLS_CHACHA20_POLY1305_SHA256 TLS_AES_128_GCM_SHA256 ECDHE-RSA-AES128-GCM-SHA256 E
$ openssl ciphers kECDHEPSK+AES128 | sed 's/:/ /g'
TLS_AES_256_GCM_SHA384 TLS_CHACHA20_POLY1305_SHA256 TLS_AES_128_GCM_SHA256 ECDHE-PSK-AES128-CBC-SHA256 E
$ openssl ciphers EECDH+aRSA+AES128:kECDHEPSK+AES128 | sed 's/:/ /g'
TLS_AES_256_GCM_SHA384 TLS_CHACHA20_POLY1305_SHA256 TLS_AES_128_GCM_SHA256 ECDHE-RSA-AES128-GCM-SHA256 E
```

You may prefer `openssl ciphers` with option `-V` for a more verbose output:

```
$ openssl ciphers -V EECDH+aRSA+AES128:kECDHEPSK+AES128
          0x13,0x02 - TLS_AES_256_GCM_SHA384  TLSv1.3 Kx=any      Au=any  Enc=AESGCM(256) Mac=AEAD
          0x13,0x03 - TLS_CHACHA20_POLY1305_SHA256 TLSv1.3 Kx=any      Au=any  Enc=CHACHA20/POLY1305(256
          0x13,0x01 - TLS_AES_128_GCM_SHA256  TLSv1.3 Kx=any      Au=any  Enc=AESGCM(128) Mac=AEAD
          0xC0,0x2F - ECDHE-RSA-AES128-GCM-SHA256 TLSv1.2 Kx=ECDH     Au=RSA  Enc=AESGCM(128) Mac=AEAD
          0xC0,0x27 - ECDHE-RSA-AES128-SHA256 TLSv1.2 Kx=ECDH     Au=RSA  Enc=AES(128)  Mac=SHA256
          0xC0,0x13 - ECDHE-RSA-AES128-SHA    TLSv1 Kx=ECDH     Au=RSA  Enc=AES(128)  Mac=SHA1
          0xC0,0x37 - ECDHE-PSK-AES128-CBC-SHA256 TLSv1 Kx=ECDHEPSK Au=PSK  Enc=AES(128)  Mac=SHA256
          0xC0,0x35 - ECDHE-PSK-AES128-CBC-SHA TLSv1 Kx=ECDHEPSK Au=PSK  Enc=AES(128)  Mac=SHA1
```

Similarly, you can test the priority strings for GnuTLS:

```
$ gnutls-cli -l --priority=NONE:+VERS-TLS1.2:+ECDHE-RSA:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+CURVE-A
Cipher suites for NONE:+VERS-TLS1.2:+ECDHE-RSA:+AES-128-GCM:+AES-128-CBC:+AEAD:+SHA256:+CURVE-ALL:+COMP-
TLS_ECDHE_RSA_AES_128_GCM_SHA256                        0xc0, 0x2f      TLS1.2
TLS_ECDHE_RSA_AES_128_CBC_SHA256                        0xc0, 0x27      TLS1.2

Protocols: VERS-TLS1.2
Ciphers: AES-128-GCM, AES-128-CBC
MACs: AEAD, SHA256
Key Exchange Algorithms: ECDHE-RSA
Groups: GROUP-SECP256R1, GROUP-SECP384R1, GROUP-SECP521R1, GROUP-X25519, GROUP-X448, GROUP-FFDHE2048, GR
PK-signatures: SIGN-RSA-SHA256, SIGN-RSA-PSS-SHA256, SIGN-RSA-PSS-RSAE-SHA256, SIGN-ECDSA-SHA256, SIGN-E
```

Switching from AES128 to AES256

Zabbix uses AES128 as the built-in default for data. Let's assume you are using certificates and want to switch to AES256, on OpenSSL 1.1.1.

This can be achieved by adding the respective parameters in `zabbix_server.conf`:

```
TLSCAFile=/home/zabbix/ca.crt
TLSCertFile=/home/zabbix/server.crt
TLSKeyFile=/home/zabbix/server.key
TLSCipherCert13=TLS_AES_256_GCM_SHA384
TLSCipherCert=EECDH+aRSA+AES256:-SHA1:-SHA384
TLSCipherPSK13=TLS_CHACHA20_POLY1305_SHA256
TLSCipherPSK=kECDHEPSK+AES256:-SHA1
TLSCipherAll13=TLS_AES_256_GCM_SHA384
TLSCipherAll=EECDH+aRSA+AES256:-SHA1:-SHA384
```

Attention:
Although only certificate-related ciphersuites will be used, `TLSCipherPSK*` parameters are defined as well to avoid their default values which include less secure ciphers for wider interoperability. PSK ciphersuites cannot be completely disabled on server/proxy.

And in `zabbix_agentd.conf`:

```
TLSConnect=cert
TLSAccept=cert
TLSCAFile=/home/zabbix/ca.crt
TLSCertFile=/home/zabbix/agent.crt
TLSKeyFile=/home/zabbix/agent.key
TLSCipherCert13=TLS_AES_256_GCM_SHA384
TLSCipherCert=EECDH+aRSA+AES256:-SHA1:-SHA384
```

Troubleshooting encryption issues General recommendations:

- Start with understanding which component acts as a TLS client and which one acts as a TLS server in problem case.
 Zabbix server, proxies and agents, depending on interaction between them, all can work as TLS servers and clients.
 For example, Zabbix server connecting to agent for a passive check, acts as a TLS client. The agent is in role of TLS server.
 Zabbix agent, requesting a list of active checks from proxy, acts as a TLS client. The proxy is in role of TLS server.
 `zabbix_get` and `zabbix_sender` utilities always act as TLS clients.
- Zabbix uses mutual authentication.
 Each side verifies its peer and may refuse connection.
 For example, Zabbix server connecting to agent can close connection immediately if agent's certificate is invalid. And vice versa - Zabbix agent accepting a connection from server can close connection if server is not trusted by agent.
- Examine logfiles in both sides - in TLS client and TLS server.
 The side which refuses connection may log a precise reason why it was refused. Other side often reports rather general error (e.g., "Connection closed by peer", "connection was non-properly terminated").
- Sometimes misconfigured encryption results in confusing error messages in no way pointing to real cause.
 In subsections below we try to provide a (far from exhaustive) collection of messages and possible causes which could help in troubleshooting.
 Please note that different crypto toolkits (OpenSSL, GnuTLS) often produce different error messages in same problem situations.
 Sometimes error messages depend even on particular combination of crypto toolkits on both sides.

Troubleshooting connection type or permission problems Server is configured to connect with PSK to agent but agent accepts only unencrypted connections

In server or proxy log (with *GnuTLS* 3.3.16)

```
Get value from agent failed: zbx_tls_connect(): gnutls_handshake() failed: \
    -110 The TLS connection was non-properly terminated.
```

In server or proxy log (with *OpenSSL* 1.0.2c)

```
Get value from agent failed: TCP connection successful, cannot establish TLS to [[127.0.0.1]:10050]: \
    Connection closed by peer. Check allowed connection types and access rights
```

One side connects with certificate but other side accepts only PSK or vice versa

In any log (with *GnuTLS*):

```
failed to accept an incoming connection: from 127.0.0.1: zbx_tls_accept(): gnutls_handshake() failed:\
    -21 Could not negotiate a supported cipher suite.
```

In any log (with *OpenSSL* 1.0.2c):

```
failed to accept an incoming connection: from 127.0.0.1: TLS handshake returned error code 1:\
    file .\ssl\s3_srvr.c line 1411: error:1408A0C1:SSL routines:ssl3_get_client_hello:no shared cipher:\
    TLS write fatal alert "handshake failure"
```

Attempting to use Zabbix sender compiled with TLS support to send data to Zabbix server/proxy compiled without TLS In connecting-side log:

Linux:

```
...In zbx_tls_init_child()
...OpenSSL library (version OpenSSL 1.1.1  11 Sep 2018) initialized
...
...In zbx_tls_connect(): psk_identity:"PSK test sender"
...End of zbx_tls_connect():FAIL error:'connection closed by peer'
...send value error: TCP successful, cannot establish TLS to [[localhost]:10051]: connection closed by pee
```

Windows:

```
...OpenSSL library (version OpenSSL 1.1.1a  20 Nov 2018) initialized
...
...In zbx_tls_connect(): psk_identity:"PSK test sender"
...zbx_psk_client_cb() requested PSK identity "PSK test sender"
...End of zbx_tls_connect():FAIL error:'SSL_connect() I/O error: [0x00000000] The operation completed succ
...send value error: TCP successful, cannot establish TLS to [[192.168.1.2]:10051]: SSL_connect() I/O erro
```

In accepting-side log:

```
...failed to accept an incoming connection: from 127.0.0.1: support for TLS was not compiled in
```

One side connects with PSK but other side uses LibreSSL or has been compiled without encryption support LibreSSL does not support PSK.

In connecting-side log:

```
...TCP successful, cannot establish TLS to [[192.168.1.2]:10050]: SSL_connect() I/O error: [0] Success
```

In accepting-side log:

```
...failed to accept an incoming connection: from 192.168.1.2: support for PSK was not compiled in
```

In Zabbix frontend:

```
Get value from agent failed: TCP successful, cannot establish TLS to [[192.168.1.2]:10050]: SSL_connect()
```

One side connects with PSK but other side uses OpenSSL with PSK support disabled In connecting-side log:

```
...TCP successful, cannot establish TLS to [[192.168.1.2]:10050]: SSL_connect() set result code to SSL_ERR
```

In accepting-side log:

```
...failed to accept an incoming connection: from 192.168.1.2: TLS handshake set result code to 1: file ssl
```

Encryption using certificates

Overview Zabbix can use RSA certificates in PEM format, signed by a public or an in-house certificate authority (CA).

Certificate verification is performed against a pre-configured CA certificate. Optionally, Certificate Revocation Lists (CRL) can be used.

Each Zabbix component can have only one certificate configured.

For more information on setting up and operating an internal CA, generating and signing certificate requests, and revoking certificates, refer to tutorials such as the OpenSSL PKI Tutorial v2.0.

Carefully consider and test your certificate extensions. For more details, see Limitations on using X.509 v3 certificate extensions.

Certificate configuration parameters The following configuration parameters are supported for setting up certificates on Zabbix components.

Parameter	Mandatory	Description
TLSCAFile	yes	Full pathname of a file containing the top-level CA(s) certificates for peer certificate verification. If using a certificate chain with multiple members, order the certificates with lower level CA(s) certificates first, followed by higher level CA(s) certificates. Certificates from multiple CAs can be included in a single file.
TLSCRLFile	no	Full pathname of a file containing Certificate Revocation Lists (CRL).
TLSCertFile	yes	Full pathname of a file containing the certificate. If using a certificate chain with multiple members, order the certificates with the server, proxy, or agent certificate first, followed by lower level CA(s) certificates, and concluded by higher level CA(s) certificates.
TLSKeyFile	yes	Full pathname of a file containing the private key. Ensure that this file is readable only by the Zabbix user by setting appropriate access rights.
TLSServerCertIssuer	no	Allowed server certificate issuer.
TLSServerCertSubject	no	Allowed server certificate subject.

Configuration examples After setting up the necessary certificates, configure Zabbix components to use certificate-based encryption.

Below are detailed steps for configuring:

- Zabbix server
- Zabbix proxy
- Zabbix agent

Zabbix server

1. Prepare the CA certificate file.

In order to verify peer certificates, Zabbix server must have access to the file containing the top-level, self-signed root CA certificates. For example, if certificates from two independent root CAs are needed, place them into a file at `/home/zabbix/zabbix_ca_file.crt`

```
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 1 (0x1)
    Signature Algorithm: sha1WithRSAEncryption
        Issuer: DC=com, DC=zabbix, O=Zabbix SIA, OU=Development group, CN=Root1 CA
            ...
        Subject: DC=com, DC=zabbix, O=Zabbix SIA, OU=Development group, CN=Root1 CA
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (2048 bit)
            ...
        X509v3 extensions:
            X509v3 Key Usage: critical
                Certificate Sign, CRL Sign
            X509v3 Basic Constraints: critical
                CA:TRUE
            ...
-----BEGIN CERTIFICATE-----
MIID2jCCAsKgAwIBAgIBATANBgkqhkiG9w0BAQUFADB+MRMwEQYKCZImiZPyLGQB
....
9wEzdN8uTrqoyU78gi12npLj08LegRKjb5hFTVm0
-----END CERTIFICATE-----
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 1 (0x1)
    Signature Algorithm: sha1WithRSAEncryption
        Issuer: DC=com, DC=zabbix, O=Zabbix SIA, OU=Development group, CN=Root2 CA
            ...
        Subject: DC=com, DC=zabbix, O=Zabbix SIA, OU=Development group, CN=Root2 CA
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (2048 bit)
            ....
        X509v3 extensions:
            X509v3 Key Usage: critical
                Certificate Sign, CRL Sign
            X509v3 Basic Constraints: critical
                CA:TRUE
            ....
-----BEGIN CERTIFICATE-----
MIID3DCCAsSgAwIBAgIBATANBgkqhkiG9w0BAQUFADB/MRMwEQYKCZImiZPyLGQB
...
vdGNYoSfvu41GQAR5Vj5FnRJRzv5XQOZ3B6894GY1zY=
-----END CERTIFICATE-----
```

2. Place the Zabbix server certificate/certificate chain into a file, for example, at `/home/zabbix/zabbix_server.crt`. The first certificate is the Zabbix server certificate, followed by the intermediate CA certificate:

```
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 1 (0x1)
    Signature Algorithm: sha1WithRSAEncryption
        Issuer: DC=com, DC=zabbix, O=Zabbix SIA, OU=Development group, CN=Signing CA
        ...
        Subject: DC=com, DC=zabbix, O=Zabbix SIA, OU=Development group, CN=Zabbix server
```

```
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (2048 bit)
                ...
        X509v3 extensions:
            X509v3 Key Usage: critical
                Digital Signature, Key Encipherment
            X509v3 Basic Constraints:
                CA:FALSE
            ...
-----BEGIN CERTIFICATE-----
MIIECDCCAvCgAwIBAgIBATANBgkqhkiG9w0BAQUFADCBgTETMBEGCgmSJomT8ixk
...
h02u1GHiy46GI+xfR3LsPwFKlkTaaLaL/6aaoQ==
-----END CERTIFICATE-----
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 2 (0x2)
    Signature Algorithm: sha1WithRSAEncryption
        Issuer: DC=com, DC=zabbix, O=Zabbix SIA, OU=Development group, CN=Root1 CA
        ...
        Subject: DC=com, DC=zabbix, O=Zabbix SIA, OU=Development group, CN=Signing CA
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (2048 bit)
            ...
        X509v3 extensions:
            X509v3 Key Usage: critical
                Certificate Sign, CRL Sign
            X509v3 Basic Constraints: critical
                CA:TRUE, pathlen:0
        ...
-----BEGIN CERTIFICATE-----
MIID4TCCAsmgAwIBAgIBAjANBgkqhkiG9w0BAQUFADB+MRMwEQYKCZImiZPyLGQB
...
dyCeWnvL7u5sd6ffo8iRny0QzbHKmQt/wUtcVIvWXdMIFJM0Hw==
-----END CERTIFICATE-----
```

Note:
Use only the attributes mentioned above for both client and server certificates to avoid affecting the certificate verification process. For example, OpenSSL might fail to establish an encrypted connection if *X509v3 Subject Alternative Name* or *Netscape Cert Type* extensions are used. For more information, see Limitations on using X.509 v3 certificate extensions.

3. Place the Zabbix server private key into a file, for example, at `/home/zabbix/zabbix_server.key`:

```
-----BEGIN PRIVATE KEY-----
MIIEwAIBADANBgkqhkiG9w0BAQEFAASCBKowggSmAgEAAoIBAQC9tIXIJoVnNXDl
...
IJLkhbybBYEf47MLhffWa7XvZTY=
-----END PRIVATE KEY-----
```

4. Edit the TLS configuration parameters in the Zabbix server configuration file:

```
TLSCAFile=/home/zabbix/zabbix_ca_file.crt
TLSCertFile=/home/zabbix/zabbix_server.crt
TLSKeyFile=/home/zabbix/zabbix_server.key
```

Zabbix proxy

1. Prepare files with the top-level CA certificates, the Zabbix proxy certificate/certificate chain, and the private key as described in the Zabbix server section. Then, edit the `TLSCAFile`, `TLSCertFile`, and `TLSKeyFile` parameters in the Zabbix proxy configuration file accordingly.

2. Edit additional TLS parameters in the Zabbix proxy configuration file:

- For active proxy: `TLSConnect=cert`
- For passive proxy: `TLSAccept=cert`

Note:
To improve proxy security, you can also set the `TLSServerCertIssuer` and `TLSServerCertSubject` parameters. For more information, see Restricting allowed certificate issuer and subject.

TLS parameters in the final proxy configuration file may look as follows:

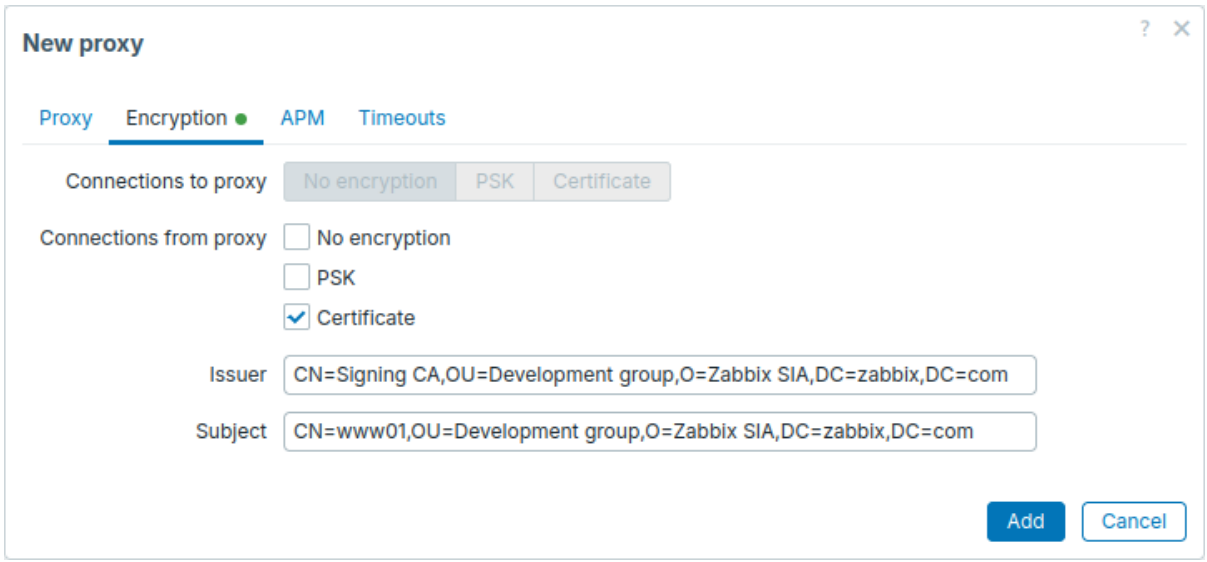
```
TLSConnect=cert
TLSAccept=cert
TLSCAFile=/home/zabbix/zabbix_ca_file.crt
TLSServerCertIssuer=CN=Signing CA,OU=Development group,O=Zabbix SIA,DC=zabbix,DC=com
TLSServerCertSubject=CN=Zabbix server,OU=Development group,O=Zabbix SIA,DC=zabbix,DC=com
TLSCertFile=/home/zabbix/zabbix_proxy.crt
TLSKeyFile=/home/zabbix/zabbix_proxy.key
```

3. Configure encryption for this proxy in Zabbix frontend:

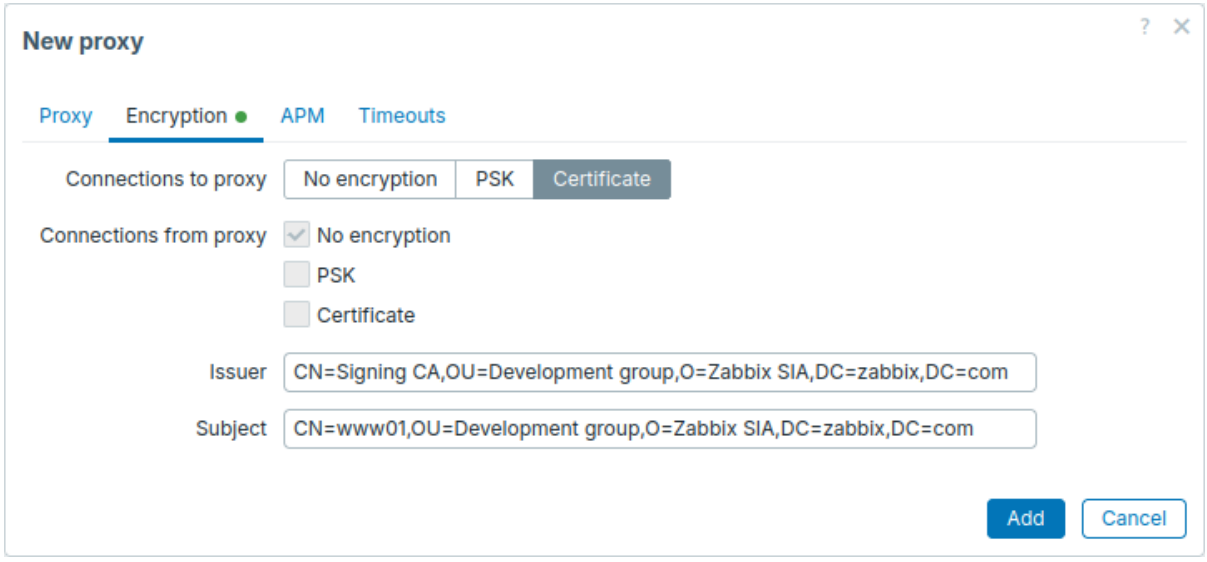
- Go to: *Administration* > *Proxies*.
- Select the proxy and click the *Encryption* tab.

In the examples below, the *Issuer* and *Subject* fields are filled in. For more information on why and how to use these fields, see Restricting allowed certificate issuer and subject.

For active proxy:



For passive proxy:



Zabbix agent

1. Prepare files with the top-level CA certificates, the Zabbix agent certificate/certificate chain, and the private key as described in the Zabbix server section. Then, edit the `TLSCAFile`, `TLSCertFile`, and `TLSKeyFile` parameters in the Zabbix agent configuration file accordingly.

2. Edit additional TLS parameters in the Zabbix agent configuration file:

- For active agent: `TLSConnect=cert`
- For passive agent: `TLSAccept=cert`

Note:
To improve agent security, you can set the `TLSServerCertIssuer` and `TLSServerCertSubject` parameters. For more information, see Restricting allowed certificate issuer and subject.

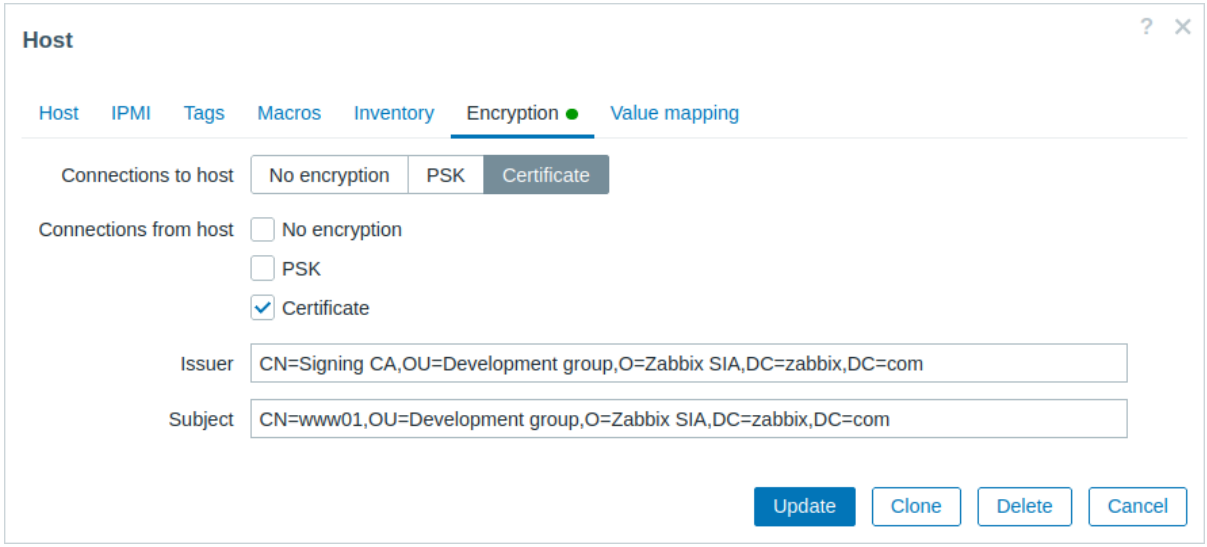
The TLS parameters in the final agent configuration file may look as follows. Note that the example assumes that the host is monitored by a proxy, hence it is specified as the certificate Subject:

```
TLSConnect=cert
TLSAccept=cert
TLSCAFile=/home/zabbix/zabbix_ca_file.crt
TLSServerCertIssuer=CN=Signing CA,OU=Development group,O=Zabbix SIA,DC=zabbix,DC=com
TLSServerCertSubject=CN=Zabbix proxy,OU=Development group,O=Zabbix SIA,DC=zabbix,DC=com
TLSCertFile=/home/zabbix/zabbix_agentd.crt
TLSKeyFile=/home/zabbix/zabbix_agentd.key
```

3. Configure encryption in Zabbix frontend for the host monitored by this agent.

- Go to: *Data collection* > *Hosts*.
- Select the host and click the *Encryption* tab.

In the example below, the *Issuer* and *Subject* fields are filled in. For more information on why and how to use these fields, see Restricting allowed certificate issuer and subject.



Zabbix web service

1. Prepare files with the top-level CA certificates, the Zabbix web service certificate/certificate chain, and the private key as described in the Zabbix server section. Then, edit the `TLSCAFile`, `TLSCertFile`, and `TLSKeyFile` parameters in the Zabbix web service configuration file accordingly.

2. Edit an additional TLS parameter in the Zabbix web service configuration file: `TLSAccept=cert`

TLS parameters in the final web service configuration file may look as follows:

```
TLSAccept=cert
TLSCAFile=/home/zabbix/zabbix_ca_file.crt
TLSCertFile=/home/zabbix/zabbix_web_service.crt
TLSKeyFile=/home/zabbix/zabbix_web_service.key
```

3. Configure Zabbix server to connect to the TLS-configured Zabbix web service by editing the `WebServiceURL` parameter in the Zabbix server configuration file:

```
WebServiceURL=https://example.com:443/report
```

Restricting allowed certificate issuer and subject When two Zabbix components (for example, server and agent) establish a TLS connection, they validate each other's certificates. If a peer certificate is signed by a trusted CA (with a pre-configured top-level certificate in `TLSCAFile`), is valid, has not expired, and passes other checks, then the communication between components can proceed. In this simplest case, the certificate issuer and subject are not verified.

However, this presents a risk: anyone with a valid certificate can impersonate anyone else (for example, a host certificate could be used to impersonate a server). While this may be acceptable in small environments where certificates are signed by a dedicated in-house CA and the risk of impersonation is low, it may not be sufficient in larger or more security-sensitive environments.

If your top-level CA issues certificates that should not be accepted by Zabbix or if you want to reduce the risk of impersonation, you can restrict allowed certificates by specifying their issuer and subject.

For example, in the Zabbix proxy configuration file, you could specify:

```
TLSServerCertIssuer=CN=Signing CA,OU=Development group,O=Zabbix SIA,DC=zabbix,DC=com
TLSServerCertSubject=CN=Zabbix server,OU=Development group,O=Zabbix SIA,DC=zabbix,DC=com
```

With these settings, an active proxy will not communicate with a Zabbix server whose certificate has a different issuer or subject. Similarly, a passive proxy will not accept requests from such a server.

Rules for matching `Issuer` and `Subject` strings

The rules for matching `Issuer` and `Subject` strings are as follows:

- `Issuer` and `Subject` strings are checked independently. Both are optional.
- An unspecified string means that any string is accepted.
- Strings are compared *as is* and must match exactly.
- UTF-8 characters are supported. However, wildcards (*) or regular expressions are not supported.
- The following RFC 4514 requirements are implemented - characters that require escaping (with a '\' backslash, U+005C):
  - anywhere in the string: '"' (U+0022), '+' (U+002B), ',' (U+002C), ';' (U+003B), '<' (U+003C), '>' (U+003E), '\\' (U+005C);
  - at the beginning of the string: space (' ', U+0020) or number sign ('#', U+0023);
  - at the end of the string: space (' ', U+0020).
- Null characters (U+0000) are not supported. If a null character is encountered, the matching will fail.
- RFC 4517 and RFC 4518 standards are not supported.

For example, if `Issuer` and `Subject` organization (`O`) strings contain trailing spaces and the `Subject` organizational unit (`OU`) string contains double quotes, these characters must be escaped:

```
TLSServerCertIssuer=CN=Signing CA,OU=Development head,O=\ Example SIA\ ,DC=example,DC=com
TLSServerCertSubject=CN=Zabbix server,OU=Development group \"5\",O=\ Example SIA\ ,DC=example,DC=com
```

Field order and formatting

Zabbix follows the recommendations of RFC 4514, which specifies a "reverse" order for these fields, starting with the lowest-level fields (`CN`), proceeding to the mid-level fields (`OU`, `O`), and concluding with the highest-level fields (`DC`).

```
TLSServerCertIssuer=CN=Signing CA,OU=Development group,O=Zabbix SIA,DC=zabbix,DC=com
TLSServerCertSubject=CN=Zabbix proxy,OU=Development group,O=Zabbix SIA,DC=zabbix,DC=com
```

In contrast, OpenSSL by default displays the `Issuer` and `Subject` strings in top-level to low-level order. In the following example, `Issuer` and `Subject` fields start with the top-level (`DC`) and end with the low-level (`CN`) field. The formatting with spaces and field separators also varies based on the options used, and thus will not match the format required by Zabbix.

```
$ openssl x509 -noout -in /home/zabbix/zabbix_proxy.crt -issuer -subject
issuer= /DC=com/DC=zabbix/O=Zabbix SIA/OU=Development group/CN=Signing CA
subject= /DC=com/DC=zabbix/O=Zabbix SIA/OU=Development group/CN=Zabbix proxy

$ openssl x509 -noout -text -in /home/zabbix/zabbix_proxy.crt
Certificate:
    ...
        Issuer: DC=com, DC=zabbix, O=Zabbix SIA, OU=Development group, CN=Signing CA
        ...
        Subject: DC=com, DC=zabbix, O=Zabbix SIA, OU=Development group, CN=Zabbix proxy
```

To format *Issuer* and *Subject* strings correctly for Zabbix, invoke OpenSSL with the following options:

```
$ openssl x509 -noout -issuer -subject \
    -nameopt esc_2253,esc_ctrl,utf8,dump_nostr,dump_unknown,dump_der,sep_comma_plus,dn_rev,sname\
    -in /home/zabbix/zabbix_proxy.crt
```

The output will then be in reverse order, comma-separated, and usable in Zabbix configuration files and frontend:

```
issuer=CN=Signing CA,OU=Development group,O=Zabbix SIA,DC=zabbix,DC=com
subject=CN=Zabbix proxy,OU=Development group,O=Zabbix SIA,DC=zabbix,DC=com
```

Limitations on using X.509 v3 certificate extensions When implementing X.509 v3 certificates within Zabbix, certain extensions may not be fully supported or could result in inconsistent behavior.

Subject Alternative Name extension

Zabbix does not support the *Subject Alternative Name* extension, which is used to specify alternative DNS names such as IP addresses or email addresses. Zabbix can only validate the value in the *Subject* field of the certificate (see Restricting Allowed Certificate Issuer and Subject). If certificates include the `subjectAltName` field, the outcome of certificate validation may vary depending on the specific crypto toolkits used to compile Zabbix components. As a result, Zabbix may either accept or reject certificates based on these combinations.

Extended Key Usage extension

Zabbix supports the *Extended Key Usage* extension. However, if used, it is generally required that both *clientAuth* (for TLS WWW client authentication) and *serverAuth* (for TLS WWW server authentication) attributes are specified. For example:

- In passive checks, where Zabbix agent operates as a TLS server, the *serverAuth* attribute must be included in the agent's certificate.
- For active checks, where the agent operates as a TLS client, the *clientAuth* attribute must be included in the agent's certificate.

While GnuTLS may issue a warning for key usage violations, it typically allows communication to proceed despite these warnings.

Name Constraints extension

Support for the *Name Constraints* extension varies among crypto toolkits. Ensure that your chosen toolkit supports this extension. This extension may restrict Zabbix from loading CA certificates if this section is marked as critical, depending on the specific toolkit in use.

Certificate Revocation Lists (CRL) If a certificate is compromised, the Certificate Authority (CA) can revoke it by including the certificate in a Certificate Revocation List (CRL). CRLs are managed through configuration files and can be specified using the `TLSCRLFile` parameter in server, proxy, and agent configuration files. For example:

```
TLSCRLFile=/home/zabbix/zabbix_crl_file.crt
```

In this case, `zabbix_crl_file.crt` may contain CRLs from multiple CAs, and could look like this:

```
-----BEGIN X509 CRL-----
MIIB/DCB5QIBATANBgkqhkiG9w0BAQUFADCBgTETMBEGCgmSJomT8ixkARkWA2Nv
...
treZeUPjb7LSmZ3K2hpbZN7SoOZcAoHQ3GWd9npuctg=
-----END X509 CRL-----
-----BEGIN X509 CRL-----
MIIB+TCB4gIBATANBgkqhkiG9w0BAQUFADB/MRMwEQYKCZImiZPyLGQBGRYDY29t
...
CAEebS2CND3ShBedZ8YSil59O6JvaDP61lR5lNs=
-----END X509 CRL-----
```

The CRL file is loaded only when Zabbix starts. To update the CRL, restart Zabbix.

> **Attention:**
> If Zabbix components are compiled with OpenSSL and CRLs are used, ensure that each top-level and intermediate CA in the certificate chains has a corresponding CRL (even if it is empty) included in the `TLSCRLFile`.

Troubleshooting certificate problems OpenSSL used with CRLs and for some CA in the certificate chain its CRL is not included in `TLSCRLFile`

In TLS server log in case of *OpenSSL* peer:

```
failed to accept an incoming connection: from 127.0.0.1: TLS handshake with 127.0.0.1 returned error code
    file s3_srvr.c line 3251: error:14089086: SSL routines:ssl3_get_client_certificate:certificate verify
    TLS write fatal alert "unknown CA"
```

In TLS server log in case of *GnuTLS* peer:

```
failed to accept an incoming connection: from 127.0.0.1: TLS handshake with 127.0.0.1 returned error code
    file rsa_pk1.c line 103: error:0407006A: rsa routines:RSA_padding_check_PKCS1_type_1:\
    block type is not 01 file rsa_eay.c line 705: error:04067072: rsa routines:RSA_EAY_PUBLIC_DECRYPT:padd
```

CRL expired or expires during server operation

OpenSSL, in server log:

- before expiration:

```
cannot connect to proxy "proxy-openssl-1.0.1e": TCP successful, cannot establish TLS to [[127.0.0.1]:20004
    SSL_connect() returned SSL_ERROR_SSL: file s3_clnt.c line 1253: error:14090086:\
    SSL routines:ssl3_get_server_certificate:certificate verify failed:\
    TLS write fatal alert "certificate revoked"
```

- after expiration:

```
cannot connect to proxy "proxy-openssl-1.0.1e": TCP successful, cannot establish TLS to [[127.0.0.1]:20004
    SSL_connect() returned SSL_ERROR_SSL: file s3_clnt.c line 1253: error:14090086:\
    SSL routines:ssl3_get_server_certificate:certificate verify failed:\
    TLS write fatal alert "certificate expired"
```

The point here is that with valid CRL a revoked certificate is reported as "certificate revoked". When CRL expires the error message changes to "certificate expired" which is quite misleading.

GnuTLS, in server log:

- before and after expiration the same:

```
cannot connect to proxy "proxy-openssl-1.0.1e": TCP successful, cannot establish TLS to [[127.0.0.1]:20004
     invalid peer certificate: The certificate is NOT trusted. The certificate chain is revoked.
```

Self-signed certificate, unknown CA

OpenSSL, in log:

```
error:'self signed certificate: SSL_connect() set result code to SSL_ERROR_SSL: file ../ssl/statem/statem_
     line 1924: error:1416F086:SSL routines:tls_process_server_certificate:certificate verify failed:\
     TLS write fatal alert "unknown CA"'
```

This was observed when server certificate by mistake had the same Issuer and Subject string, although it was signed by CA. Issuer and Subject are equal in top-level CA certificate, but they cannot be equal in server certificate. (The same applies to proxy and agent certificates.)

To check whether a certificate contains the same Issuer and Subject entries, run:

```
openssl x509 -in <yourcertificate.crt> -noout -text
```

It is acceptable for the root (top-level) certificate to have identical values for Issuer and Subject.

Encryption using pre-shared keys

Overview Each pre-shared key (PSK) in Zabbix actually is a pair of:

- non-secret PSK identity string,
- secret PSK string value.

PSK identity string is a non-empty UTF-8 string. For example, "PSK ID 001 Zabbix agentd". It is a unique name by which this specific PSK is referred to by Zabbix components. Do not put sensitive information in PSK identity string - it is transmitted over the network unencrypted.

PSK value is a hard to guess string of hexadecimal digits, for example, e560cb0d918d26d31b4f642181f5f570ad89a390931102e5391d08

Warning:
Never use the example PSK values provided in the documentation in production environments. Instead, generate a unique, cryptographically secure PSK.

Size limits There are size limits for PSK identity and value in Zabbix, in some cases a crypto library can have lower limit:

Component	PSK identity max size	PSK value min size	PSK value max size
Zabbix	128 UTF-8 characters	128-bit (16-byte PSK, entered as 32 hexadecimal digits)	2048-bit (256-byte PSK, entered as 512 hexadecimal digits)
GnuTLS	128 bytes (may include UTF-8 characters)	-	2048-bit (256-byte PSK, entered as 512 hexadecimal digits)
OpenSSL 1.0.x, 1.1.0	127 bytes (may include UTF-8 characters)	-	2048-bit (256-byte PSK, entered as 512 hexadecimal digits)
OpenSSL 1.1.1	127 bytes (may include UTF-8 characters)	-	512-bit (64-byte PSK, entered as 128 hexadecimal digits)
OpenSSL 1.1.1a and later	127 bytes (may include UTF-8 characters)	-	2048-bit (256-byte PSK, entered as 512 hexadecimal digits)

Attention:
Zabbix frontend allows configuring up to 128-character long PSK identity string and 2048-bit long PSK regardless of crypto libraries used.
 If some Zabbix components support lower limits, it is the user's responsibility to configure PSK identity and value with allowed length for these components.
 Exceeding length limits results in communication failures between Zabbix components.

Before Zabbix server connects to agent using PSK, the server looks up the PSK identity and PSK value configured for that agent in database (actually in configuration cache). Upon receiving a connection the agent uses PSK identity and PSK value from its configuration file. If both parties have the same PSK identity string and PSK value the connection may succeed.

Attention:
Each PSK identity must be paired with only one value. It is the user's responsibility to ensure that there are no two PSKs with the same identity string but different values. Failing to do so may lead to unpredictable errors or disruptions of communication between Zabbix components using PSKs with this PSK identity string.

Generating PSK For example, a 256-bit (32 bytes) PSK can be generated using the following commands:

- with *OpenSSL*:

```
$ openssl rand -hex 32
af8ced32dfe8714e548694e2d29e1a14ba6fa13f216cb35c19d0feb1084b0429
```

- with *GnuTLS*:

```
$ psktool -u psk_identity -p database.psk -s 32
Generating a random key for user 'psk_identity'
Key stored to database.psk

$ cat database.psk
psk_identity:9b8eafedfaae00cece62e85d5f4792c7d9c9bcc851b23216a1d300311cc4f7cb
```

Note that "psktool" above generates a database file with a PSK identity and its associated PSK. Zabbix expects just a PSK in the PSK file, so the identity string and colon (':') should be removed from the file.

Configuring PSK for server-agent communication (example) On the agent host, write the PSK value generated above into a file, for example, `/home/zabbix/zabbix_agentd.psk`. The file must contain the PSK value on the first line.

Set access rights to PSK file - it must be readable only by Zabbix user.

Edit TLS parameters in the Zabbix agent configuration file (`zabbix_agentd.conf`). For example, set:

```
TLSConnect=psk
TLSAccept=psk
TLSPSKFile=/home/zabbix/zabbix_agentd.psk
TLSPSKIdentity=PSK 001
```

The agent will connect to server (active checks) and accept from server and `zabbix_get` only connections using PSK. PSK identity will be "PSK 001".

Restart the agent. Now you can test the connection using `zabbix_get`, for example:

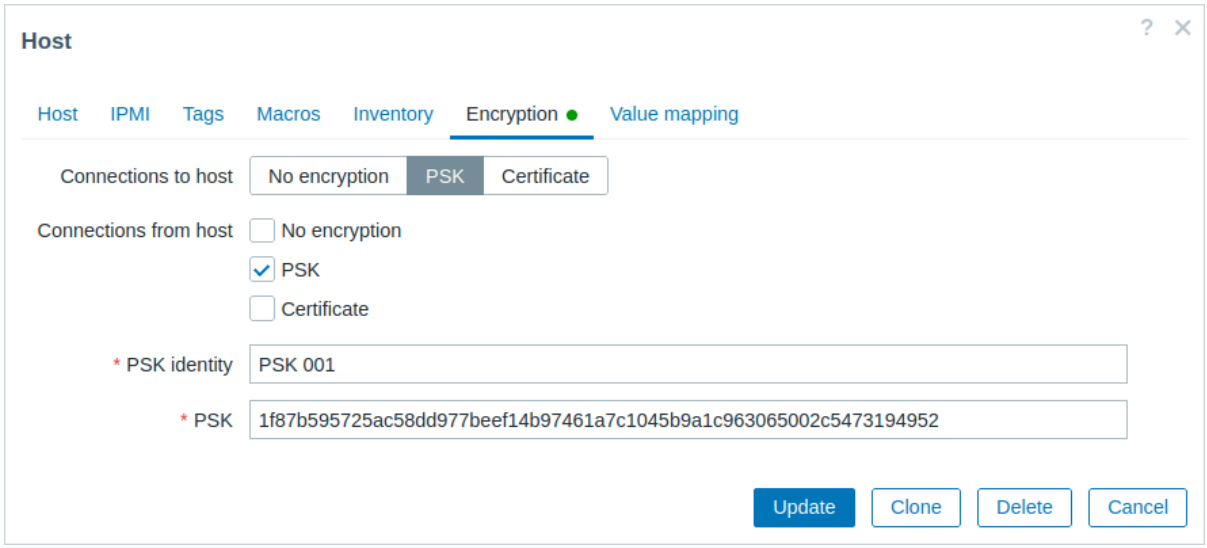
```
zabbix_get -s 127.0.0.1 -k "system.cpu.load[all,avg1]" --tls-connect=psk --tls-psk-identity="PSK 001" --tl
```

(To minimize downtime see how to change connection type in Connection encryption management).

Configure PSK encryption for this agent in Zabbix frontend:

1. Go to: *Data collection* > *Hosts*.
2. Select the host and click the *Encryption* tab.
3. Configure the PSK encryption details.

Example:



All mandatory input fields are marked with a red asterisk.

When configuration cache is synchronized with database the new connections will use PSK. Check server and agent logfiles for error messages.

Configuring PSK for server - active proxy communication (example) On the proxy, write the PSK value into a file, for example, `/home/zabbix/zabbix_proxy.psk`. The file must contain PSK in the first text string, for example:

```
e560cb0d918d26d31b4f642181f5f570ad89a390931102e5391d08327ba434e9
```

Set access rights to PSK file - it must be readable only by Zabbix user.

Edit TLS parameters in the Zabbix proxy configuration file (`zabbix_proxy.conf`). For example, set:

```
TLSConnect=psk
TLSPSKFile=/home/zabbix/zabbix_proxy.psk
TLSPSKIdentity=PSK 002
```

The proxy will connect to server using PSK. PSK identity will be "PSK 002".

(To minimize downtime see how to change connection type in Connection encryption management).

Configure PSK for this proxy in Zabbix frontend. Go to *Administration* > *Proxies*, select the proxy, go to *Encryption* tab. In "Connections from proxy" mark PSK. Paste into "PSK identity" field "PSK 002" and "e560cb0d918d26d31b4f642181f5f570ad89a390931102e5391d083 into "PSK" field. Click "Update".

Restart proxy. It will start using PSK-based encrypted connections to server. Check server and proxy logfiles for error messages.

For a passive proxy the procedure is very similar. The only difference - set `TLSAccept=psk` in the Zabbix proxy configuration file and set *Connections to proxy* in Zabbix frontend to PSK.

Troubleshooting PSK problems PSK contains an odd number of hex-digits

Proxy or agent does not start, message in the proxy or agent log:

```
invalid PSK in file "/home/zabbix/zabbix_proxy.psk"
```

PSK identity string longer than 128 bytes is passed to GnuTLS

In TLS client side log:

```
gnutls_handshake() failed: -110 The TLS connection was non-properly terminated.
```

In TLS server side log.

```
gnutls_handshake() failed: -90 The SRP username supplied is illegal.
```

Too long PSK value used with OpenSSL 1.1.1

In connecting-side log:

```
...OpenSSL library (version OpenSSL 1.1.1  11 Sep 2018) initialized
...
...In zbx_tls_connect(): psk_identity:"PSK 1"
...zbx_psk_client_cb() requested PSK identity "PSK 1"
...End of zbx_tls_connect():FAIL error:'SSL_connect() set result code to SSL_ERROR_SSL: file ssl\statem\ex
```

In accepting-side log:

```
...Message from 123.123.123.123 is missing header. Message ignored.
```

This problem typically arises when upgrading OpenSSL from 1.0.x or 1.1.0 to 1.1.1 and if the PSK value is longer than 512-bit (64-byte PSK, entered as 128 hexadecimal digits).

See also: Value size limits

Secure connection to the database

Overview

This section provides Zabbix setup steps and configuration examples for secure TLS connections between:

Database	Zabbix components
MySQL	Zabbix frontend, Zabbix server, Zabbix proxy
PostgreSQL	Zabbix frontend, Zabbix server, Zabbix proxy

To set up connection encryption within the DBMS, see official vendor documentation for details:

- MySQL: source and replica replication database servers.
- MySQL: group replication, etc. database servers.
- PostgreSQL encryption options.

All examples are based on the GA releases of MySQL CE (8.0) and PgSQL (13) available through official repositories using CentOS 8.

Requirements

The following is required to set up encryption:

- Developer-supported operating system with OpenSSL >=1.1.X or alternative.

Note:
It is recommended to avoid OS in the end-of-life status, especially in the case of new installations

- Database engine (RDBMS) installed and maintained from the official repository provided by developer. Operating systems often shipped with outdated database software versions for which encryption support is not implemented, for example RHEL 7 based systems and PostgreSQL 9.2, MariaDB 5.5 without encryption support.

Terminology

Setting this option enforces to use TLS connection to database from Zabbix server/proxy and frontend to database:

- `required` - connect using TLS as transport mode without identity checks
- `verify_ca` - connect using TLS and verify certificate
- `verify_full` - connect using TLS, verify certificate and verify that database identity (CN) specified by DBHost matches its certificate

Zabbix configuration

Frontend to the database

A secure connection to the database can be configured during frontend installation:

- Mark the *Database TLS encryption* checkbox in the Configure DB connection step to enable transport encryption.
- Mark the *Verify database certificate* checkbox that appears when *TLS encryption* field is checked to enable encryption with certificates.

Note:
For MySQL, the *Database TLS encryption* checkbox is disabled, if *Database host* is set to localhost, because connection that uses a socket file (on Unix) or shared memory (on Windows) cannot be encrypted.
 For PostgreSQL, the *TLS encryption* checkbox is disabled, if the value of the *Database host* field begins with a slash or the field is empty.

The following parameters become available in the TLS encryption in certificates mode (if both checkboxes are marked):

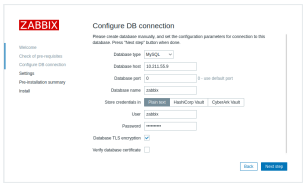
Parameter	Description
Database TLS CA file	Specify the full path to a valid TLS certificate authority (CA) file.
Database TLS key file	Specify the full path to a valid TLS key file.
Database TLS certificate file	Specify the full path to a valid TLS certificate file.
Database host verification	Mark this checkbox to activate host verification. Disabled for MySQL because the PHP MySQL library does not allow the peer certificate validation step to be skipped.
Database TLS cipher list	Specify a custom list of valid ciphers. The format of the cipher list must conform to the OpenSSL standard. Available for MySQL only.

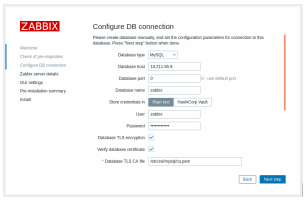
Attention:
TLS parameters must point to valid files. If they point to non-existent or invalid files, it will lead to the authorization error.
 If certificate files are writable, the frontend generates a warning in the System information report that "TLS certificate files must be read-only." (displayed only if the PHP user is the owner of the certificate).
Certificates protected by passwords are not supported.

Use cases

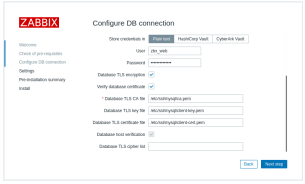
Zabbix frontend uses GUI interface to define possible options: `required`, `verify_ca`, `verify_full`. Specify required options in the installation wizard step *Configure DB connections*. These options are mapped to the frontend configuration file (`zabbix.conf.php`) in the following manner:

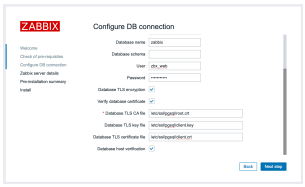
GUI settings	Configuration file	Description	Result

	...
	// Used for TLS connection.
	`$DB['ENCRYPTION'] = true;`
	`$DB['KEY_FILE'] = '';`
	`$DB['CERT_FILE'] = '';`
	`$DB['CA_FILE'] = '';`
	`$DB['VERIFY_HOST'] = false;`
	`$DB['CIPHER_LIST'] = '';`
	...
		1. Check *Database TLS encryption*
		2. Leave *Verify database certificate* unchecked
			Enable `required` mode.

	...
	`$DB['ENCRYPTION'] = true;`
	`$DB['KEY_FILE'] = '';`
	`$DB['CERT_FILE'] = '';`
	`$DB['CA_FILE'] = '/etc/ssl/mysql/ca.pem';`
	`$DB['VERIFY_HOST'] = false;`
	`$DB['CIPHER_LIST'] = '';`
	...
		1. Check *Database TLS encryption* and *Verify database certificate*
		2. Specify path to *Database TLS CA file*
			Enable `verify_ca` mode.

GUI settings	Configuration file	Description	Result

	...
	// Used for TLS connection with strictly defined Cipher list.
	`$DB['ENCRYPTION'] = true;`
	`$DB['KEY_FILE'] = '<key_file_path>';`
	`$DB['CERT_FILE'] = '<key_file_path>';`
	`$DB['CA_FILE'] = '<key_file_path>';`
	`$DB['VERIFY_HOST'] = true;`
	`$DB['CIPHER_LIST'] = '<cipher_list>';`
	...
	
	Or:
	
	...
	// Used for TLS connection without Cipher list defined - selected by MySQL server
	`$DB['ENCRYPTION'] = true;`
	`$DB['KEY_FILE'] = '<key_file_path>';`
	`$DB['CERT_FILE'] = '<key_file_path>';`
	`$DB['CA_FILE'] = '<key_file_path>';`
	`$DB['VERIFY_HOST'] = true;`
	`$DB['CIPHER_LIST'] = '';`
	...
		1. Check *Database TLS encryption* and *Verify database certificate*
		2. Specify path to *Database TLS key file*
		3. Specify path to *Database TLS CA file*
		4. Specify path to *Database TLS certificate file*
		5. Specify *Database TLS cipher list* (optional)
			Enable `verify_full` mode for MySQL.

GUI settings	Configuration file	Description	Result
	... $DB['ENCRYPTION'] = true; $DB['KEY_FILE'] = '<key_file_path>'; $DB['CERT_FILE'] = '<key_file_path>'; $DB['CA_FILE'] = '<key_file_path>'; $DB['VERIFY_HOST'] = true; $DB['CIPHER_LIST'] = ' '; ...	1. Check *Database TLS encryption* and *Verify database certificate* 2. Specify path to *Database TLS key file* 3. Specify path to *Database TLS CA file* 4. Specify path to *Database TLS certificate file* 5. Check *Database host verification*	Enable `verify_full` mode for PostgreSQL.

See also: Encryption configuration examples for MySQL, Encryption configuration examples for PostgreSQL.

Zabbix server/proxy configuration

Secure connections to the database can be configured with the respective parameters in the Zabbix server and/or proxy configuration file.

Configuration	Result
None	Connection to the database without encryption.
Set `DBTLSConnect=required`	Server/proxy make a TLS connection to the database. An unencrypted connection is not allowed.
1. Set `DBTLSConnect=verify_ca` 2. Set `DBTLSCAFile` - specify the TLS certificate authority file	Server/proxy make a TLS connection to the database after verifying the database certificate.
1. Set `DBTLSConnect=verify_full` 2. Set `DBTLSCAFile` - specify TLS certificate authority file	Server/proxy make a TLS connection to the database after verifying the database certificate and the database host identity.
1. Set `DBTLSCAFile` - specify TLS certificate authority file 2. Set `DBTLSCertFile` - specify the client public key certificate file 3. Set `DBTLSKeyFile` - specify the client private key file	Server/proxy provide a client certificate while connecting to the database.
Set `DBTLSCipher` - the list of encryption ciphers that the client permits for connections using TLS protocols up to TLS 1.2 or `DBTLSCipher13` - the list of encryption ciphers that the client permits for connections using TLS 1.3 protocol	(MySQL) TLS connection is made using a cipher from the provided list. (PostgreSQL) Setting this option will be considered as an error.

MySQL encryption configuration

Overview

This section provides several encryption configuration examples for CentOS 8.2 and MySQL 8.4.0 and can be used as a quickstart guide for encrypting the connection to the database.

Attention:
If MySQL host is set to localhost, encryption options will not be available. In this case a connection between Zabbix frontend and the database uses a socket file (on Unix) or shared memory (on Windows) and cannot be encrypted.

> **Note:**
> List of encryption combinations is not limited to the ones listed on this page. There are a lot more combinations available.

Pre-requisites

Install MySQL database from the official repository.

See MySQL documentation for details on how to use MySQL repo.

MySQL server is ready to accept secure connections using a self-signed certificate.

To see which users are using an encrypted connection, run the following query (Performance Schema should be turned ON):

```
SELECT sbt.variable_value AS tls_version, t2.variable_value AS cipher, processlist_user AS user, processli
FROM performance_schema.status_by_thread  AS sbt
JOIN performance_schema.threads AS t ON t.thread_id = sbt.thread_id
JOIN performance_schema.status_by_thread AS t2 ON t2.thread_id = t.thread_id
WHERE sbt.variable_name = 'Ssl_version' and t2.variable_name = 'Ssl_cipher'
ORDER BY tls_version;
```

Transport-only encryption

MySQL configuration

Modern versions of the database are ready out-of-the-box for `required` encryption mode. A server-side certificate will be created after initial setup and launch.

Create users and roles for the main components (replace <strong_password> with a strong, unique password):

```
mysql> CREATE USER
 'zbx_srv'@'%' IDENTIFIED WITH caching_sha2_password BY '<strong_password>',
 'zbx_web'@'%' IDENTIFIED WITH caching_sha2_password BY '<strong_password>'
 REQUIRE SSL
 PASSWORD HISTORY 5;

mysql> CREATE ROLE 'zbx_srv_role', 'zbx_web_role';

mysql> GRANT SELECT, UPDATE, DELETE, INSERT, CREATE, DROP, ALTER, INDEX, REFERENCES ON zabbix.* TO 'zbx_sr
mysql> GRANT SELECT, UPDATE, DELETE, INSERT ON zabbix.* TO 'zbx_web_role';

mysql> GRANT 'zbx_srv_role' TO 'zbx_srv'@'%';
mysql> GRANT 'zbx_web_role' TO 'zbx_web'@'%';

mysql> SET DEFAULT ROLE 'zbx_srv_role' TO 'zbx_srv'@'%';
mysql> SET DEFAULT ROLE 'zbx_web_role' TO 'zbx_web'@'%';
```

Note that the X.509 protocol is not used to check identity, but the user is configured to use only encrypted connections. See MySQL documentation for more details about configuring users.

Run to check connection (socket connection cannot be used to test secure connections):

```
mysql -u zbx_srv -p -h 10.211.55.9 --ssl-mode=REQUIRED
```

Check current status and available cipher suites:

```
mysql> status
--------------
mysql Ver 8.4.0 for Linux on x86_64 (MySQL Community Server - GPL)

Connection id: 62
Current database:
Current user: zbx_srv@bfdb.local
SSL: Cipher in use is TLS_AES_256_GCM_SHA384


mysql> SHOW SESSION STATUS LIKE 'Ssl_cipher_list'\G;
*************************** 1. row ***************************
Variable_name: Ssl_cipher_list
Value: TLS_AES_256_GCM_SHA384:TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256:TLS_AES_128_CCM_SHA256:E
```

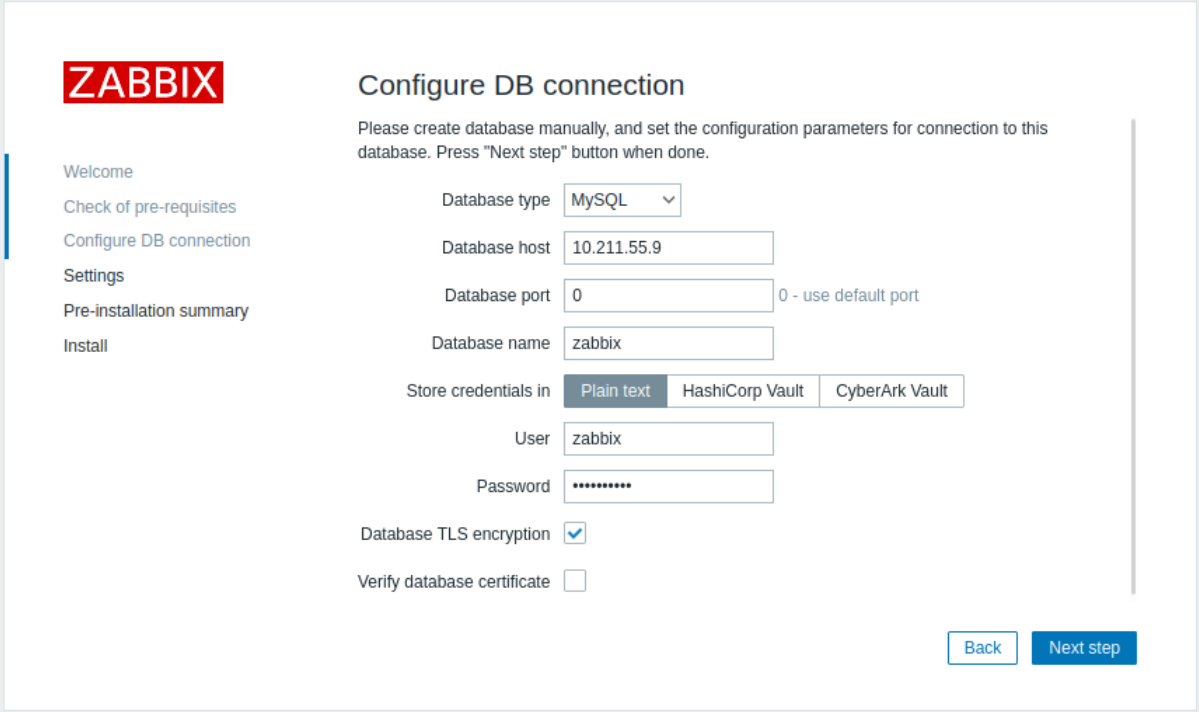
```
1 row in set (0.00 sec)

ERROR:
No query specified
```

Frontend

To enable transport-only encryption for connections between Zabbix frontend and the database:

- Check *Database TLS encryption*
- Leave *Verify database certificate* unchecked



Server

To enable transport-only encryption for connections between server and the database, edit the server configuration file:

```
...
DBHost=10.211.55.9
DBName=zabbix
DBUser=zbx_srv
DBPassword=<strong_password>
DBTLSConnect=required
...
```

Encryption with certificate authority verification

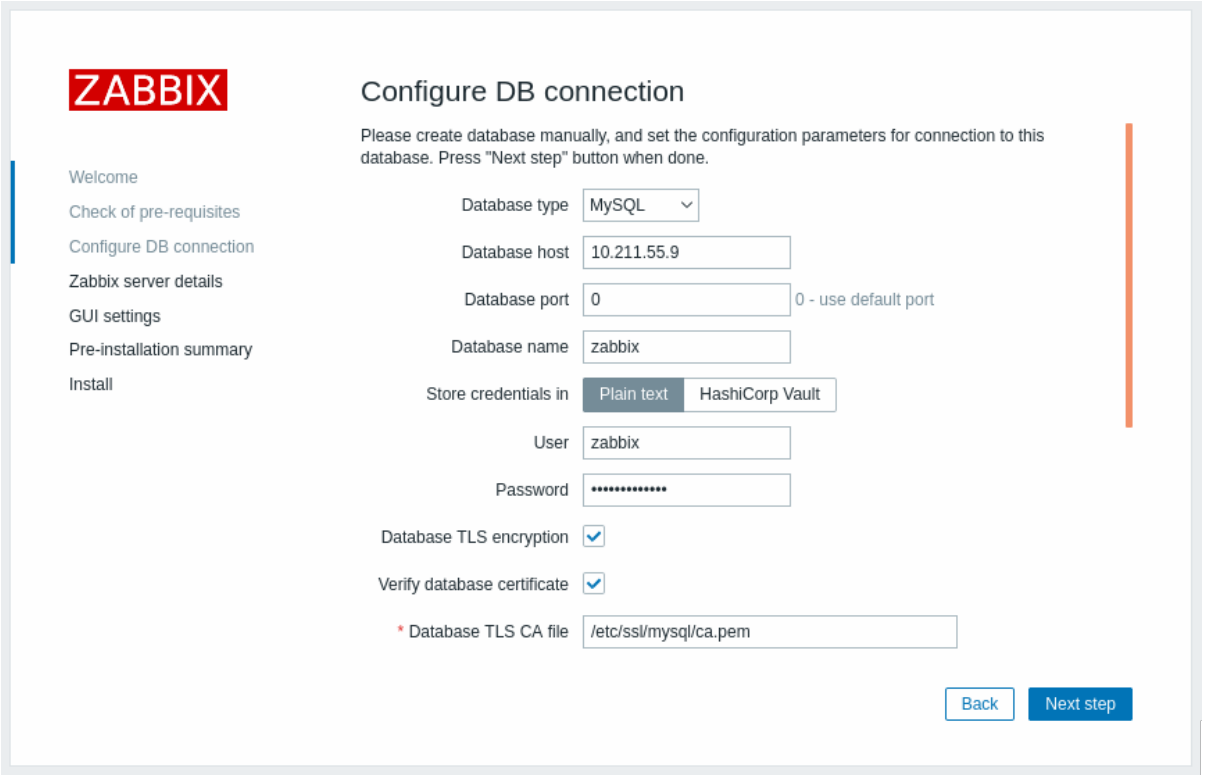
Copy required MySQL CA to the Zabbix frontend server, assign proper permissions to allow the webserver to read this file.

Note:
This mode doesn't work on RHEL 7 due to older MySQL libraries.

Frontend

To enable encryption with certificate verification for connections between Zabbix frontend and the database:

- Check *Database TLS encryption* and *Verify database certificate*
- Specify path to Database TLS CA file



Alternatively, this can be set in the frontend configuration file (`zabbix.conf.php`):

```
...
$DB['ENCRYPTION'] = true;
$DB['KEY_FILE'] = '';
$DB['CERT_FILE'] = '';
$DB['CA_FILE'] = '/etc/ssl/mysql/ca.pem';
$DB['VERIFY_HOST'] = false;
$DB['CIPHER_LIST'] = '';
...
```

Troubleshoot user using command-line tool to check if connection is possible for required user:

```
    mysql -u zbx_web -p -h 10.211.55.9 --ssl-mode=REQUIRED --ssl-ca=/var/lib/mysql/ca.pem
```

Server

To enable encryption with certificate verification for connections between Zabbix server and the database, edit the server configuration file:

```
...
DBHost=10.211.55.9
DBName=zabbix
DBUser=zbx_srv
DBPassword=<strong_password>
DBTLSConnect=verify_ca
DBTLSCAFile=/etc/ssl/mysql/ca.pem
...
```

Encryption with full verification

MySQL configuration

Set MySQL CE server configuration option (*/etc/my.cnf.d/server-tls.cnf*) to:

```
[mysqld]
...
#### in this examples keys are located in the MySQL CE datadir directory
ssl_ca=ca.pem
ssl_cert=server-cert.pem
ssl_key=server-key.pem

require_secure_transport=ON
```

```
tls_version=TLSv1.3
...
```

Keys for the MySQL CE server and client (Zabbix frontend) should be created manually according to the MySQL CE documentation: Creating SSL and RSA certificates and keys using MySQL orCreating SSL certificates and keys using openssl

Attention:
MySQL server certificate should contain the Common Name field set to the FQDN name as Zabbix frontend will use the DNS name to communicate with the database or IP address of the database host.

Create MySQL user:

```
mysql> CREATE USER
  'zbx_srv'@'%' IDENTIFIED WITH caching_sha2_password BY '<strong_password>',
  'zbx_web'@'%' IDENTIFIED WITH caching_sha2_password BY '<strong_password>'
  REQUIRE X509
  PASSWORD HISTORY 5;
```

Check if it is possible to log in with that user:

```
mysql -u zbx_web -p -h 10.211.55.9 --ssl-mode=VERIFY_IDENTITY --ssl-ca=/var/lib/mysql/ca.pem --ssl-cert=/v
```

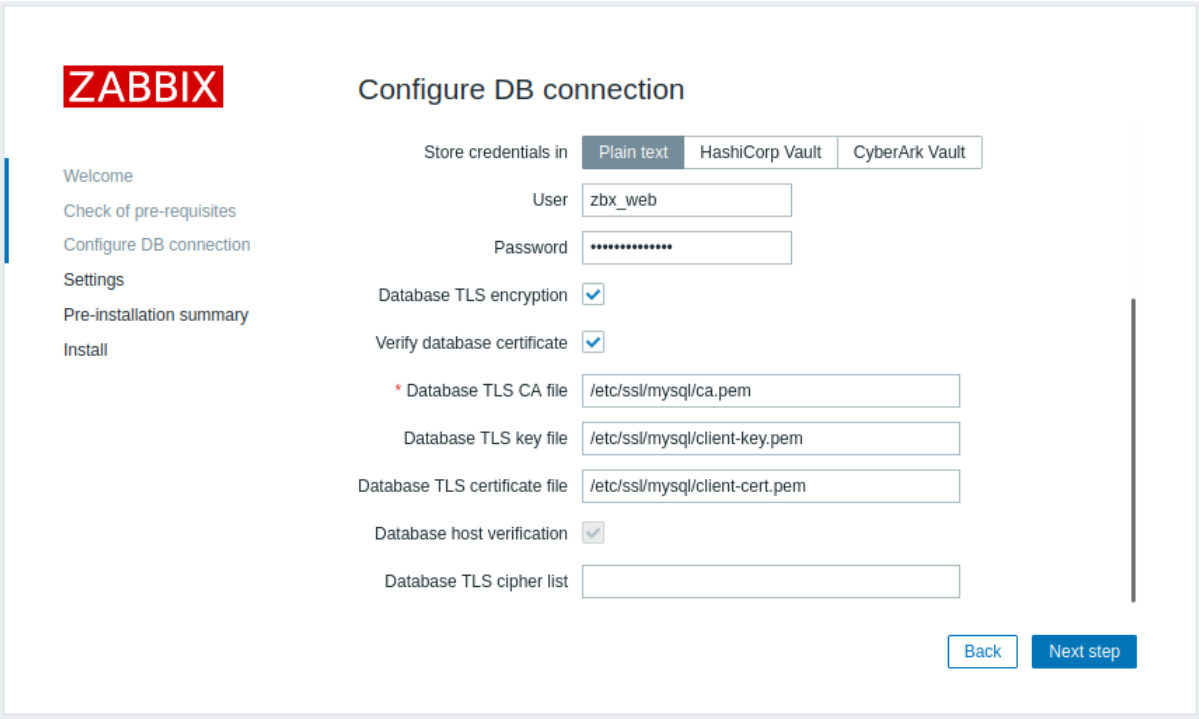
Frontend

To enable encryption with full verification for connections between Zabbix frontend and the database:

- Check *Database TLS encryption* and *Verify database certificate*
- Specify path to *Database TLS key file*
- Specify path to *Database TLS CA file*
- Specify path to *Database TLS certificate file*

Note that *Database host verification* is checked and grayed out - this step cannot be skipped for MySQL.

Warning:
If *Database TLS cipher list* field is left empty, the common ciphers permitted by both frontend (client) and server will be enabled. Alternatively, the ciphers can be set explicitly, in conformance with the cipher configuration requirements.



Alternatively, this can be set in the frontend configuration file (`zabbix.conf.php`):

```
...
// Used for TLS connection with strictly defined Cipher list.
$DB['ENCRYPTION'] = true;
$DB['KEY_FILE'] = '/etc/ssl/mysql/client-key.pem';
```

```
$DB['CERT_FILE'] = '/etc/ssl/mysql/client-cert.pem';
$DB['CA_FILE'] = '/etc/ssl/mysql/ca.pem';
$DB['VERIFY_HOST'] = true;
$DB['CIPHER_LIST'] = 'TLS_AES_256_GCM_SHA384:TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256:TLS_AES_1
...
// or
...
// Used for TLS connection without Cipher list defined - selected by MySQL server
$DB['ENCRYPTION'] = true;
$DB['KEY_FILE'] = '/etc/ssl/mysql/client-key.pem';
$DB['CERT_FILE'] = '/etc/ssl/mysql/client-cert.pem';
$DB['CA_FILE'] = '/etc/ssl/mysql/ca.pem';
$DB['VERIFY_HOST'] = true;
$DB['CIPHER_LIST'] = '';
...
```

Server

To enable encryption with full verification for connections between Zabbix server and the database, edit the server configuration file:

```
...
DBHost=10.211.55.9
DBName=zabbix
DBUser=zbx_srv
DBPassword=<strong_password>
DBTLSConnect=verify_full
DBTLSCAFile=/etc/ssl/mysql/ca.pem
DBTLSCertFile=/etc/ssl/mysql/client-cert.pem
DBTLSKeyFile=/etc/ssl/mysql/client-key.pem
...
```

PostgreSQL encryption configuration

Overview

This section provides several encryption configuration examples for CentOS 8.2 and PostgreSQL 13.

> **Note:**
> Connection between Zabbix frontend and PostgreSQL cannot be encrypted (parameters in GUI are disabled), if the value of *Database host* field begins with a slash or the field is empty.

Pre-requisites

Install the PostgreSQL database using the official repository.

PostgreSQL is not configured to accept TLS connections out-of-the-box. Please follow instructions from PostgreSQL documentation for certificate preparation with postgresql.conf and also for user access control through `ph_hba.conf`.

By default, the PostgreSQL socket is bound to the localhost, for the network remote connections to allow listening on the real network interface.

PostgreSQL settings for all modes can look like this:

`/var/lib/pgsql/13/data/postgresql.conf`:

```
...
ssl = on
ssl_ca_file = 'root.crt'
ssl_cert_file = 'server.crt'
ssl_key_file = 'server.key'
ssl_ciphers = 'HIGH:MEDIUM:+3DES:!aNULL'
ssl_prefer_server_ciphers = on
ssl_min_protocol_version = 'TLSv1.3'
...
```

For access control adjust `/var/lib/pgsql/13/data/pg_hba.conf`:

```
...
### require
hostssl all all 0.0.0.0/0 md5

### verify CA
hostssl all all 0.0.0.0/0 md5 clientcert=verify-ca

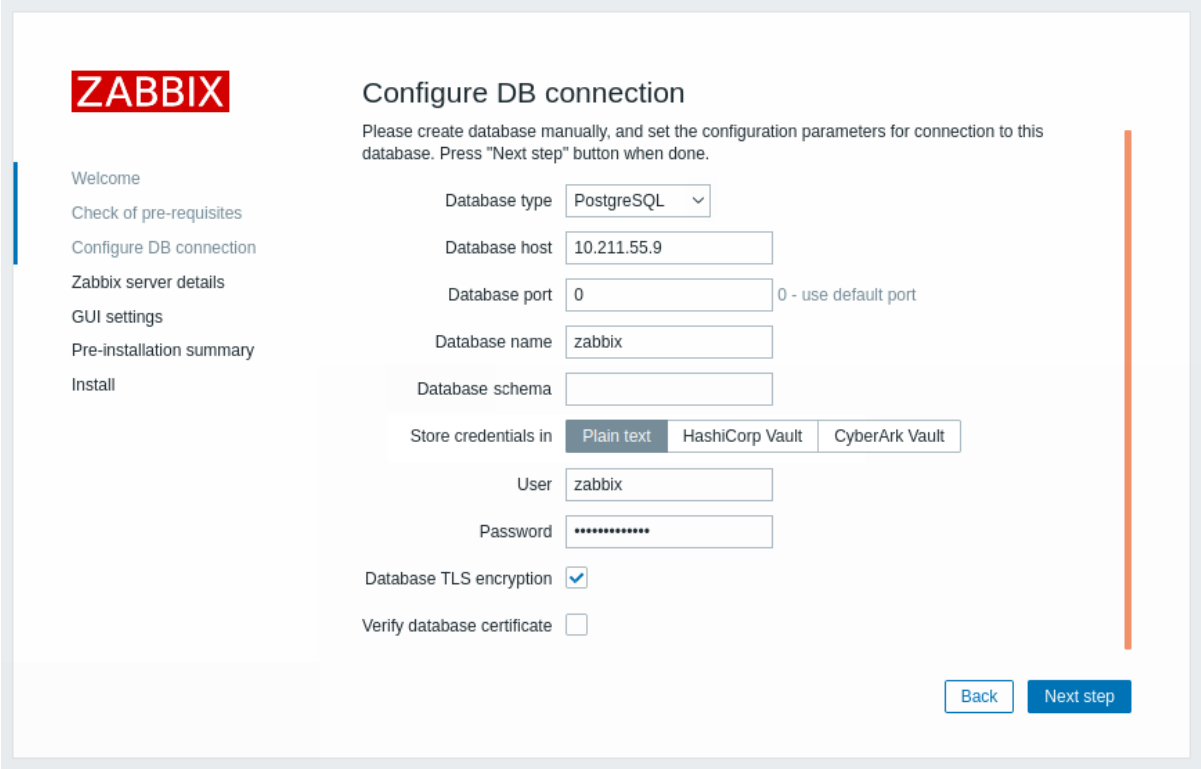
### verify full
hostssl all all 0.0.0.0/0 md5 clientcert=verify-full
...
```

Transport-only encryption

Frontend

To enable transport-only encryption for connections between Zabbix frontend and the database:

- Check *Database TLS encryption*
- Leave *Verify database certificate* unchecked



Server

To enable transport-only encryption for connections between server and the database, edit the server configuration file:

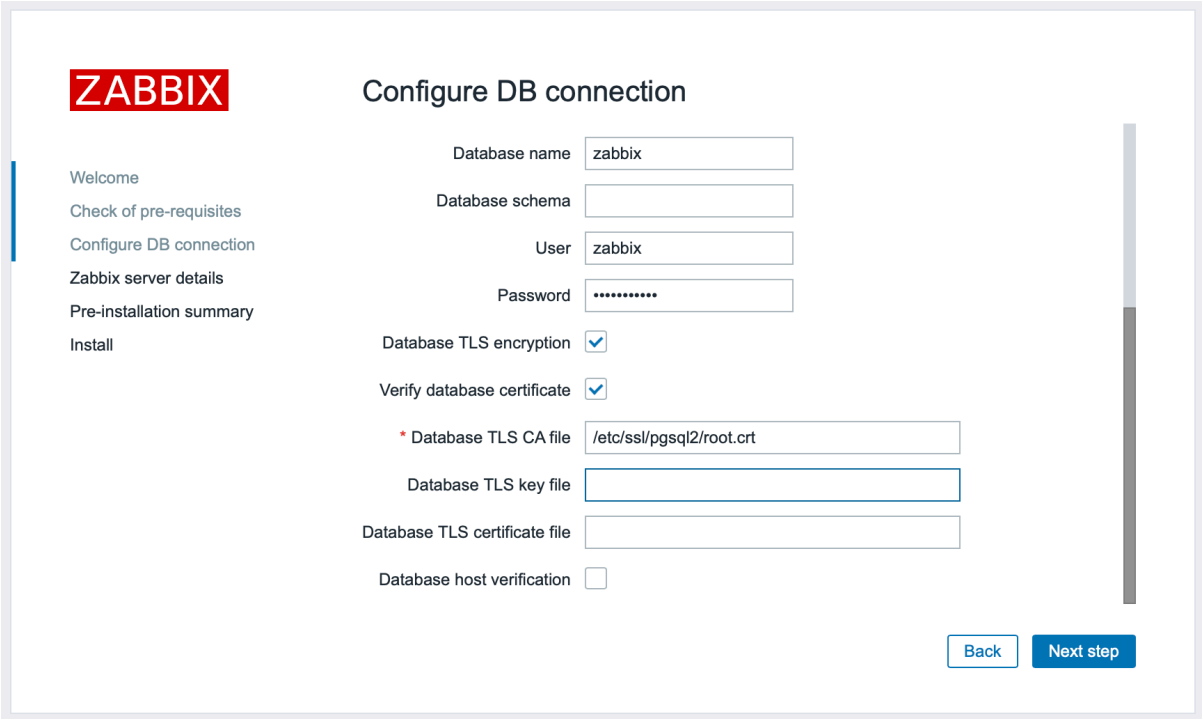
```
...
DBHost=10.211.55.9
DBName=zabbix
DBUser=zbx_srv
DBPassword=<strong_password>
DBTLSConnect=required
...
```

Encryption with certificate authority verification

Frontend

To enable encryption with certificate authority verification for connections between Zabbix frontend and the database:

- Check *Database TLS encryption* and *Verify database certificate*
- Specify path to *Database TLS CA file*



Alternatively, this can be set in the frontend configuration file (`zabbix.conf.php`):

```
...
$DB['ENCRYPTION'] = true;
$DB['KEY_FILE'] = '';
$DB['CERT_FILE'] = '';
$DB['CA_FILE'] = '/etc/ssl/pgsql/root.crt';
$DB['VERIFY_HOST'] = false;
$DB['CIPHER_LIST'] = '';
...
```

Server

To enable encryption with certificate verification for connections between Zabbix server and the database, edit the server configuration file (replace `<strong_password>` with a strong, unique password):

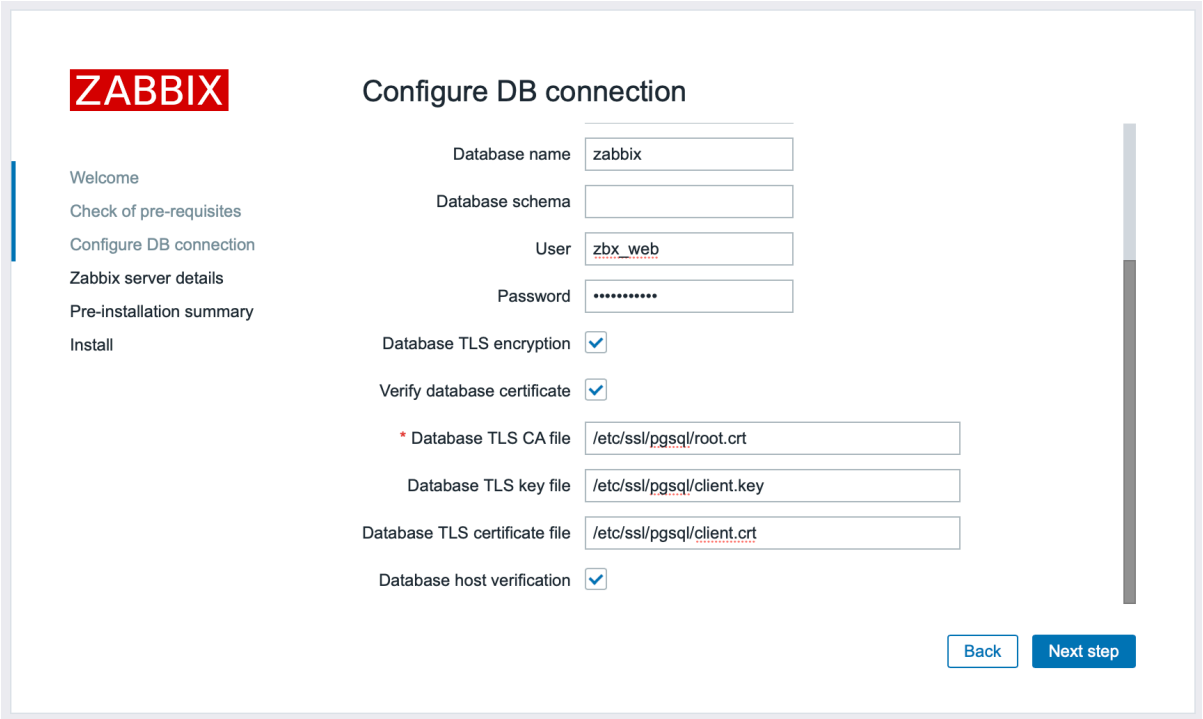
```
...
DBHost=10.211.55.9
DBName=zabbix
DBUser=zbx_srv
DBPassword=<strong_password>
DBTLSConnect=verify_ca
DBTLSCAFile=/etc/ssl/pgsql/root.crt
...
```

Encryption with full verification

Frontend

To enable encryption with certificate and database host identity verification for connections between Zabbix frontend and the database:

- Check *Database TLS encryption* and *Verify database certificate*
- Specify path to *Database TLS key file*
- Specify path to *Database TLS CA file*
- Specify path to *Database TLS certificate file*
- Check *Database host verification*



Alternatively, this can be set in the frontend configuration file (`zabbix.conf.php`):

```
$DB['ENCRYPTION'] = true;
$DB['KEY_FILE'] = '';
$DB['CERT_FILE'] = '';
$DB['CA_FILE'] = '/etc/ssl/pgsql/root.crt';
$DB['VERIFY_HOST'] = true;
$DB['CIPHER_LIST'] = '';
...
```

Server

To enable encryption with certificate and database host identity verification for connections between Zabbix server and the database, edit the server configuration file:

```
...
DBHost=10.211.55.9
DBName=zabbix
DBUser=zbx_srv
DBPassword=<strong_password>
DBTLSConnect=verify_full
DBTLSCAFile=/etc/ssl/pgsql/root.crt
DBTLSCertFile=/etc/ssl/pgsql/client.crt
DBTLSKeyFile=/etc/ssl/pgsql/client.key
...
```

Secure connection between frontend and server

Overview

This section provides Zabbix setup steps and configuration examples for secure TLS connections between Zabbix frontend and Zabbix server.

Configuration

By default, communication between Zabbix frontend and Zabbix server is unencrypted. For better security, enable TLS on both sides. Below is an example of the simplest way to do this.

1. Generate certificates and keys.

Create a working directory:

```
sudo mkdir -p /etc/zabbix/ssl && cd /etc/zabbix/ssl
```

Create a CA certificate (adjust the `MyZabbixCA` value to match the actual common name):

```
sudo openssl genrsa -out ca.key 4096
sudo openssl req -new -x509 -days 3650 -key ca.key -out ca.crt -subj "/CN=MyZabbixCA/"
```

Generate a private key and certificate for Zabbix server (adjust the `zabbix-server.example.com` value to match the actual common name):

```
sudo openssl genrsa -out server.key 2048
sudo openssl req -new -key server.key -out server.csr -subj "/CN=zabbix-server.example.com/"
sudo openssl x509 -req -days 365 -in server.csr -CA ca.crt -CAkey ca.key -CAcreateserial -sha256 -out serv
```

Generate a private key and certificate for Zabbix frontend (adjust the `zabbix-frontend-node` value to match the actual common name):

```
sudo openssl genrsa -out frontend.key 2048
sudo openssl req -new -key frontend.key -out frontend.csr -subj "/CN=zabbix-frontend-node/"
sudo openssl x509 -req -days 365 -in frontend.csr -CA ca.crt -CAkey ca.key -CAcreateserial -sha256 -out fr
```

2. Set proper permissions.

For Zabbix server (adjust ownership/group according to your distribution's Zabbix server daemon user):

```
sudo chown root:zabbix /etc/zabbix/ssl/server.{crt,key} /etc/zabbix/ssl/ca.crt
sudo chmod 640 /etc/zabbix/ssl/server.key
sudo chmod 644 /etc/zabbix/ssl/server.crt /etc/zabbix/ssl/ca.crt
```

For frontend (adjust ownership/group according to your distribution's webserver user):

```
sudo chown root:www-data /etc/zabbix/ssl/frontend.{crt,key}
sudo chmod 640 /etc/zabbix/ssl/frontend.key
sudo chmod 644 /etc/zabbix/ssl/frontend.crt
```

3. Configure Zabbix server.

In the server configuration file (`zabbix_server.conf`) add:

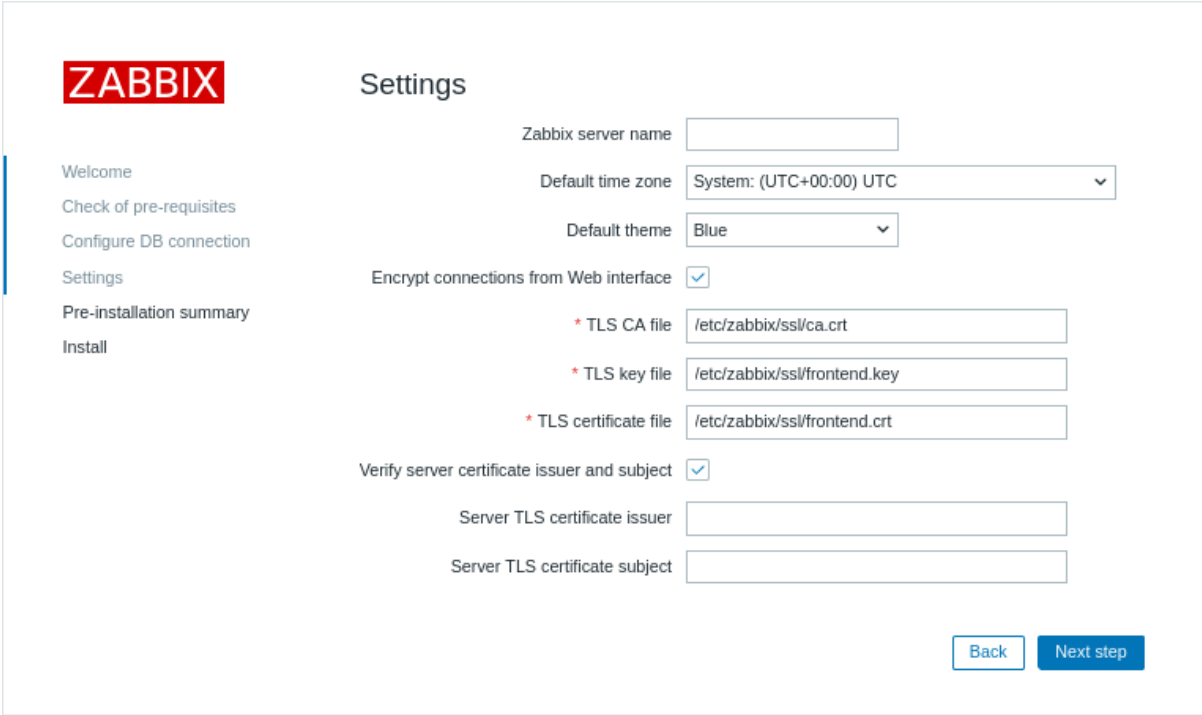
```
TLSFrontendAccept=cert
TLSCertFile=/etc/zabbix/ssl/server.crt
TLSKeyFile=/etc/zabbix/ssl/server.key
TLSCAFile=/etc/zabbix/ssl/ca.crt
### Optionally:
### TLSFrontendCertIssuer=CN=MyZabbixCA
### TLSFrontendCertSubject=CN=zabbix-frontend-node
```

Then restart the server:

```
sudo systemctl restart zabbix-server
```

4. Configure Zabbix frontend.

During web interface installation, enable the *Encrypt connections from Web interface* option (and *Verify server certificate issuer and subject* option, if necessary) and fill in the *TLS CA file*, *TLS key file*, *TLS certificate file* fields (and *Server TLS certificate issuer* and *Server TLS certificate subject* fields, if necessary):



Parameter	Description
TLS CA file	Specify the full path to the Certificate Authority (CA) certificate file used to verify the server's certificate.
TLS key file	Specify the full path to the client private key file corresponding to the client certificate.
TLS certificate file	Specify the full path to the client certificate file if mutual TLS authentication is required.
Server TLS certificate issuer	Specify a Distinguished Name (DN) of the issuer to match against the server's certificate.
Server TLS certificate subject	Specify a Distinguished Name (DN) of the subject to match against the server's certificate.

On existing installations, edit the following parameters in the frontend configuration file:

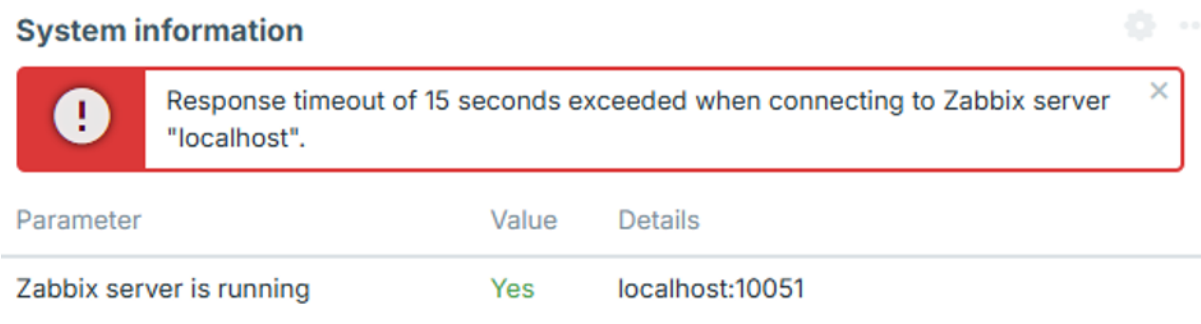
```
$ZBX_SERVER_TLS['ACTIVE'] = 'true';
$ZBX_SERVER_TLS['CA_FILE'] = '/etc/zabbix/ssl/ca.crt';
$ZBX_SERVER_TLS['KEY_FILE'] = '/etc/zabbix/ssl/frontend.key';
$ZBX_SERVER_TLS['CERT_FILE'] = '/etc/zabbix/ssl/frontend.crt';
// Optionally:
// $ZBX_SERVER_TLS['CERTIFICATE_ISSUER']  = 'CN=MyZabbixCA';
// $ZBX_SERVER_TLS['CERTIFICATE_SUBJECT'] = 'CN=zabbix-server.example.com';
```

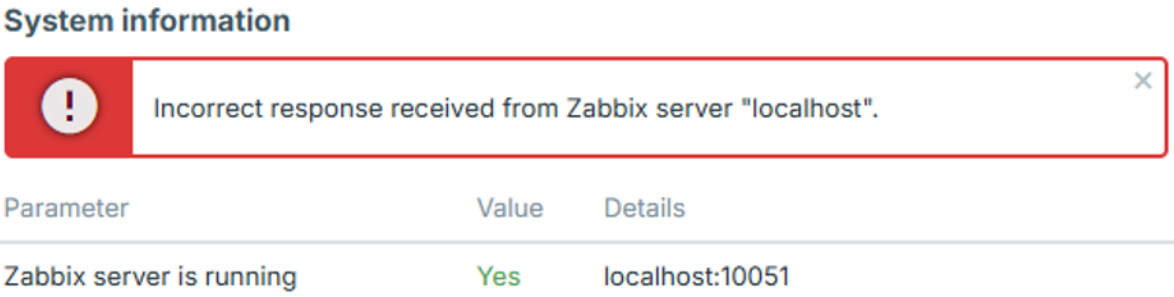
5. Verify encryption by confirming there are no error messages in Zabbix frontend or Zabbix server log file:

```
tail -f /var/log/zabbix/zabbix_server.log
```

Troubleshooting

Incorrect or incomplete TLS configuration can cause errors in the Zabbix frontend or indicate that the Zabbix server is running but the connection is not established successfully:





To identify the exact configuration issue, set the Zabbix server log level to 4, reproduce the error in the frontend, and check the Zabbix server log for the corresponding error message.

The following cases describe common errors and their possible causes.

Frontend IP address is not allowed

The Zabbix server log contains an error similar to:

```
54283:20260819:105926.400 frontend connection from "127.0.0.1" is not allowed by FrontendAllowedIP
```

The connection is dropped after the TCP handshake because the frontend IP address is not authorized.

Add the frontend IP address to the `FrontendAllowedIP` parameter in `zabbix_server.conf` and restart the Zabbix server.

Unencrypted connection is not allowed

The Zabbix server log contains an error similar to:

```
56884:20260819:111154.285 frontend connection of type "unencrypted" is not allowed by TLSFrontendAccept
```

The Zabbix server is configured to accept only TLS connections, for example with `TLSFrontendAccept=cert`, but the frontend is attempting to connect without encryption.

Enable frontend encryption and configure the required TLS files in `zabbix.conf.php`, as described in step 4 of the Configuration.

Frontend attempts TLS while server expects unencrypted

The Zabbix server log contains an error similar to:

```
61091:20260819:115916.019 End of zbx_tls_accept():FAIL error:'TLS handshake set result code to 1: file ../
```

The exact error message depends on the OpenSSL version.

The Zabbix server is configured to accept only unencrypted connections, for example with `TLSFrontendAccept=unencrypted`, while the frontend is attempting to establish a TLS connection.

Configure the Zabbix server to accept TLS connections by setting `TLSFrontendAccept=cert` and defining `TLSCertFile` and `TLSKeyFile` in `zabbix_server.conf`. Alternatively, disable frontend encryption in `zabbix.conf.php`.

Frontend certificate is not trusted

The Zabbix server log contains an error similar to:

```
62248:20260819:124203.812 End of zbx_tls_accept():FAIL error:'self-signed certificate: TLS handshake set r
```

The frontend presented a certificate, but the Zabbix server does not trust the Certificate Authority (CA) that signed it.

The same problem can also appear in the frontend system information as:

```
fread(): SSL operation failed with code 1. OpenSSL Error messages:

error:0A000418:SSL routines::tlsv1 alert unknown ca [zabbix.php:17 → require_once() → ZBase->run() → ZBase
```

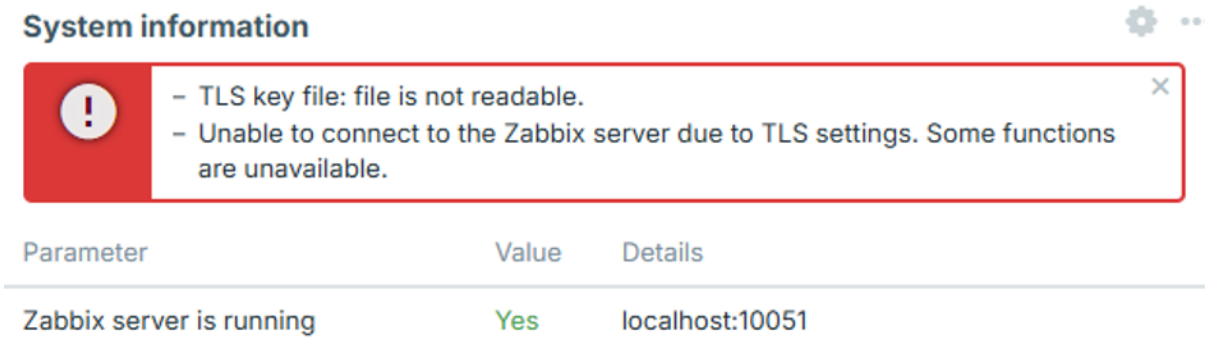
The frontend may also display ”Incorrect response received from Zabbix server ”localhost””.

Verify that the frontend certificate is signed by the CA specified by `TLSCAFile` in `zabbix_server.conf`.

TLS key file is not readable

The frontend may display ”TLS key file: file is not readable.”

It may also display ”Unable to connect to the Zabbix server due to TLS settings. Some functions are unavailable.”.

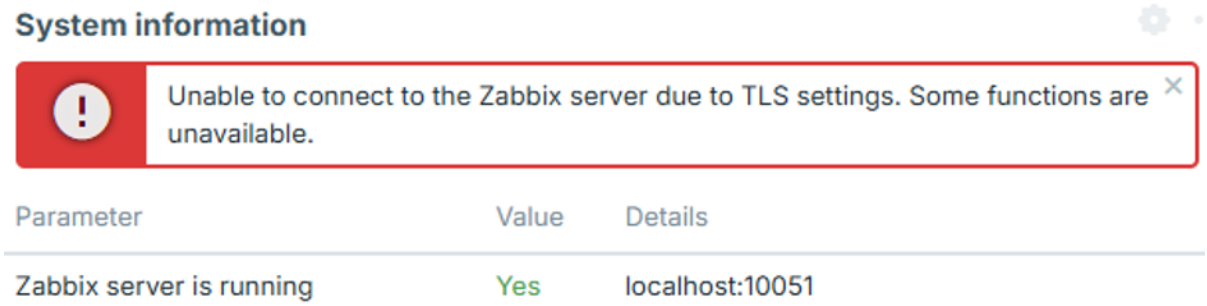


The web server user does not have sufficient permissions to read one or more certificate, key, or CA files configured in `zabbix.conf.php`. This can occur when certificate files are stored in a restricted directory, such as /root/ or a user's home directory.

Move the certificate files to a directory accessible to the web server and set appropriate permissions as described in step 2 of the Configuration.

TLS files are invalid or corrupted

The frontend may display "Unable to connect to the Zabbix server due to TLS settings. Some functions are unavailable.".



This can occur when the configured certificate or key files exist and are readable, but their contents are empty, corrupted, or otherwise invalid. In this case, the frontend does not establish a connection to the Zabbix server, and no corresponding connection attempt appears in the Zabbix server log.

Verify that the certificate and key files contain valid, properly formatted data and that the private key corresponds to the certificate.

Access control

Overview

This section contains best practices for setting up access control in a secure way.

Principle of least privilege

User accounts, at all times, should run with as few privileges as possible. This means that user accounts in Zabbix frontend, database users, or the user for Zabbix server/proxy/agent processes should only have the privileges that are essential for performing the intended functions.

Attention:
Giving extra privileges to the 'zabbix' user will allow it to access configuration files and execute operations that can compromise the infrastructure security.

When configuring user account privileges, Zabbix frontend user types should be considered. Note that although the *Admin* user type has fewer privileges than the *Super Admin* user type, it can still manage configuration and execute custom scripts.

Note:
Some information is available even for non-privileged users. For example, while *Alerts* → *Scripts* is available only for *Super Admin* users, scripts can also be retrieved through Zabbix API. In this case, limiting script permissions and excluding sensitive information from scripts (for example, access credentials) can help avoid exposing sensitive information available in global scripts.

Secure user for Zabbix agent

By default, Zabbix server, proxy, and agent (or agent 2) processes share one `zabbix` user. To prevent Zabbix agent/agent 2 (running on the same machine as server/proxy) from accessing sensitive details in the server/proxy configuration (for example, database credentials), the agent should be run under a different user:

For Zabbix agent:

1. Create a secure group and user (e.g., `zabbix-agent`).
2. Set this user in the agent configuration file User parameter.
3. Restart agent to drop privileges to the new user.

For Zabbix agent 2, the configuration must be applied at the service level, because the agent 2 configuration file does not support the `User` parameter. For an example, see ZBX-26442.

Revoke write access to SSL configuration (Windows)

If you have compiled Zabbix agent on Windows, with OpenSSL located in an unprotected directory (e.g., `C:\zabbix`, `c:\openssl-64bit`, `C:\OpenSSL-Win64-111-static`, or `C:\dev\openssl`), make sure to revoke write access from non-administrator users to this directory. Otherwise, the agent loads SSL settings from a path that can be modified by unprivileged users, resulting in a potential security vulnerability.

Hardening security of Zabbix components

Some functionality can be switched off to harden the security of Zabbix components:

- Global script execution on Zabbix server can be disabled by setting `EnableGlobalScripts=0` in server configuration.
- Global script execution on Zabbix proxy is disabled by default (can be enabled by setting `EnableRemoteCommands=1` in proxy configuration).
- Global script execution on Zabbix agents is disabled by default (can be enabled by adding an `AllowKey=system.run[<command>,*]` or `AllowKeyRegexp=^system\.run\[.*\]$` parameter for each allowed command in agent configuration).
- Access to the configuration of user HTTP authentication is enabled by default; can be disabled by setting `$ZBX_FEATURE_FLAGS['http = false` in frontend configuration.
- Access to the modules configuration is enabled by default; can be disabled by setting `$ZBX_FEATURE_FLAGS['modules_config_ena = false` in frontend configuration.
- Editing of specific media types can be restricted by setting `$ZBX_FEATURE_FLAGS['media_type_denylist'] = [<value1>, <value2>, ...]` in frontend configuration. Possible values: `'email'`, `'script'`, `'sms'`, `'webhook'`, and `'push'`. One or more values can be specified.

Reinstalling the frontend by running `setup.php` is restricted if the `zabbix.conf.php` configuration file exists.

UNC path access on Windows by Zabbix agent

Zabbix agents on Windows follow UNC paths (SMB shares like `\\server\share\file.txt`) in items like `vfs.file.*`, `vfs.dir.*`, `modbus.get`, and `perf_counter*`. This can be a security risk in some contexts.

When Windows is asked to access a UNC path, it tries to authenticate on that server. This means that a malicious request to Zabbix agent can expose the NTLM hash to the requesters server. Users can mitigate this by configuring the AllowKey, AllowKeyRegexp, DenyKey, and DenyKeyRegexp configuration parameters as needed.

Securing MySQL/MariaDB

Overview

This section contains best practices for securing a MySQL/MariaDB database.

> **Note:**
> For a basic setup, see the default MySQL/MariaDB database creation instructions, which include creating the 'zabbix' user with full privileges on the Zabbix database. This user is the database owner that also has the necessary privileges for modifying the database structure when upgrading Zabbix.

To improve security, creating additional database roles and users with minimal privileges is recommended. These roles and users should be configured based on the principle of least privilege, that is, they should only have privileges that are essential for performing the intended functions.

Creating user roles

Create the following roles with the corresponding privileges:

- **zbx_srv** - role for running Zabbix server and proxy:

```
CREATE ROLE 'zbx_srv';
GRANT DELETE, INSERT, SELECT, UPDATE ON zabbix.* TO 'zbx_srv';
FLUSH PRIVILEGES;
```

- **zbx_web** - role for running Zabbix frontend and API:

```
CREATE ROLE 'zbx_web';
GRANT DELETE, INSERT, SELECT, UPDATE ON zabbix.* TO 'zbx_web';
FLUSH PRIVILEGES;
```

- **zbx_bckp** - role for table backup:

```
CREATE ROLE 'zbx_bckp';
GRANT LOCK TABLES, TRIGGER, SELECT ON zabbix.* TO 'zbx_bckp';
GRANT process ON *.* TO 'zbx_bckp';
FLUSH PRIVILEGES;
```

Note:
Table restoration and upgrade should be performed by the database owner.

- **zbx_part** - role with a reduced set of privileges for database partitioning; note that this role can be created only after the database has been created, as it grants privileges on specific database tables:

```
CREATE ROLE 'zbx_part';
GRANT SELECT, ALTER, DROP ON zabbix.history TO 'zbx_part';
GRANT SELECT, ALTER, DROP ON zabbix.history_uint TO 'zbx_part';
GRANT SELECT, ALTER, DROP ON zabbix.history_str TO 'zbx_part';
GRANT SELECT, ALTER, DROP ON zabbix.history_text TO 'zbx_part';
GRANT SELECT, ALTER, DROP ON zabbix.history_log TO 'zbx_part';
GRANT SELECT, ALTER, DROP ON zabbix.trends TO 'zbx_part';
GRANT SELECT, ALTER, DROP ON zabbix.trends_uint TO 'zbx_part';
-- For MariaDB: skip the next line (GRANT session_variables_admin ON *.* TO 'zbx_part';)
GRANT session_variables_admin ON *.* TO 'zbx_part';
GRANT SELECT ON zabbix.dbversion TO 'zbx_part';
GRANT SELECT, DELETE ON zabbix.housekeeper TO 'zbx_part';
FLUSH PRIVILEGES;
```

After creating roles, they can be assigned to users.

Assigning user roles

To assign the created user roles, create users and assign the relevant roles to them. Replace <user>, <host>, <role>, and <password> with the appropriate values. Ensure that <password> is strong and unique.

```
CREATE USER '<user>'@'<host>' IDENTIFIED BY '<password>';
GRANT '<role>' TO '<user>'@'<host>';
SET DEFAULT ROLE '<role>' TO '<user>'@'<host>';
-- For MariaDB: SET DEFAULT ROLE '<role>' FOR '<user>'@'<host>'
FLUSH PRIVILEGES;
```

For example, to create and assign the role for running Zabbix server and proxy:

```
CREATE USER 'usr_srv'@'localhost' IDENTIFIED BY 'strong-password';
GRANT 'zbx_srv' TO 'usr_srv'@'localhost';
SET DEFAULT ROLE ALL TO 'usr_srv'@'localhost';
FLUSH PRIVILEGES;
```

Securing PostgreSQL/TimescaleDB

Overview

This section contains best practices for securing a PostgreSQL database.

Note:
For a basic setup, see the default PostgreSQL database creation instructions, which include creating the 'zabbix' user with full privileges on the Zabbix database. This user is the database owner that also has the necessary privileges for modifying the database structure when upgrading Zabbix.

To improve security, configuring a secure schema usage pattern, as well as creating additional database roles and users with minimal privileges is recommended. These roles and users should be configured based on the principle of least privilege, that is, they should only have privileges that are essential for performing the intended functions.

Database setup

Create the user that will be the database owner, and create the Zabbix database; the database owner is the user that is specified on database creation:

```
createuser -U postgres -h localhost --pwprompt usr_owner
createdb -U postgres -h localhost -O usr_owner -E Unicode -T template0 zabbix
```

Note:
A clean install or upgrade of the database has to be performed by the database owner. This is because the right to drop a database object or alter its definition is a privilege that is inherent to the database owner and cannot be granted or revoked.

Attention:
The following commands on this page must be executed while the connection to PostgreSQL is made specifically to the `zabbix` database.

Create the `zabbix` schema and set the database owner (`usr_owner`) to be the owner of this schema:

```
CREATE SCHEMA zabbix AUTHORIZATION usr_owner;
```

Configure a secure schema usage pattern:

```
REVOKE CREATE ON SCHEMA public FROM PUBLIC;
REVOKE ALL ON DATABASE zabbix FROM PUBLIC;
-- Note: search_path should point to the "zabbix" schema:
ALTER ROLE ALL IN DATABASE zabbix SET search_path = "zabbix";
```

After setting up the database, proceed to creating user roles.

Creating user roles

Create the following roles with the corresponding privileges:

- **zbx_srv** - role for running Zabbix server and proxy:

```
CREATE ROLE zbx_srv;
GRANT CONNECT ON DATABASE zabbix TO zbx_srv;
GRANT USAGE ON SCHEMA zabbix TO zbx_srv;
ALTER DEFAULT PRIVILEGES FOR ROLE usr_owner IN SCHEMA zabbix GRANT DELETE, INSERT, SELECT, UPDATE ON TABLE
ALTER DEFAULT PRIVILEGES FOR ROLE usr_owner IN SCHEMA zabbix GRANT SELECT, UPDATE, USAGE ON sequences TO z
```

- **zbx_web** - role for running Zabbix frontend and API:

```
CREATE ROLE zbx_web;
GRANT CONNECT ON DATABASE zabbix TO zbx_web;
GRANT USAGE ON SCHEMA zabbix TO zbx_web;
ALTER DEFAULT PRIVILEGES FOR ROLE usr_owner IN SCHEMA zabbix GRANT DELETE, INSERT, SELECT, UPDATE ON TABLE
ALTER DEFAULT PRIVILEGES FOR ROLE usr_owner IN SCHEMA zabbix GRANT SELECT, UPDATE, USAGE ON sequences TO z
```

- **zbx_bckp** - role for table backup:

```
CREATE ROLE zbx_bckp;
GRANT CONNECT ON DATABASE zabbix TO zbx_bckp;
GRANT USAGE ON SCHEMA zabbix TO zbx_bckp;
ALTER DEFAULT PRIVILEGES FOR ROLE usr_owner IN SCHEMA zabbix GRANT SELECT ON TABLES TO zbx_bckp;
ALTER DEFAULT PRIVILEGES FOR ROLE usr_owner IN SCHEMA zabbix GRANT SELECT, UPDATE, USAGE ON sequences TO z
```

> **Note:**
> Table restoration is possible only by the database owner.

After creating roles, they can be assigned to users.

Assigning user roles

To assign the created user roles, create users and assign the relevant roles to them. Replace <user>, <role>, and <password> with the appropriate values. Ensure that <password> is strong and unique.

```
CREATE USER <user> WITH ENCRYPTED password '<password>';
GRANT <role> TO <user>;
```

For example, to create and assign the role for running Zabbix server and proxy:

```
CREATE USER usr_srv WITH ENCRYPTED password 'strong-password';
GRANT zbx_srv TO usr_srv;
```

Database partitioning with TimescaleDB

Database partitioning is facilitated by TimescaleDB. To utilize TimescaleDB, Zabbix server requires database owner privileges.

If the PostgreSQL zabbix schema has already been created in the zabbix database, you can enable TimescaleDB with the following command:

```
echo "CREATE EXTENSION IF NOT EXISTS timescaledb WITH SCHEMA zabbix CASCADE;" | sudo -u postgres psql zabb
```

Cryptography

Overview

This section contains best practices for setting up cryptography in a secure way.

Setting up SSL for Zabbix frontend

On RHEL-based systems, install the mod_ssl package:

```
dnf install mod_ssl
```

Create a directory for SSL keys:

```
mkdir -p /etc/httpd/ssl/private
chmod 700 /etc/httpd/ssl/private
```

Create the SSL certificate:

```
openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout /etc/httpd/ssl/private/apache-selfsigned.key -
```

Fill out the prompts appropriately. The most important line is the one that requests the Common Name. You must enter the domain name you want to be associated with your server. You can enter the public IP address instead if you do not have a domain name.

```
Country Name (2 letter code) [XX]:
State or Province Name (full name) []:
Locality Name (eg, city) [Default City]:
Organization Name (eg, company) [Default Company Ltd]:
Organizational Unit Name (eg, section) []:
Common Name (eg, your name or your server's hostname) []:example.com
Email Address []:
```

Edit the Apache SSL configuration file (/etc/httpd/conf.d/ssl.conf):

```
DocumentRoot "/usr/share/zabbix"
ServerName example.com:443
SSLCertificateFile /etc/httpd/ssl/apache-selfsigned.crt
SSLCertificateKeyFile /etc/httpd/ssl/private/apache-selfsigned.key
```

Restart the Apache service to apply the changes:

```
systemctl restart httpd.service
```

Securing web server

Overview

This section contains best practices for setting up the web server in a secure way.

Enforcing root URL redirect to Zabbix SSL

On RHEL-based systems, add a virtual host to Apache configuration (`/etc/httpd/conf/httpd.conf`) and set a permanent redirect for document root to Zabbix SSL URL. Note that `example.com` should be replaced with the actual name of the server.

```
### Add lines:

<VirtualHost *:*>
    ServerName example.com
    Redirect permanent / https://example.com
</VirtualHost>
```

Restart the Apache service to apply the changes:

```
systemctl restart httpd.service
```

Enabling HTTP Strict Transport Security (HSTS) on the web server

To protect Zabbix frontend against protocol downgrade attacks, we recommend enabling the HSTS policy on the web server.

To enable the HSTS policy for your Zabbix frontend in Apache configuration, follow these steps:

1. Locate your virtual host's configuration file:

- `/etc/httpd/conf/httpd.conf` on RHEL-based systems
- `/etc/apache2/sites-available/000-default.conf` on Debian/Ubuntu

2. Add the following directive to your virtual host's configuration file:

```
<VirtualHost *:*>
    Header set Strict-Transport-Security "max-age=31536000"
</VirtualHost>
```

3. Restart the Apache service to apply the changes:

```
### On RHEL-based systems:
systemctl restart httpd.service

### On Debian/Ubuntu
systemctl restart apache2.service
```

Enforcing Secure and SameSite session cookies in Zabbix

When configuring Zabbix, it is essential to enforce secure and SameSite attributes for session cookies to enhance security and prevent cross-site request forgery (CSRF) attacks. However, enforcing `SameSite=Strict` may cause issues in certain scenarios, such as:

- Dashboard URL widgets displaying "user not logged in" when embedding same-domain iframes.
- Users accessing the dashboard via HTTP instead of HTTPS may face login issues.
- Inability to share URLs to specific Zabbix menu sections or hosts.

To mitigate these issues, users should have a way to adjust the SameSite policy.

1. Secure cookies

Setting the `secure` flag ensures that cookies are only transmitted over HTTPS, preventing exposure over unencrypted connections.

To enable secure cookies in Zabbix, add or modify the following setting in the web server configuration:

For Apache:

```
Header always edit Set-Cookie ^(.*)$ $1;Secure
```

For Nginx:

```
proxy_cookie_path / "/; Secure";
```

Ensure that your Zabbix frontend is accessed via HTTPS; otherwise, cookies with the Secure flag will not be sent.

2. Configuring the SameSite attribute

Web server settings can also enforce the SameSite attribute:

For Apache:

```
<IfModule mod_headers.c>
    Header onsuccess edit Set-Cookie (.*) "$1; SameSite=Strict"
</IfModule>
```

For Nginx (version 1.19.3+):

```
proxy_cookie_flags ~ samesite=Strict; # Replace ~ with 'zbx_session' for specificity
```

Enabling Content Security Policy (CSP) on the web server

To protect Zabbix frontend against Cross Site Scripting (XSS), data injection, and similar attacks, we recommend enabling Content Security Policy on the web server. To do so, configure the web server to return the HTTP header.

Attention:
The following CSP header configuration is only for the default Zabbix frontend installation and for cases when all content originates from the site's domain (excluding subdomains). A different CSP header configuration may be required if you are, for example, configuring the *URL* widget to display content from the site's subdomains or external domains, switching from *OpenStreetMap* to another map engine, or adding external CSS or widgets. If you're using the Duo Universal Prompt multi-factor authentication method, make sure to add "duo.com" to the CSP directive in your virtual host's configuration file.

To enable CSP for your Zabbix frontend in Apache configuration, follow these steps:

1. Locate your virtual host's configuration file:

- `/etc/httpd/conf/httpd.conf` on RHEL-based systems
- `/etc/apache2/sites-available/000-default.conf` on Debian/Ubuntu

2. Add the following directive to your virtual host's configuration file:

```
<VirtualHost *:*>
    Header set Content-Security-Policy: "default-src 'self' *.openstreetmap.org; script-src 'self' 'unsafe
</VirtualHost>
```

3. Restart the Apache service to apply the changes:

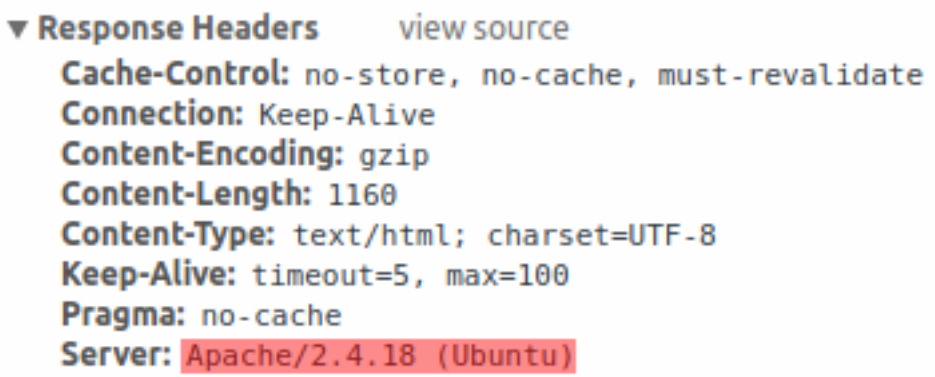
```
### On RHEL-based systems:
systemctl restart httpd.service

### On Debian/Ubuntu
systemctl restart apache2.service
```

Disabling web server information exposure

To improve security, disabling all web server signatures is recommended.

By default, the web server is exposing the software signature:



The signature can be disabled by adding the following parameters to the Apache configuration file:

```
ServerSignature Off
ServerTokens Prod
```

PHP signature (X-Powered-By HTTP header) can be disabled by changing the `php.ini` configuration file (by default, the signature is disabled):

```
expose_php = Off
```

Web server restart is required for configuration file changes to be applied.

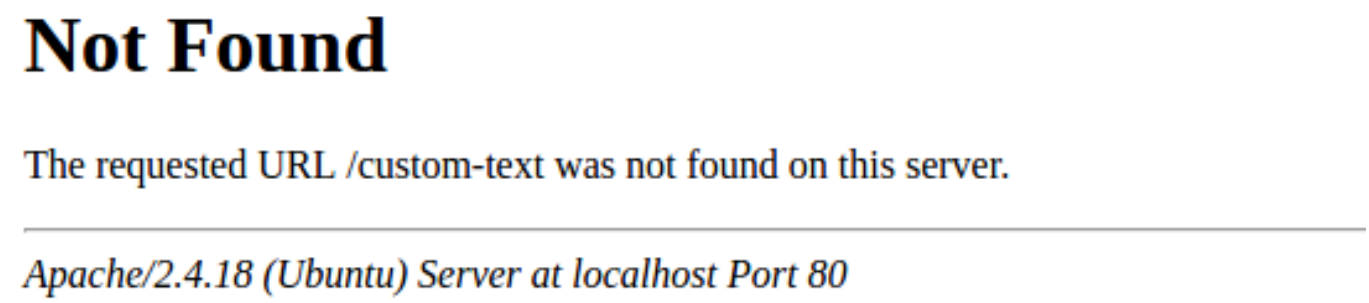
For additional security, you can use the *mod_security* tool with Apache (package *libapache2-mod-security2*). This tool allows removing the server signature instead of removing only the version from the server signature. The server signature can be changed to any value by setting "SecServerSignature" to any desired value after installing *mod_security*.

Please refer to the documentation of your web server to find help on how to remove/change software signatures.

Disabling web server default error pages

To avoid information exposure, disabling the default error pages is recommended.

By default, a web server uses built-in error pages:



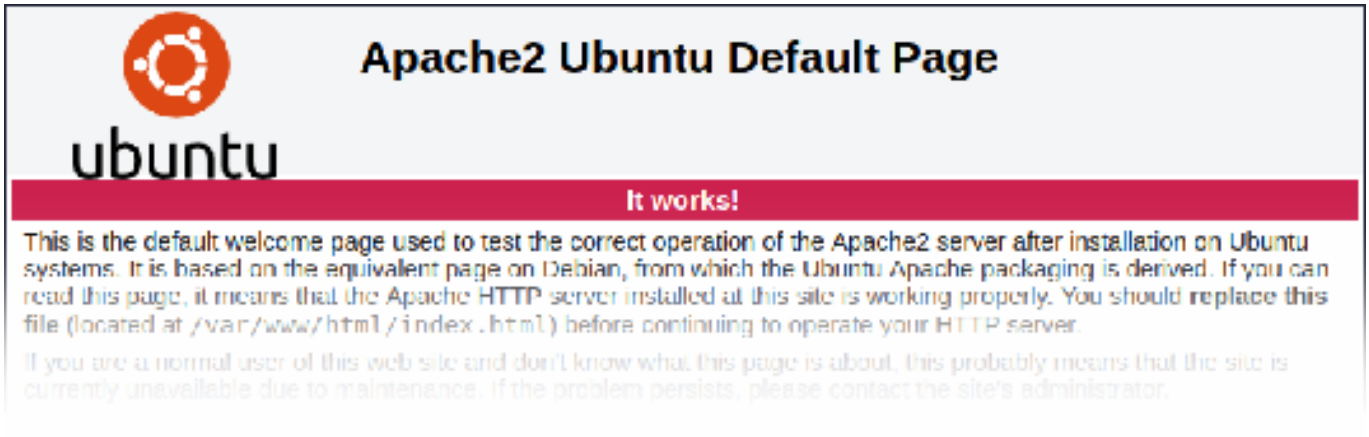
These default error pages should be replaced/removed. For example, the "ErrorDocument" directive can be used to define a custom error page/text for the Apache web server.

Please refer to the documentation of your web server to find help on how to replace/remove default error pages.

Removing web server test page

To avoid information exposure, removing the web server test page is recommended.

By default, the Apache web server webroot contains the `index.html` test page:



Please refer to the documentation of your web server to find help on how to remove default test pages.

Set X-Frame-Options HTTP response header

By default, Zabbix is configured with the *Use X-Frame-Options HTTP header* parameter set to `SAMEORIGIN`. This means that content can only be loaded in a frame that has the same origin as the page itself.

Zabbix frontend elements that pull content from external URLs (namely, the URL dashboard widget) display retrieved content in a sandbox with all sandboxing restrictions enabled.

These settings enhance the security of the Zabbix frontend and provide protection against XSS and clickjacking attacks. *Super admin* users can modify the *Use iframe sandboxing* and *Use X-Frame-Options HTTP header* parameters as needed. Please carefully weigh the risks and benefits before changing default settings. Turning iframe sandboxing or X-Frame-Options HTTP header off completely is not recommended.

Hiding the file with list of common passwords

To increase the complexity of password brute force attacks, limiting access to the `ui/data/top_passwords.txt` file is recommended. This file contains a list of the most common and context-specific passwords and prevents users from setting such passwords (if the *Avoid easy-to-guess passwords* parameter is enabled in the password policy).

To limit access to the `top_passwords.txt` file, modify your web server configuration.

On Apache, file access can be limited using the `.htaccess` file:

```
<Files "top_passwords.txt">
    Order Allow,Deny
    Deny from all
</Files>
```

On NGINX, file access can be limited using the `location` directive:

```
location = /data/top_passwords.txt {
    deny all;
    return 404;
}
```

Storage of secrets

Overview Zabbix can be configured to retrieve sensitive information from a secure vault. The following secret management services are supported: HashiCorp Vault KV Secrets Engine - Version 2, CyberArk Vault CV12.

Secrets can be used for retrieving:

- user macro values
- database access credentials

Zabbix provides read-only access to the secrets in a vault, assuming that secrets are managed by someone else.

For information about specific vault provider configuration, see:

- HashiCorp configuration
- CyberArk configuration

Caching of secret values By default, vault secret macro values are retrieved by Zabbix server on every refresh of configuration data and then stored in the configuration cache. Zabbix proxy receives values of vault secret macros from Zabbix server on each configuration sync and stores them in its own configuration cache.

> **Attention:**
> Encryption must be enabled between Zabbix server and proxy; otherwise, a server warning message is logged.

It is also possible to configure that macro values are retrieved by Zabbix server and Zabbix proxy independently.

To manually trigger refresh of cached secret values from a vault, use the 'secrets_reload' command-line option.

For Zabbix frontend database credentials caching is disabled by default, but can be enabled by setting the option `$DB['VAULT_CACHE'] = true` in the frontend configuration file. The credentials will be stored in a local cache using the filesystem temporary file directory. The web server must allow writing in a private temporary folder (for example, for Apache the configuration option `PrivateTmp=True` must be set). To control how often the data cache is refreshed/invalidated, use the `ZBX_DATA_CACHE_TTL` constant .

TLS configuration To configure TLS for communication between Zabbix components and the vault, add a certificate signed by a certificate authority (CA) to the system-wide default CA store. To use another location, specify the directory in the SSLCALocation Zabbix server/proxy configuration parameter, place the certificate file inside that directory, then run the CLI command:

```
c_rehash .
```

CyberArk configuration

This section explains how to configure Zabbix to retrieve secrets from CyberArk Vault CV12.

The vault should be installed and configured as described in the official CyberArk documentation.

To learn about configuring TLS in Zabbix, see *Storage of secrets*.

Database credentials

Access to a secret with database credentials is configured for each Zabbix component separately.

Server and proxies

To obtain database credentials from the vault for Zabbix server or proxy, specify the following configuration parameters in the configuration file:

- `Vault` - which vault provider should be used;
- `VaultURL` - vault server HTTP(S) URL;
- `VaultDBPath` - query to the vault secret containing database credentials which will be retrieved by keys "Content" and "UserName" (this option can only be used if DBUser and DBPassword are not specified);
- `VaultTLSCertFile`, `VaultTLSKeyFile` - SSL certificate and key file names; setting up these options is not mandatory, but highly recommended;
- `VaultPrefix` - custom prefix for the vault path or query, depending on the vault; if not specified, the most suitable default will be used.

Attention:
`Vault`, `VaultURL`, `VaultTLSCertFile`, `VaultTLSKeyFile`, and `VaultPrefix` configuration parameters are also used for vault authentication when processing secret vault macros by Zabbix server (and Zabbix proxy, if configured). Zabbix server and proxies will not open vault secret macros that contain DB credentials from VaultDBPath.

Zabbix server and Zabbix proxy read the vault-related configuration parameters from `zabbix_server.conf` and `zabbix_proxy.conf` files upon startup.

Example

1. In `zabbix_server.conf`, specify the following parameters:

```
Vault=CyberArk
VaultURL=https://127.0.0.1:1858
VaultDBPath=AppID=zabbix_server&Query=Safe=passwordSafe;Object=zabbix_server_database
VaultTLSCertFile=cert.pem
VaultTLSKeyFile=key.pem
VaultPrefix=/AIMWebService/api/Accounts?
```

2. Zabbix will send the following API request to the vault:

```
curl \
--header "Content-Type: application/json" \
--cert cert.pem \
--key key.pem \
https://127.0.0.1:1858/AIMWebService/api/Accounts?AppID=zabbix_server&Query=Safe=passwordSafe;Object=zabbi
```

3. The vault response will contain the keys "Content" and "UserName":

```
{
    "Content": <password>,
    "UserName": <username>,
    "Address": <address>,
    "Database": <Database>,
    "PasswordChangeInProcess":<PasswordChangeInProcess>
}
```

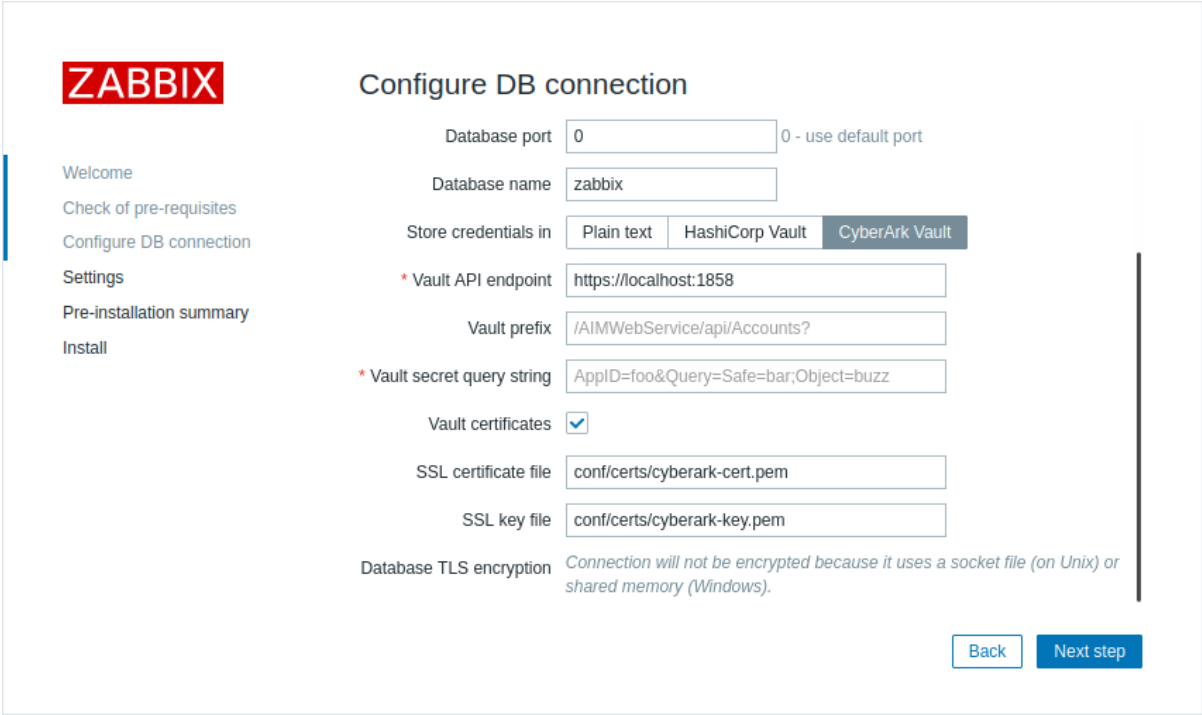
4. As a result, Zabbix will use the following credentials for database authentication:

- Username: <username>
- Password: <password>

Frontend

To obtain database credentials from the vault for Zabbix frontend, specify the following parameters during frontend installation.

1. At the *Configure DB Connection* step, set the *Store credentials in* parameter to *CyberArk Vault*.



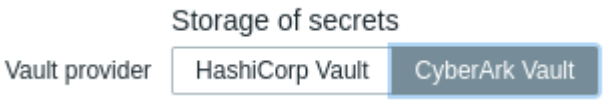
2. Then, fill in the additional parameters:

Parameter	Mandatory	Default value	Description
Vault API endpoint	yes	https://localhost:1858	Specify the URL for connecting to the vault in the format `scheme://host:port`
Vault prefix	no	/AIMWebService/api/Accounts?	Provide a custom prefix for the vault path or query. If not specified, the default is used.
Vault secret query string	yes		A query, which specifies from where database credentials should be retrieved. Example: `AppID=foo&Query=Safe=bar;Object=buzz`
Vault certificates	no		After marking the checkbox, additional parameters will appear, allowing you to configure client authentication. While this parameter is optional, it is highly recommended to enable it for communication with the CyberArk Vault.
SSL certificate file	no	conf/certs/cyberark-cert.pem	Path to the SSL certificate file. The file must be in PEM format. If the certificate file also contains the private key, leave the SSL key file parameter empty.
SSL key file	no	conf/certs/cyberark-key.pem	Name of the SSL private key file used for client authentication. The file must be in PEM format.

User macro values

To use CyberArk Vault for storing *Vault secret* user macro values, make sure that:

- Zabbix server is configured to work with CyberArk Vault;
- the *Vault provider* parameter in *Administration > General > Other* is set to "CyberArk Vault".



Note:
Zabbix server (and Zabbix proxy, if configured) require access to *Vault secret* macro values from the vault. Zabbix frontend does not need such access.

The macro value should contain a query (as `query:key`).

See *Vault secret macros* for detailed information on macro value processing by Zabbix.

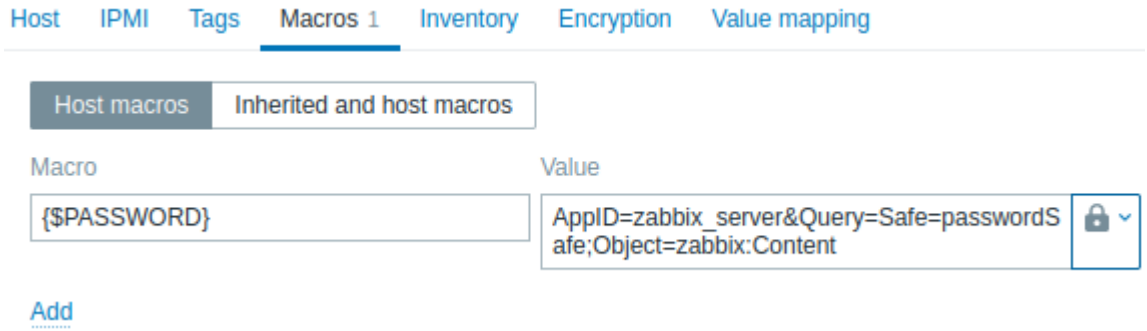
Query syntax

The colon symbol (`:`) is reserved for separating the query from the key.

If a query itself contains a forward slash or a colon, these symbols should be URL-encoded (`/` is encoded as `%2F`, `:` is encoded as `%3A`).

Example

1. In Zabbix, add a user macro `{$PASSWORD}` of type *Vault secret* and with the value `AppID=zabbix_server&Query=Safe=passwordS`



2. Zabbix will send the following API request to the vault:

```
curl \
--header "Content-Type: application/json" \
--cert cert.pem \
--key key.pem \
https://127.0.0.1:1858/AIMWebService/api/Accounts?AppID=zabbix_server&Query=Safe=passwordSafe;Object=zabbi
```

3. The vault response will contain the key `Content`:

```
{
    "Content": <password>,
    "UserName": <username>,
    "Address": <address>,
    "Database" :<Database>,
    "PasswordChangeInProcess":<PasswordChangeInProcess>
}
```

4. As a result, Zabbix will resolve the macro `{$PASSWORD}` to the value - `<password>`

Update existing configuration

To update an existing configuration for retrieving secrets from a CyberArk Vault:

1. Update the Zabbix server or proxy configuration file parameters as described in the *Database credentials* section.
2. Update the Zabbix frontend configuration file parameters:

```
$DB['VAULT']                    = 'CyberArk';
$DB['VAULT_URL']                = 'https://127.0.0.1:1858';
$DB['VAULT_DB_PATH']            = 'AppID=foo&Query=Safe=bar;Object=buzz';
$DB['VAULT_TOKEN']              = '';
$DB['VAULT_CERT_FILE']          = 'conf/certs/cyberark-cert.pem';
$DB['VAULT_KEY_FILE']           = 'conf/certs/cyberark-key.pem';
$DB['VAULT_PREFIX']             = '';
```

3. Configure user macros as described in the *User macro values* section, if necessary.

To update an existing configuration for retrieving secrets from a HashiCorp Vault, see *HashiCorp configuration*.

HashiCorp configuration

Overview

This section explains how to configure Zabbix for retrieving secrets from HashiCorp Vault KV Secrets Engine - Version 2.

The vault should be deployed and configured as described in the official HashiCorp documentation.

To learn about configuring TLS in Zabbix, see *Storage of secrets*.

Retrieving database credentials

To retrieve a secret with database credentials successfully, it is required to configure both:

- Zabbix server/proxy
- Zabbix frontend

Server/proxy

To configure Zabbix server or proxy, specify the following configuration parameters in the configuration file:

- `Vault` - which vault provider should be used.
- `VaultToken` - vault authentication token (see Zabbix server/proxy configuration file for details); cannot be used with `VaultAppRoleID`.
- `VaultAppRoleID` and `VaultAppSecretID` - vault authentication role ID and secret ID (see Zabbix server/proxy configuration file for details); cannot be used with `VaultToken`.
- `VaultURL` - vault server HTTP(S) URL.
- `VaultDBPath` - path to the vault secret containing database credentials (this option can only be used if DBUser and DBPassword are not specified). Zabbix server or proxy will retrieve the credentials by keys "password" and "username".
- `VaultPrefix` - custom prefix for the vault path or query, depending on the vault; if not specified, the most suitable default will be used. Note that the prefix must start and end with `/`.

Attention:
`Vault`, `VaultToken` (or `VaultAppRoleID`/`VaultAppSecretID`), `VaultURL`, and `VaultPrefix` configuration parameters are also used for vault authentication when processing secret vault macros by Zabbix server (and Zabbix proxy, if configured). Zabbix server and proxies will not open vault secret macros that contain DB credentials from VaultDBPath.

Using different tokens for different proxies is strongly recommended.

Zabbix server and Zabbix proxy read the vault-related configuration parameters from `zabbix_server.conf` and `zabbix_proxy.conf` upon startup. Additionally, Zabbix server and Zabbix proxy will read the `VAULT_TOKEN` environment variable once during startup and will unset it so that it would not be available through forked scripts; it is an error if both `VaultToken` and `VAULT_TOKEN` parameters contain a value.

Example

1. Configure the Zabbix server parameters.

Attention:
Do not use the example token, role ID, or secret ID values from the documentation in production environments.

- If using `VAULT_TOKEN` environment variable, specify the following parameters in Zabbix server configuration file (`zabbix_server.conf`) and set the variable:

```
Vault=HashiCorp
VaultURL=https://127.0.0.1:8200
VaultDBPath=database
VaultPrefix=/v1/secret/data/zabbix/
```

Set the `VAULT_TOKEN` environment variable for the Zabbix server before it starts:

```
export VAULT_TOKEN=hvs.CAESIIG_PILmULFYOsEyWHxkZ2mF2a8VPKNLE8eHqd4autYGGh4KHGh2cy5aeTY0NFNSaUp3ZnpWbDF1RUN
zabbix_server -c /etc/zabbix/zabbix_server.conf
```

- If using `VaultToken` parameter, specify the following parameters in Zabbix server configuration file:

Attention:
If you use the `VaultToken` parameter, the token is stored in plain text in the configuration file. Set the file permissions to 0600 so that only the user running Zabbix server or proxy can read it.

```
Vault=HashiCorp
VaultToken=hvs.CAESIIG_PILmULFYOsEyWHxkZ2mF2a8VPKNLE8eHqd4autYGGh4KHGh2cy5aeTY0NFNSaUp3ZnpWbDF1RUNjUkNTZEg
VaultURL=https://127.0.0.1:8200
```

```
VaultDBPath=database
VaultPrefix=/v1/secret/data/zabbix/
```

- If using `VaultAppRoleID` parameter, specify the following parameters in Zabbix server configuration file:

> **Attention:**
> If you use the `VaultAppSecretID` parameter, the secret ID is stored in plain text in the configuration file. Set the file permissions to 0600 so that only the user running Zabbix server or proxy can read it.

```
Vault=HashiCorp
VaultAppRoleID=6f1567f1-538d-7d70-e8e4-5d3e59471b2e
VaultAppSecretID=89dac8b1-f7d6-96a7-2ecf-f94d2de3d6f1
VaultURL=https://127.0.0.1:8200
VaultDBPath=database
VaultPrefix=/v1/secret/data/zabbix/
```

2. Run the following CLI commands to create the required secret in the vault:

```
#### Enable "secret/" mount point if not already enabled; note that "kv-v2" must be used.
vault secrets enable -path=secret/ kv-v2

#### Put new secrets with keys username and password under mount point "secret/" and path "zabbix/database
vault kv put -mount=secret zabbix/database username=zabbix password=<password>

#### Test that secret is successfully added.
vault kv get secret/zabbix/database

#### Finally test with Curl; note that "data" need to be manually added after mount point and "/v1" before
curl --header "X-Vault-Token: <VaultToken>" https://127.0.0.1:8200/v1/secret/data/zabbix/database
```

3. As a result, Zabbix server will retrieve the following credentials for database authentication:

- Username: zabbix
- Password: <password>

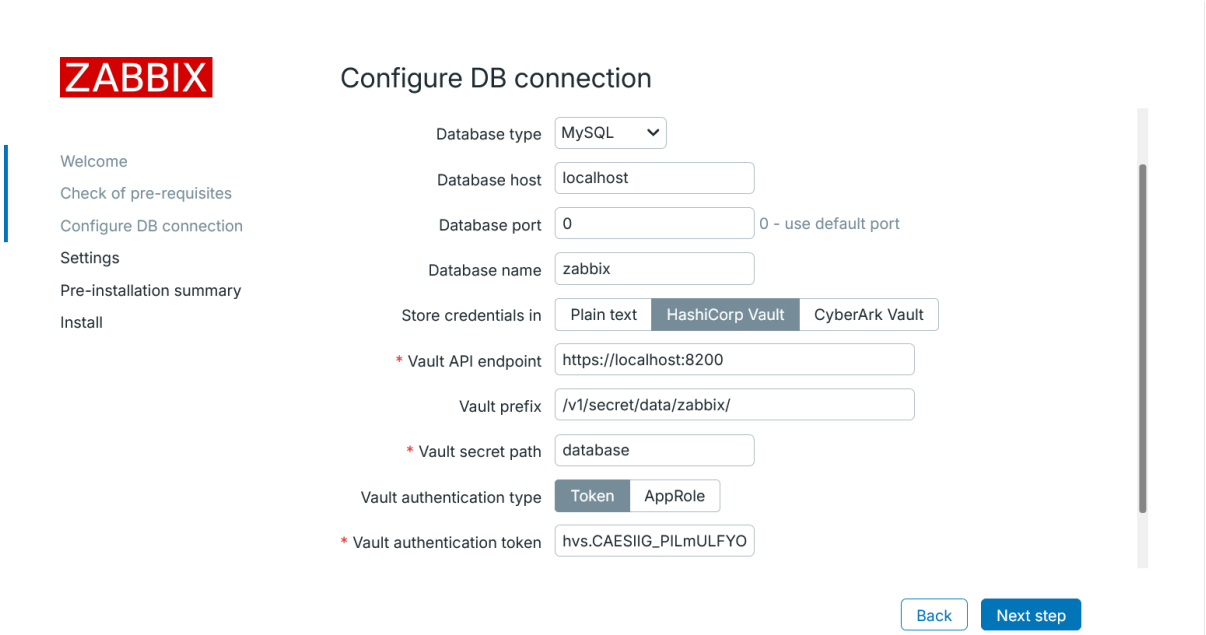
Frontend

Zabbix frontend can be configured to retrieve database credentials from the vault either during frontend installation or by updating the frontend configuration file (`zabbix.conf.php`).

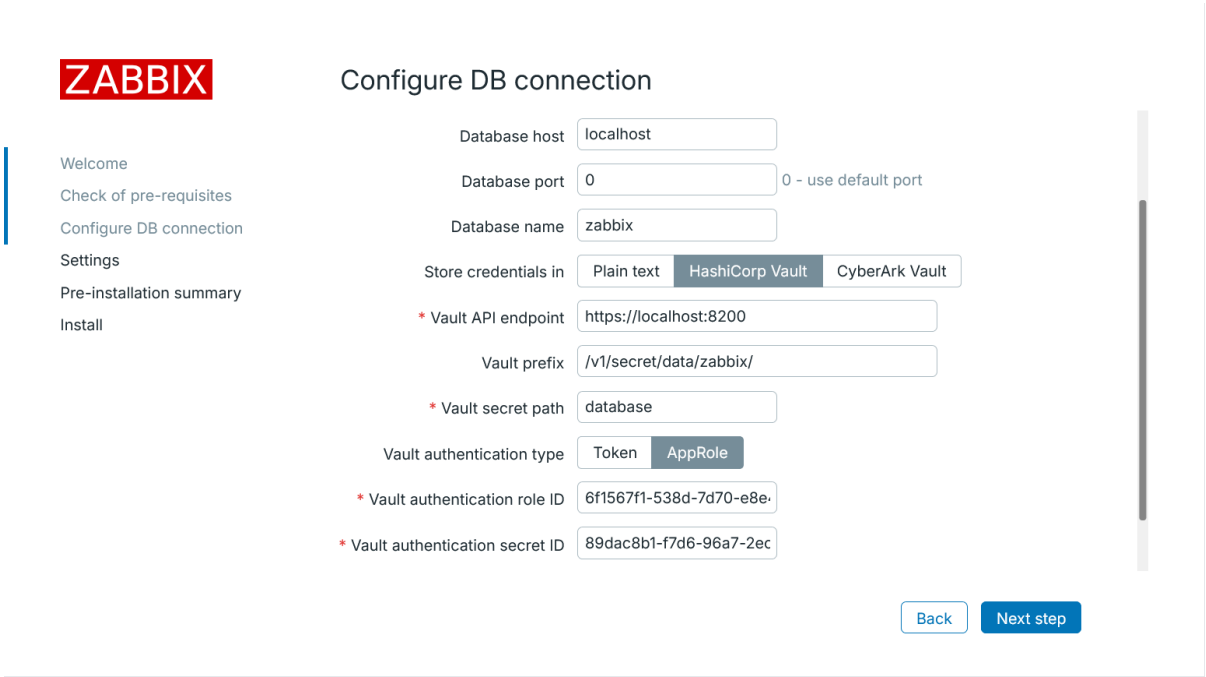
> **Attention:**
> If the vault credentials have changed since the previous frontend installation, update existing configuration.

During **frontend installation** the connection parameters must be specified at the *Configure DB Connection* step.

Configuration example using token authentication:



Configuration example using AppRole authentication:



Parameter	Mandatory	Default value	Description
Store credentials in	no		Select **HashiCorp Vault**.
Vault API endpoint	yes	https://localhost:8200	Specify the URL for connecting to the vault in the format `scheme://host:port`
Vault prefix	no	/v1/secret/data/	Provide a custom prefix for the vault path or query. If not specified, the default is used. Example: `/v1/secret/data/zabbix/`
Vault secret path	yes		A path to the secret from where credentials for the database shall be retrieved by the keys ”password” and ”username”. Example: `database`
Vault authentication type	no		Select vault authentication type: **Token** - use token authentication; **AppRole** - use AppRole authentication method.

Parameter	Mandatory	Default value	Description
Vault authentication token	yes		Provide an authentication token for read-only access to the secret path. See HashiCorp documentation for information about creating tokens and vault policies. This field is only available if *Token* is selected as the vault authentication type.
Vault authentication role ID	yes		Provide the role ID of the AppRole. See HashiCorp documentation for information about AppRoles. This field is only available if *AppRole* is selected as the vault authentication type.
Vault authentication secret ID	yes		Provide the secret ID issued against the AppRole. See HashiCorp documentation for information about AppRoles. This field is only available if *AppRole* is selected as the vault authentication type.

Retrieving user macro values

To use HashiCorp Vault for storing *Vault secret* user macro values, make sure that:

- Zabbix server/proxy is configured to work with HashiCorp Vault.
- The *Vault provider* parameter in *Administration > General > Other* is set to "HashiCorp Vault" (default).



Note:
Zabbix server (and Zabbix proxy, if configured) require access to *Vault secret* macro values from the vault. Zabbix frontend does not need such access.

The macro value should contain a reference path (as `path:key`, for example, `macros:password`). The authentication token specified during Zabbix server/proxy configuration (by the `VaultToken` parameter) must provide read-only access to this path.

See *Vault secret macros* for detailed information on macro value processing by Zabbix.

Path syntax

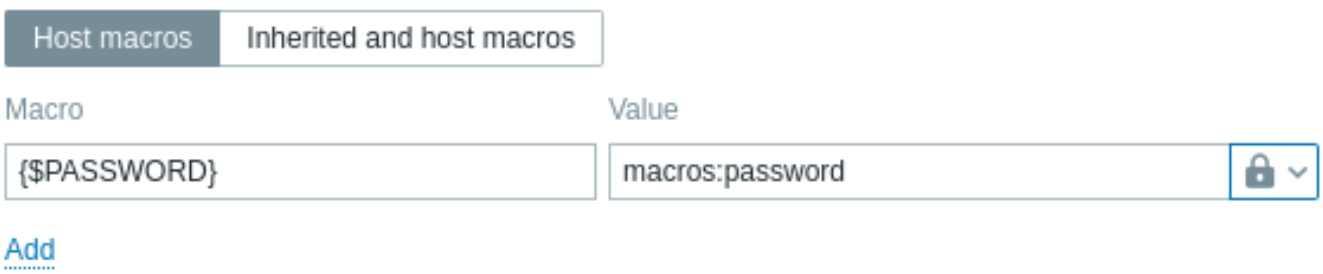
The symbols forward slash ("/") and colon (":") are reserved.

A forward slash can only be used to separate a mount point from a path (e.g., *secret/zabbix* where the mount point is "secret" and the path is "zabbix"). In the case of Vault macros, a colon can only be used to separate a path/query from a key.

It is possible to URL-encode the forward slash and colon symbols if there is a need to create a mount point with the name that is separated by a forward slash (e.g., *foo/bar/zabbix*, where the mount point is "foo/bar" and the path is "zabbix", can be encoded as "foo%2Fbar/zabbix") and if a mount point name or path need to contain a colon.

Example

1. In Zabbix, add a user macro {$PASSWORD} of type "Vault secret" and with the value `macros:password`



2. Run the following CLI commands to create required secret in the vault:

```
#### Enable "secret/" mount point if not already enabled; note that "kv-v2" must be used.
vault secrets enable -path=secret/ kv-v2

#### Put new secret with key "password" under mount point "secret/" and path "zabbix/macros".
vault kv put -mount=secret zabbix/macros password=<password>

#### Test that secret is successfully added.
vault kv get secret/zabbix/macros

#### Finally test with Curl; note that "data" need to be manually added after mount point and "/v1" before
curl --header "X-Vault-Token: <VaultToken>" https://127.0.0.1:8200/v1/secret/data/zabbix/macros
```

3. As a result, Zabbix will resolve the macro {$PASSWORD} to the value: <password>

Updating existing configuration

To update an existing configuration for retrieving secrets from a HashiCorp Vault:

1. Update the Zabbix server or proxy configuration file parameters as described in the *Database credentials* section.
2. Update the Zabbix frontend configuration file parameters:

Using Vault token:

```
$DB['VAULT']                    = 'HashiCorp';
$DB['VAULT_URL']                = 'https://localhost:8200';
$DB['VAULT_DB_PATH']            = 'database';
$DB['VAULT_TOKEN']              = 'hvs.CAESIIG_PILmULFYOsEyWHxkZ2mF2a8VPKNLE8eHqd4autYGGh4KHGh2cy5aeTY0NFN
$DB['VAULT_PREFIX']             = '/v1/secret/data/zabbix/';
```

Or, using Vault AppRole method:

```
$DB['VAULT']                    = 'HashiCorp';
$DB['VAULT_URL']                = 'https://localhost:8200';
$DB['VAULT_DB_PATH']            = 'database';
$DB['VAULT_APP_ROLE_ID']        = '6f1567f1-538d-7d70-e8e4-5d3e59471b2e';
$DB['VAULT_APP_SECRET_ID']      = '89dac8b1-f7d6-96a7-2ecf-f94d2de3d6f1';
$DB['VAULT_PREFIX']             = '/v1/secret/data/zabbix/';
```

3. Configure user macros as described in the *User macro values* section, if necessary.

To update an existing configuration for retrieving secrets from a CyberArk Vault, see *CyberArk configuration*.

Troubleshooting

1. In Zabbix server/proxy log file:

```
237839:20260729:130916.425 Vault token renew started to fail: cannot login into HashiCorp vault with AppRo
```

Probable cause: HashiCorp Vault has been (re)started, but has not been unsealed (required in default configuration, not required if Auto Unseal is configured).

2. In Zabbix server/proxy log file:

```
235545:20260729:125617.944 Vault token renew started to fail: cannot login into HashiCorp vault with AppRo
```

or

```
235545:20260729:125718.537 Vault token renew failed: cannot login into HashiCorp vault with AppRole method
```

Probable cause: VaultAppSecretID has expired or reached its usage limit in HashiCorp vault.

Google Chrome TLS certificate trust

Overview

This page provides Zabbix setup steps and configuration examples for using Google Chrome with Zabbix frontend or zabbix-web-service endpoints secured by a self-signed certificate or a private certificate authority.

These instructions assume that the target web server is already configured for HTTPS. For configuring TLS on the Zabbix frontend, see Secure connection to the frontend.

Configuration

Google Chrome on Linux uses a per-user NSS certificate database for trusted certificates. To make Chrome trust a self-signed certificate or a private certificate authority, add the required certificates to the NSS database of the user account that runs Chrome.

1. Install required packages.

For Debian/Ubuntu:

```
sudo apt install ca-certificates libnss3-tools
```

For RHEL-based systems:

```
sudo dnf install ca-certificates nss-tools
```

Or:

```
sudo yum install ca-certificates nss-tools
```

2. Prepare certificate files.

For each certificate needed: create a PEM-based `.crt` file for the self-signed certificate, or create separate PEM-based `.crt` files for the root certificate authority certificate and, if required, each intermediate certificate.

3. Create the NSS database directory.

Use the home directory of the user that runs Chrome.

For example:

```
sudo mkdir -p /var/lib/zabbix/.pki/nssdb
sudo -u zabbix certutil -N -d sql:/var/lib/zabbix/.pki/nssdb
sudo chown -R zabbix:zabbix /var/lib/zabbix/.pki/nssdb
```

If Chrome on the target system uses a different NSS database location, use that directory instead.

Make sure the certificate files are readable by the user that runs Google Chrome. If necessary, copy them to a location accessible to that user.

4. Import the certificates.

If using a self-signed certificate:

```
sudo -u zabbix certutil -d sql:/var/lib/zabbix/.pki/nssdb \
  -A -t "P,," \
  -n "Zabbix self-signed certificate" \
  -i /path/to/self-signed.crt
```

If using a private certificate authority, import the root certificate authority certificate into the database:

```
sudo -u zabbix certutil -d sql:/var/lib/zabbix/.pki/nssdb \
  -A -t "C,," \
  -n "Zabbix root certificate authority" \
  -i /path/to/root-ca.crt
```

If any intermediate certificates are used, import each one separately:

```
sudo -u zabbix certutil -d sql:/var/lib/zabbix/.pki/nssdb \
  -A -t ",," \
  -n "Zabbix intermediate certificate" \
  -i /path/to/intermediate-ca.crt
```

5. Verify the certificate database.

```
sudo -u zabbix certutil -d sql:/var/lib/zabbix/.pki/nssdb -L
```

6. Test Chrome.

Google Chrome uses the certificate database of the user account that launches it. If the user account does not have a writable home directory, configure a writable home and XDG directories before running the test.

Example (replace `google-chrome` with `chromium` if Chromium is used.):

```
sudo -u zabbix env \
  HOME=/var/lib/zabbix-home \
  XDG_CONFIG_HOME=/var/lib/zabbix-home/.config \
  XDG_DATA_HOME=/var/lib/zabbix-home/.local/share \
  XDG_CACHE_HOME=/var/lib/zabbix-home/.cache \
  google-chrome --headless --disable-gpu \
  --user-data-dir=/var/lib/zabbix-home/chrome-profile \
  --dump-dom https://FQDN_OF_ZABBIX_SERVER
```

If the configuration is correct, the command returns the HTML of the requested page.

If Chrome does not trust the certificate, the command returns a certificate error page containing messages such as `NET::ERR_CERT_AUTHORITY_INVALID`.

Troubleshooting

Error	Possible cause
`NET::ERR_CERT_AUTHORITY_INVALID`	The certificate was imported into the wrong user profile, the root certificate authority certificate is missing, or an intermediate certificate is missing.
`NET::ERR_CERT_COMMON_NAME_INVALID`	The certificate does not match the hostname used in the browser URL.
`NET::ERR_CERT_DATE_INVALID`	The certificate is expired, not yet valid, or the system clock is incorrect.
Chrome still shows a certificate warning after import	Chrome is using a different NSS database for the current user, or the certificate trust flags are incorrect.

Security best practices

Overview This section contains best practices for setting up Zabbix in a secure way.

The practices in this section are not required for the functioning of Zabbix but are recommended for better system security.

For best practices on keeping Zabbix mobile app and Zabbix bridge adapter secure, see Secure mobile access.

UTF-8 encoding UTF-8 is the only encoding supported by Zabbix. It is known to work without any security flaws. Users should be aware that there are known security issues if using some of the other encodings.

Windows installer paths When using Windows installers, it is recommended to use the default paths provided by the installer. Using custom paths without proper permissions could compromise the security of the installation.

Macros in user-defined global scripts To enhance security, it is recommended to use macro functions instead of plain macros in user-defined global scripts, as macros are not automatically escaped.

Zabbix Security Advisories and CVE database See *Zabbix Security Advisories and CVE database*.

HTML email templates When creating or editing message templates used for HTML emails, always wrap every macro with the `htmlencode()` macro function. For example:

```
<b>Problem started</b> at {{EVENT.TIME}.htmlencode()} on {{EVENT.DATE}.htmlencode()}<br><b>Problem name:</
{{EVENT.NAME}.htmlencode()}<br><b>Host:</b> {{HOST.NAME}.htmlencode()}
<br><b>Severity:</b>
{{EVENT.SEVERITY}.htmlencode()}<br><b>Operational
data:</b> {{EVENT.OPDATA}.htmlencode()}<br><b>Original problem ID:</b> {{EVENT.ID}.htmlencode()}<br>{{TRIG
```

Using `htmlencode()` ensures that any HTML characters in macro values are escaped and prevents injection of HTML into notifications (for example, an attacker inserting a malicious/phishing link into a notification).

Note: the default HTML email messages provided by Zabbix already apply `htmlencode()` to macros. This guidance applies when editing existing templates or creating new ones—verify that macros are encoded before using a template to send HTML email.

Web interface

Overview For an easy access to Zabbix from anywhere and from any platform, the web-based interface is provided.

Note:
If using more than one frontend instance make sure that the locales and libraries (LDAP, SAML etc.) are installed and configured identically for all frontends.

Frontend help A help link is provided in Zabbix frontend sections and forms in the upper-right corner, linking directly to the corresponding parts of the documentation.

For form fields with a help icon, hovering over the icon displays the corresponding hint.

Kiosk mode For frontend sections with the *Kiosk mode* button (located in the upper-right corner), click the button to switch the section to kiosk mode. In kiosk mode, only the section content is displayed.

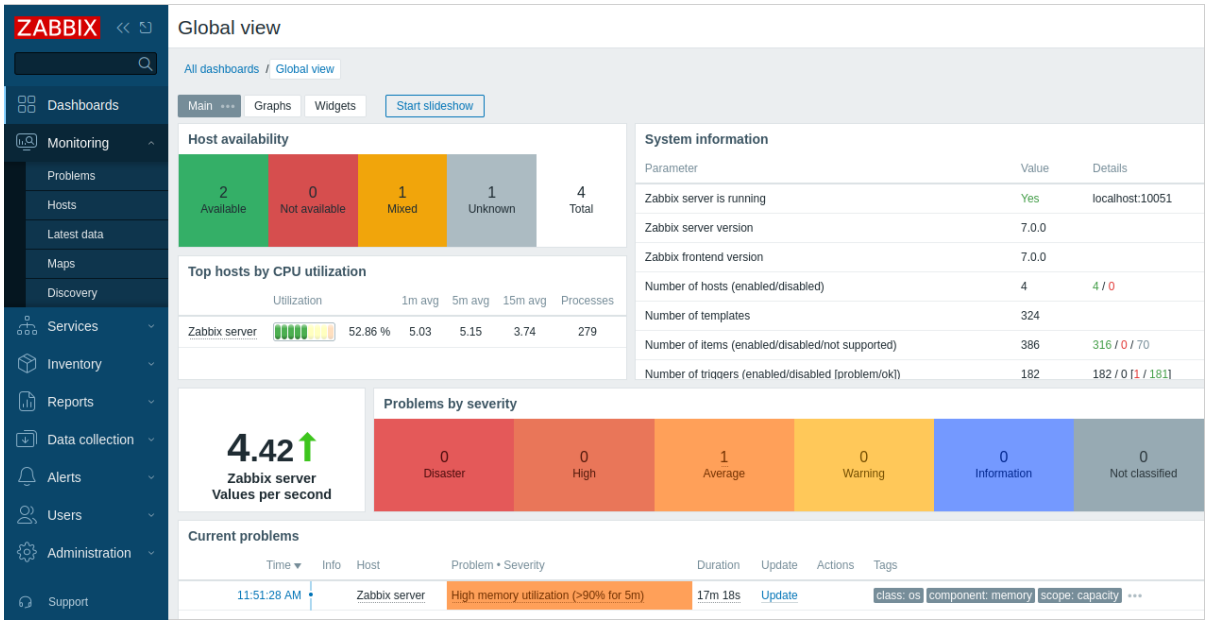
To exit kiosk mode, hover over the content and click the *Normal view* button that appears in the upper-right corner.

Menu

Overview

A vertical menu in a sidebar provides access to various Zabbix frontend sections.

The menu is dark blue in the default theme.

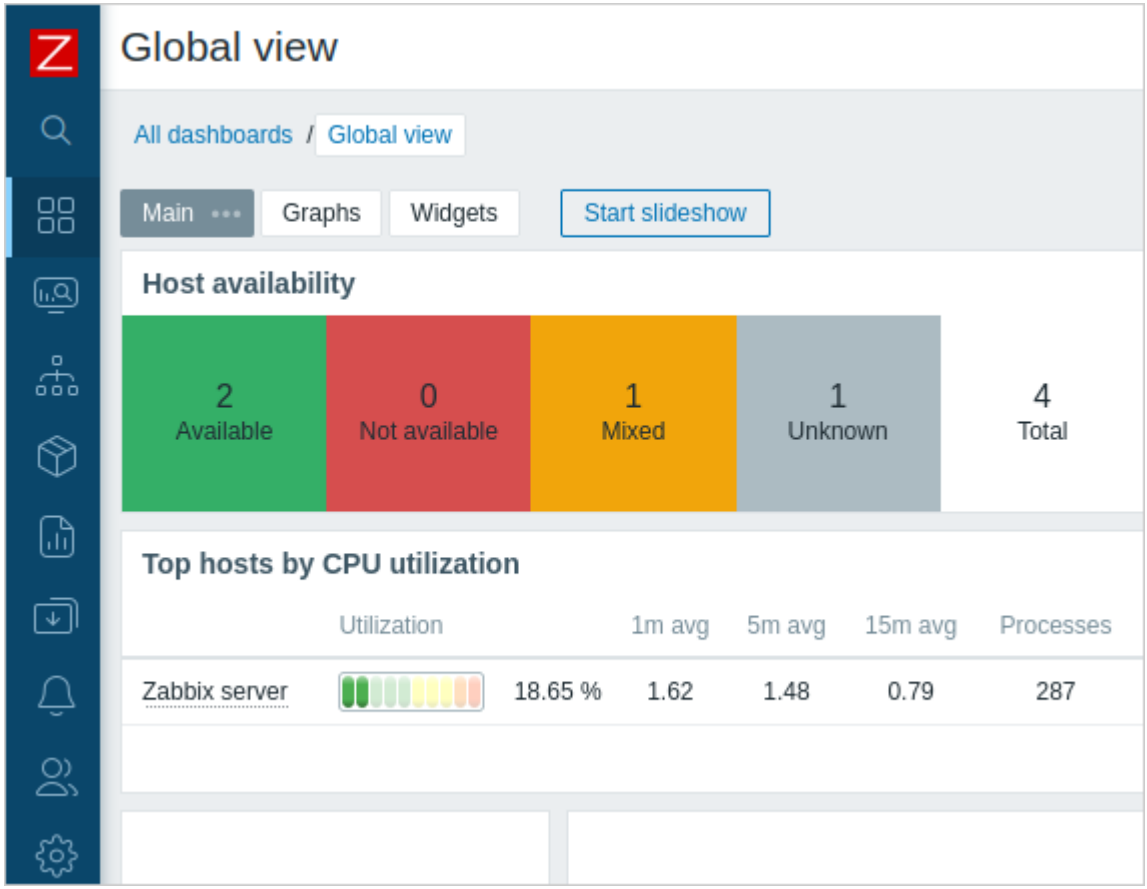


Working with the menu

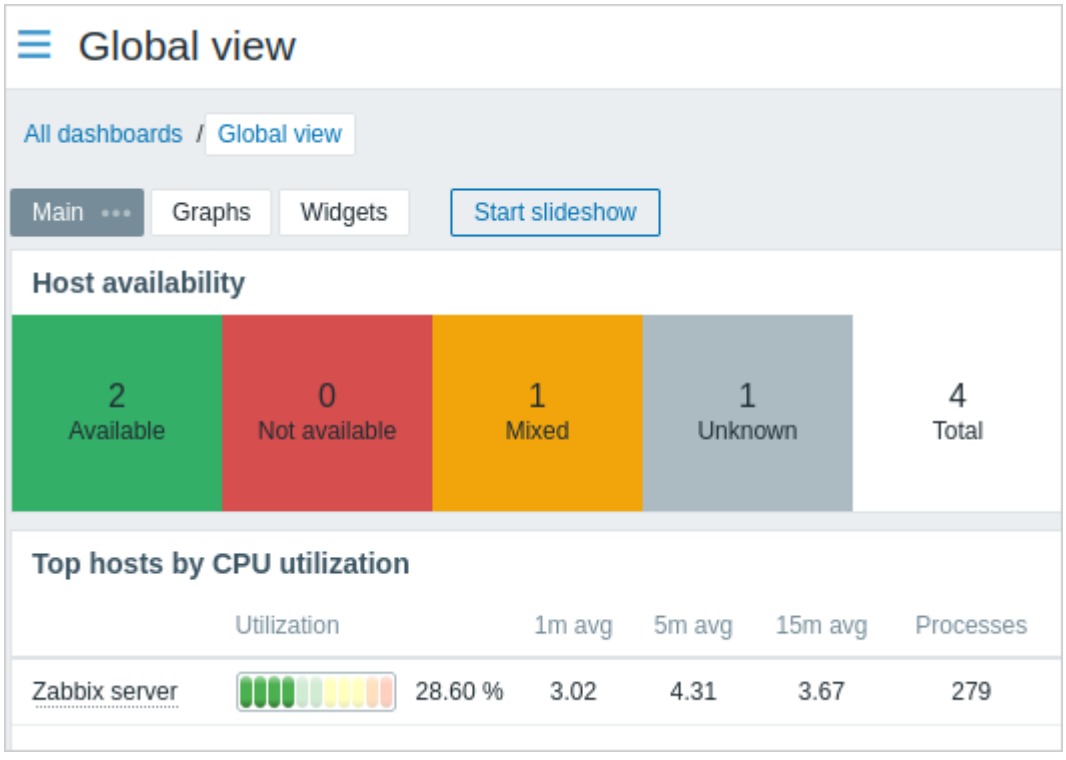
A global search box is located below the Zabbix logo.

The menu can be collapsed or hidden completely:

- To collapse, click next to Zabbix logo. In the collapsed menu only the icons are visible.



- To hide, click  next to Zabbix logo. In the hidden menu everything is hidden.



Collapsed menu

When the menu is collapsed to icons only, a full menu reappears as soon as the mouse cursor is placed upon it. Note that it reappears over page content; to move page content to the right you have to click the expand button. If the mouse cursor again is placed outside the full menu, the menu will collapse again after two seconds.

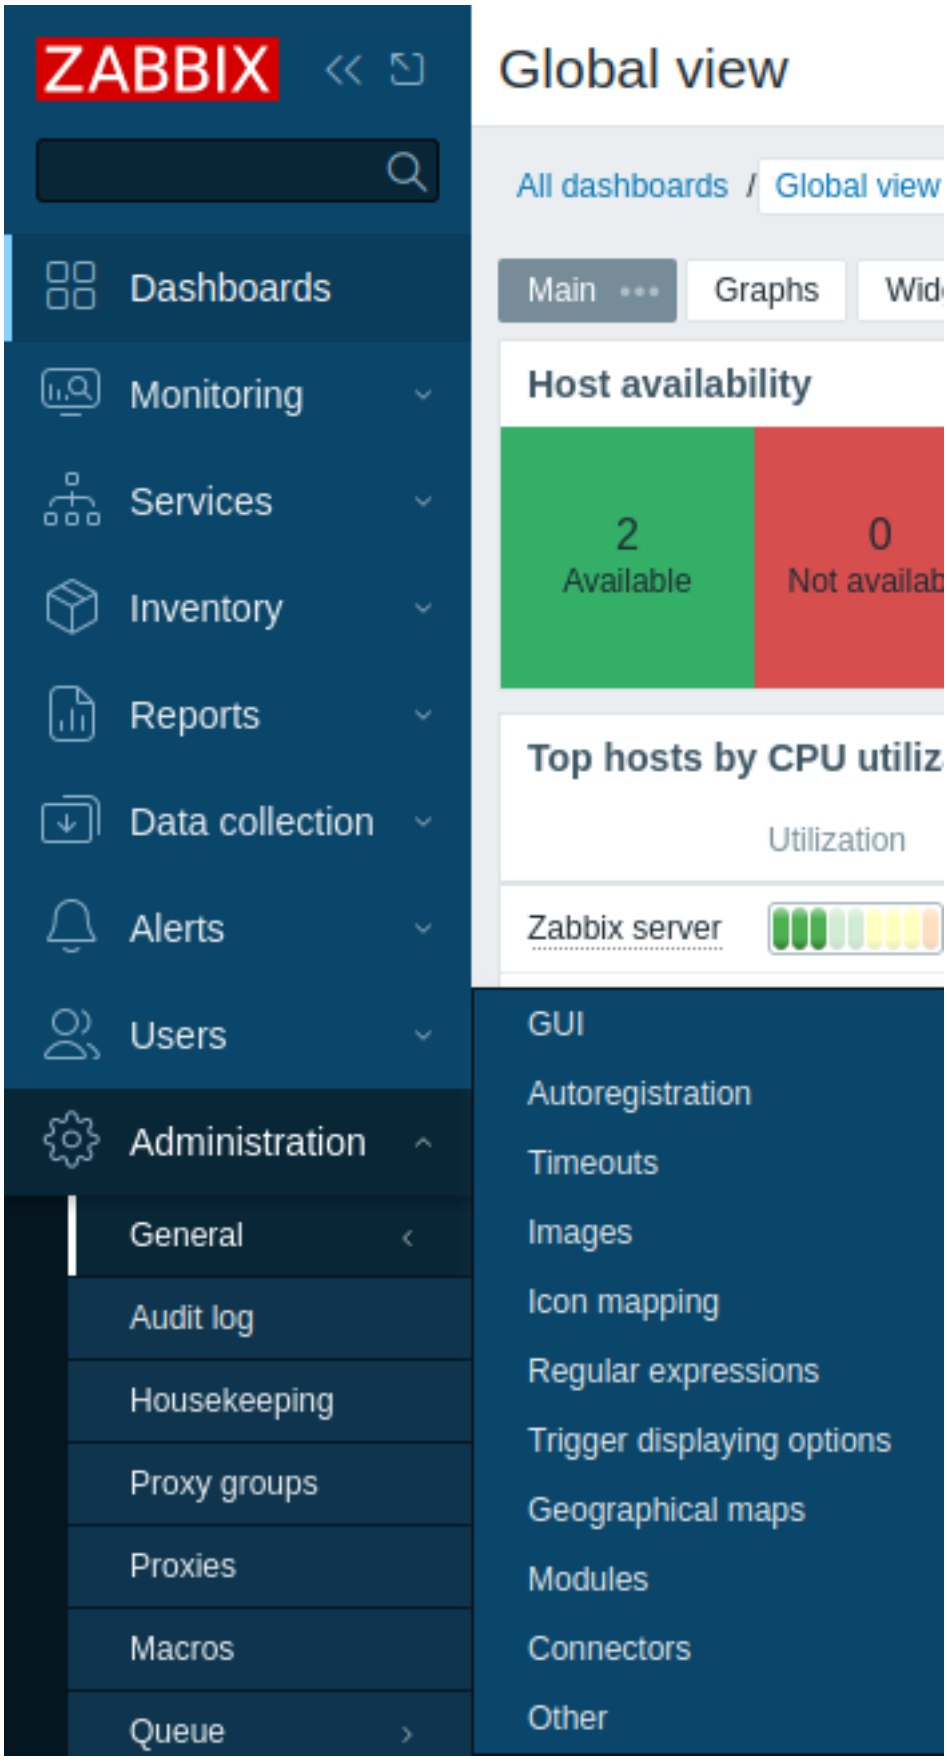
You can also make a collapsed menu reappear fully by hitting the Tab key. Hitting the Tab key repeatedly allows you to focus on the next menu element.

Hidden menu

Even when the menu is hidden completely, a full menu is just one mouse click away, by clicking the burger icon . Note that it reappears over page content; to move page content to the right you have to unhide the menu by clicking the show sidebar button.

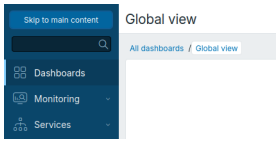
Menu levels

There are up to three levels in the menu.



Accessibility: Skip navigation

The *Skip to main content* link allows keyboard and screen reader users to skip repetitive sidebar navigation and move directly to the main content.



It remains hidden during normal use and appears only when focused, such as when the Tab key is pressed after the page loads or when it is accessed using a screen reader.

Context menus

In addition to the main menu, Zabbix provides context menus for host, item, and event for quick access to frequently used entities, such as the latest values, simple graph, configuration form, related scripts or external links.

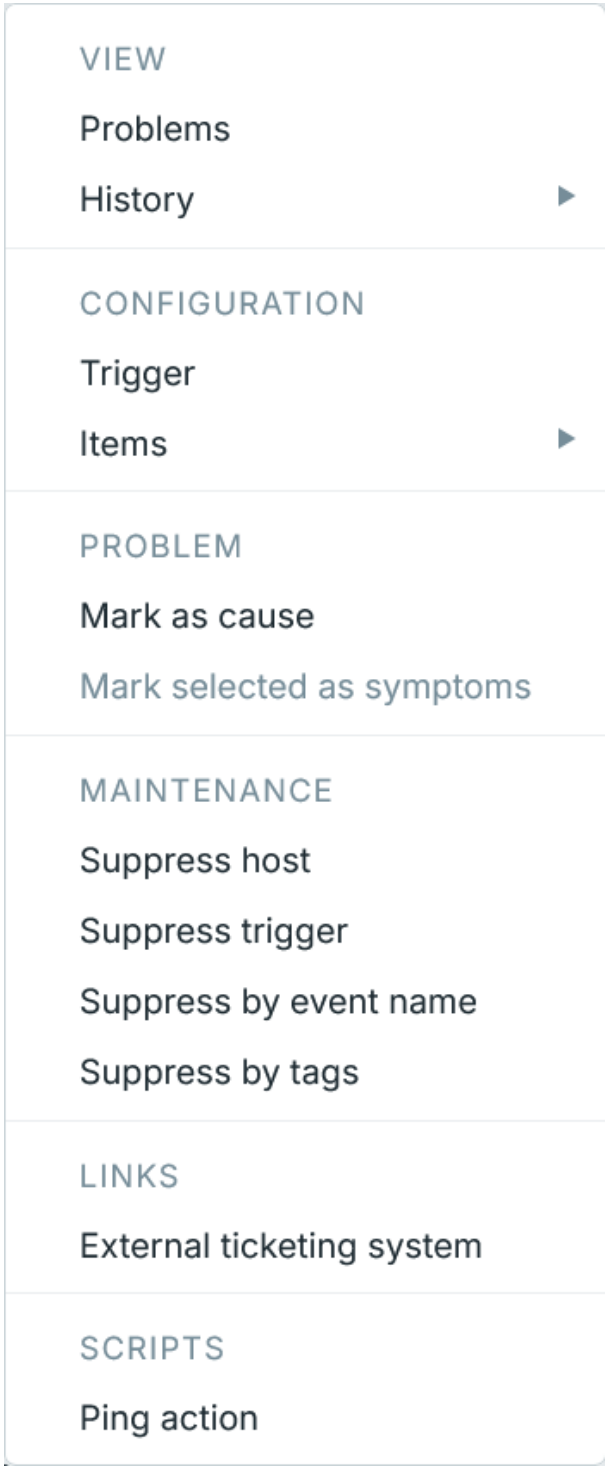
The context menus are accessible by clicking the host, item or problem/trigger name in supported locations.

Event menu

Overview

The event menu contains shortcuts to actions or frontend sections that are frequently required for an event.

The event menu can be opened by clicking the event name.



Content

The event context menu has several sections: *View*, *Actions*, *Configuration*, *Problem*, *Maintenance*, *Links*, and *Scripts*. For the entities that are not configured, links are disabled and displayed in gray. The sections *Scripts* and *Links* are displayed if their entities are configured.

The *View* section contains links to:

- **Problems** - opens the list of unresolved problems of the underlying trigger;
- **History** - leads to the *Latest data* graph/item history for the underlying item(s). If a trigger uses several items, links will be available for each of them.

The *Actions* section is available in *Trigger overview* widgets only. It contains a link to:

- **Update problem** - opens the problem update screen.

The *Configuration* section contains links to the configuration of:

- **Trigger** that fired the problem;
- **Items** used in the trigger expression.

Note:
Note that configuration section is available only for Admin and Super admin users.

The *Problem* section contains the options to:

- **Mark as cause** - mark the problem as cause;
- **Mark selected as symptoms** - mark the selected problems as symptoms of this problem.

The *Maintenance* section contains options to start ad-hoc maintenance:

- **Suppress host** - start a maintenance period for the host;
- **Suppress trigger** - start a maintenance period for the trigger;
- **Suppress by event name** - start a maintenance period for the event name;
- **Suppress by tags** - start a maintenance period for events with the corresponding tags.

When you select the respective option, a maintenance period configuration form will open.

A one-time maintenance period of one hour (configurable in user settings and profile), starting from now, will be pre-selected. The respective host/trigger/event name/tags will be preselected, too. The maintenance name will be pre-filled, in the format `Ad-hoc: <event name>`.

The *Links* section contains links to:

- access a configured trigger URL;
- access custom links configured in Global scripts (with scope 'Manual event action' and type 'URL');
- access a configured external ticket for the problem (see the *Include event menu entry* option when configuring webhooks).

The *Scripts* section contains links to execute a global script (with scope *Manual event action*). This feature may be handy for running scripts used for managing problem tickets in external systems.

Supported locations

The event context menu is accessible by clicking a problem or event name in various frontend sections, for example:

- Dashboards widgets, such as *Problems* widget, *Trigger overview* widget, etc.
- *Monitoring* > *Problems*
- *Monitoring* > *Problems* > Event details
- *Reports* > *Top 100 triggers* (global scripts and access to external ticket are not supported in this location)

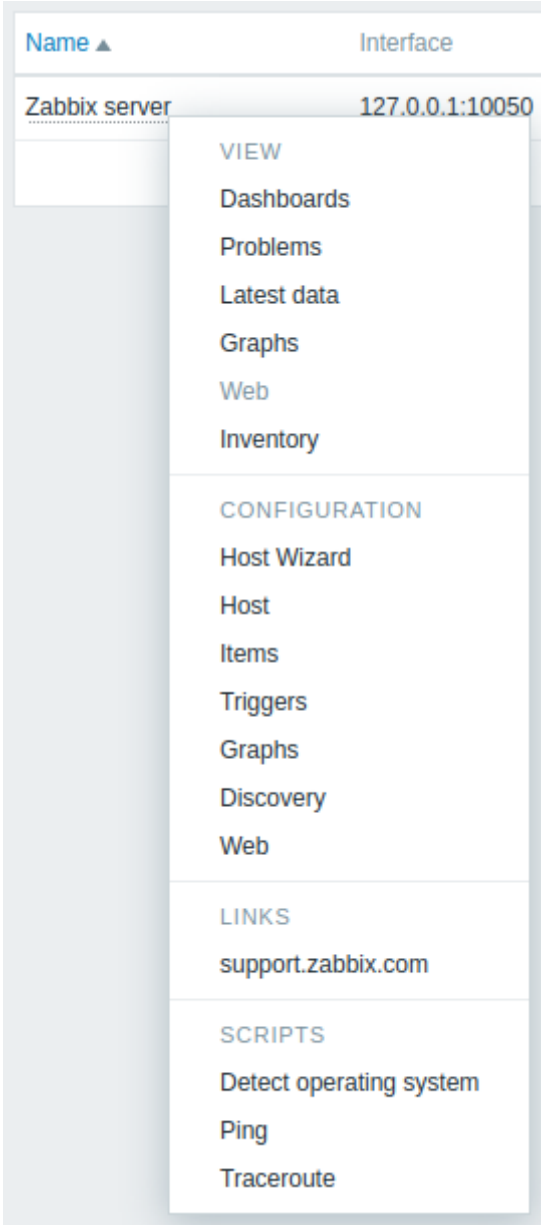
Host menu

Overview

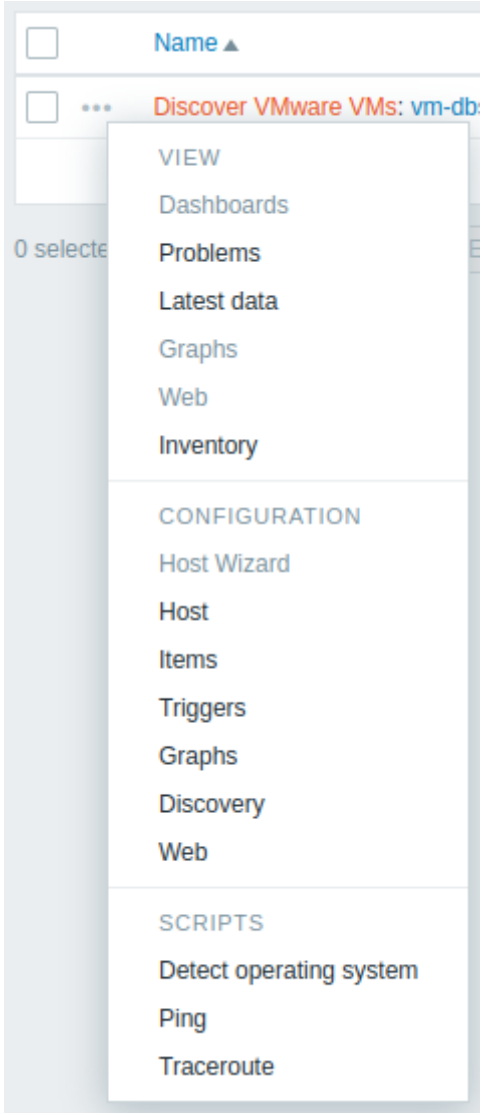
The host menu contains shortcuts to actions or frontend sections that are frequently required for a host.

The host menu can be opened by clicking the host name or the three dots icon, depending on the frontend section, for example:

Monitoring > *Latest data*



Data collection > *Hosts*



Content

The host context menu has four sections: *View*, *Configuration*, *Links*, and *Scripts*. For the entities that are not configured, links are disabled and displayed in gray. The sections *Scripts* and *Links* are displayed if their entities are configured.

View section contains links to:

- **Dashboards** - opens widgets and graphs.
- **Problems** - opens the *Problems* section with the list of unresolved problems of the underlying trigger.
- **Latest data** - opens the *Latest data* section with the list of all the latest data of the current host.
- **Graphs** - opens simple graphs of the current host.
- **Web**- opens the link to the configured web scenarios.
- **Inventory** - opens the link to the inventory of the current host.

Configuration section contains links to:

- **Host Wizard** - opens Host Wizard for the current host (disabled for discovered hosts).
- **Host** - configuration form of the current host.
- **Items** - the list of the current host items.
- **Triggers** - the list of the current host triggers.
- **Graphs** - simple graphs of the current host.
- **Discovery** - the list of low-level discovery rules of the current host.
- **Web** - the list of web scenarios of the current host.

Note:
Note that configuration section is available only for Admin and Super admin users.

Links section contains links to:

- access a configured trigger URL.
- access custom links configured in Global scripts (with scope *Manual host action* and type 'URL').

Scripts section allows execution of global scripts configured for the current host. These scripts need to have their scope defined as *Manual host action* to be available in the host menu.

Supported locations

The host menu is accessible by clicking a host name in various frontend sections, for example:

- Dashboards widgets, such as *Problems*, *Top items*, *Trigger overview*, etc.
- *Monitoring* > *Problems*
- *Monitoring* > *Problems* > Event details
- *Monitoring* > *Hosts*
- *Monitoring* > *Hosts* > *Web Monitoring*
- *Monitoring* > *Latest data*
- *Monitoring* > *Maps*
- *Inventory* > *Hosts*
- *Reports* > *Top 100 triggers*

The host menu is accessible by clicking the three dots icon in Data collection > Hosts.

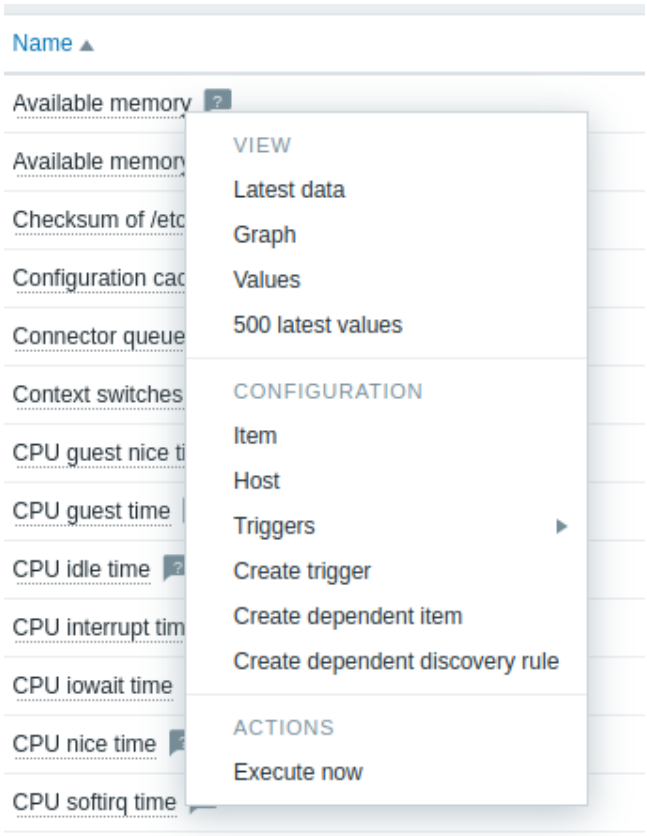
Item menu

Overview

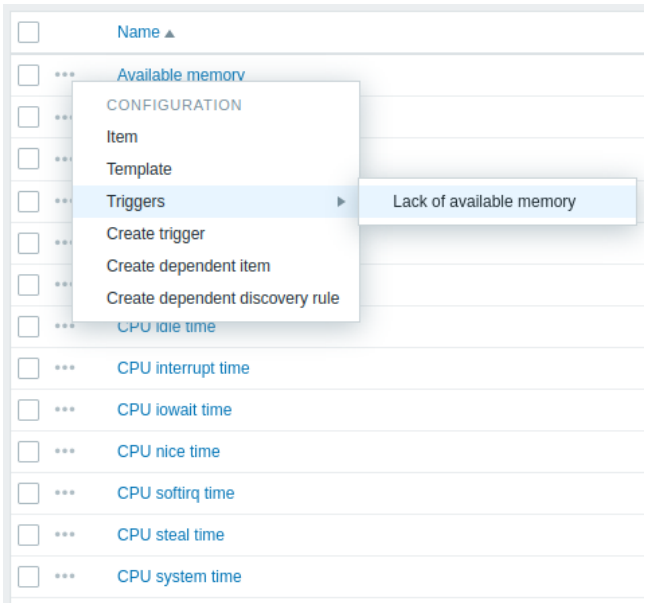
The item menu contains shortcuts to actions or frontend sections that are frequently required for an item.

The item menu can be opened by clicking the item name or the three-dot icon, depending on the frontend section, for example:

Monitoring > *Latest data*



Data Collection > *Templates* > *Items*



Content

The item menu has three sections: *View*, *Configuration*, and *Actions*.

The *View* section contains the following options:

- **Latest data** - opens the *Latest data* section filtered by the current host and item;
- **Graph** - opens a simple graph of the current item;
- **Values** - opens the list of all values received for the current item within the past 60 minutes;
- **500 latest values** - opens the list of 500 latest values for the current item.

The *Configuration* section (available only for *Admin* and *Super admin* type users) contains the following options:

- **Item** - opens the item configuration form of the current item;
- **Template** - opens the template configuration form of the template that the item belongs to (available only when the item menu is accessed from *Data collection* > *Templates* > *Items*);
- **Host** - opens the host configuration form of the host that the item belongs to;
- **Triggers** - on mouseover, opens a list of item's triggers, if any; clicking a trigger opens its trigger configuration form;
- **Create trigger** - opens the trigger configuration form to create a trigger for this item;
- **Create dependent item** - opens the item configuration form to create a dependent item, with the current item as the master item;
- **Create dependent discovery rule** - opens the discovery rule configuration form to create a dependent discovery rule, with the current item as the master item.

The *Actions* section contains the following option:

- **Execute now** - executes a check for a new item value immediately.

Supported locations

The item menu is accessible by clicking an item name in various frontend sections, for example:

- *Monitoring* > *Latest data*
- *Data collection* > *Hosts* > *Items*
- *Data collection* > *Hosts* > *Discovery rules* > *Item prototypes*

The item menu is accessible by clicking an item value in the *Top items* dashboard widget.

Frontend sections

Menu structure The Zabbix frontend menu has the following structure:

- Dashboards
- Monitoring
  - Problems
  - Hosts
  - Latest data
  - Maps
  - Discovery
- APM
  - Traces
  - Metrics
  - Logs
- Services
  - Services
  - SLA
  - SLA report
- Inventory
  - Overview
  - Hosts
- Reports
  - System information
  - Scheduled reports
  - Availability report
  - Top 100 triggers
  - Audit log
  - Action log
  - Notifications

- Data collection
  - Template groups
  - Host groups
  - Templates
  - Hosts
  - Maintenance
  - Event processing
  - Discovery
- Alerts
  - Actions
    - Trigger actions
    - Service actions
    - Discovery actions
    - Autoregistration actions
    - Internal actions
  - Media types
  - Scripts
- Users
  - User groups
  - User roles
  - Users
  - API tokens
  - Authentication
  - Devices
- Administration
  - General
    - GUI
    - Autoregistration
    - Timeouts
    - Images
    - Icon mapping
    - Regular expressions
    - Trigger displaying options
    - Geographical maps
    - Modules
    - Connectors
    - Other
  - Audit log
  - Housekeeping
  - Proxy groups
  - Proxies
  - Macros
  - Queue
    - Queue overview
    - Queue overview by proxy
    - Queue details

Dashboards

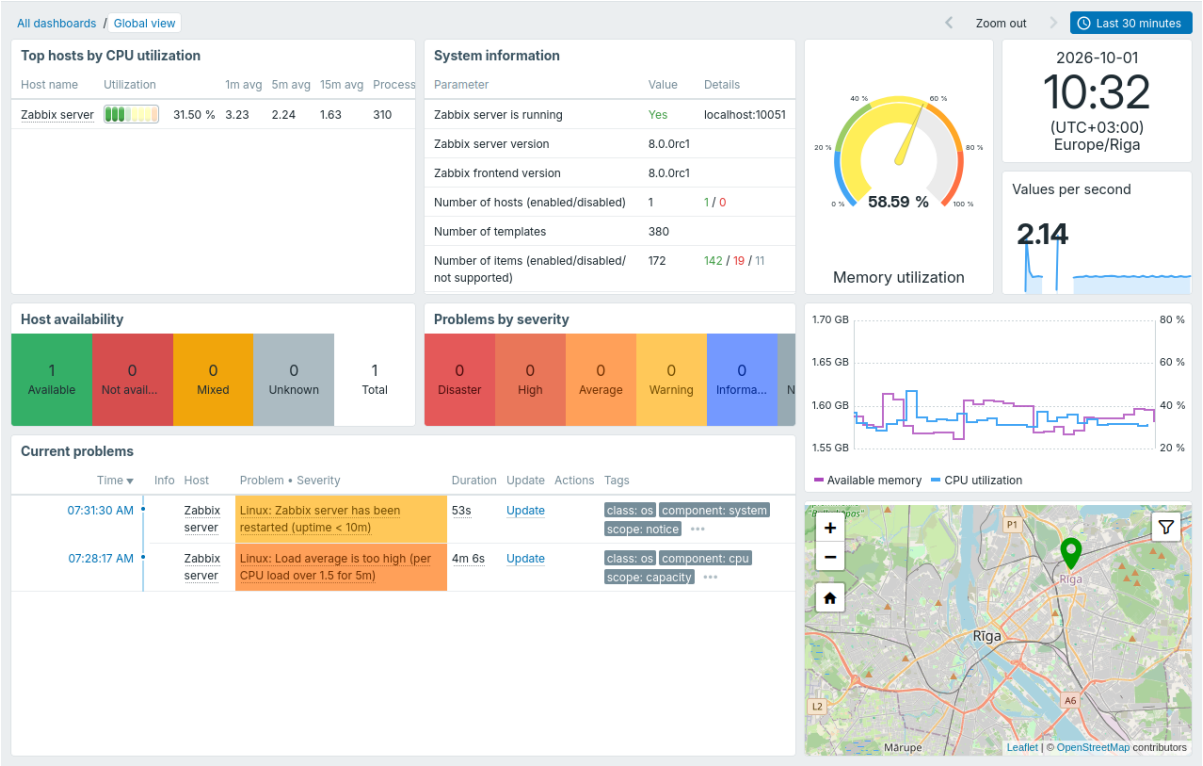
Dashboards give you an overview of your monitored environment. You can build and customize them to meet your needs.

The building blocks are widgets: a graph, a map, a list of recent problems, and so on, each designed to display a particular type of information.

Widgets can be organized into multiple dashboard pages. Such pages can be viewed as a slideshow. You can also navigate to any page by clicking its tab.

You can create several dashboards, though only one is shown at a time.

Zabbix comes with some default dashboards. One of them is the *Global view* dashboard, which shows the health and status of your monitored hosts, including the Zabbix server:

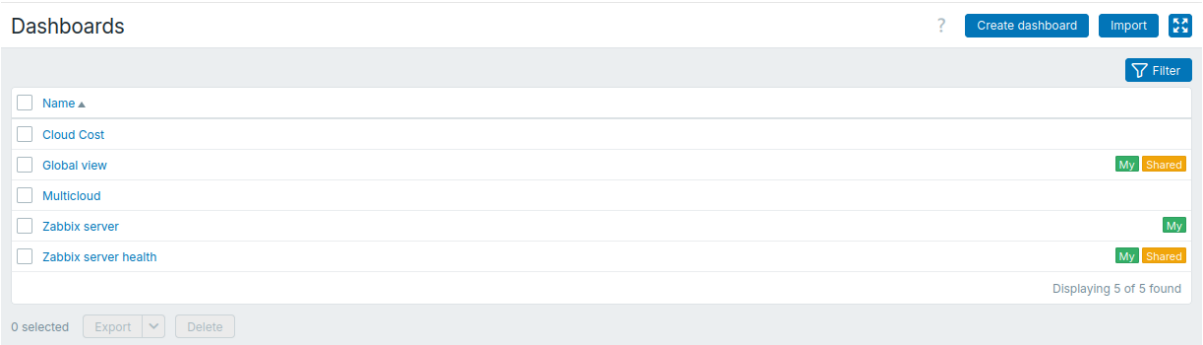


You can find more dashboards in the Dashboards repository. These dashboards are designed to work with existing Zabbix templates and can be downloaded and imported into your Zabbix environment. A detailed description of each dashboard, including prerequisites, is available in the dashboard's README file.

Dashboard list

To open the list of dashboards, either:

- Click *Dashboards* in the menu.
- Click *All dashboards* while viewing a dashboard to open the dashboard list.



Access to the *Dashboards* menu depends on the user role and the *Dashboard* permission under *Access to UI elements*.

From here, you can: - Create dashboards. - Delete dashboards. You can delete a dashboard by selecting its checkbox and clicking the *Delete* button. - Import dashboards. - Export dashboards.

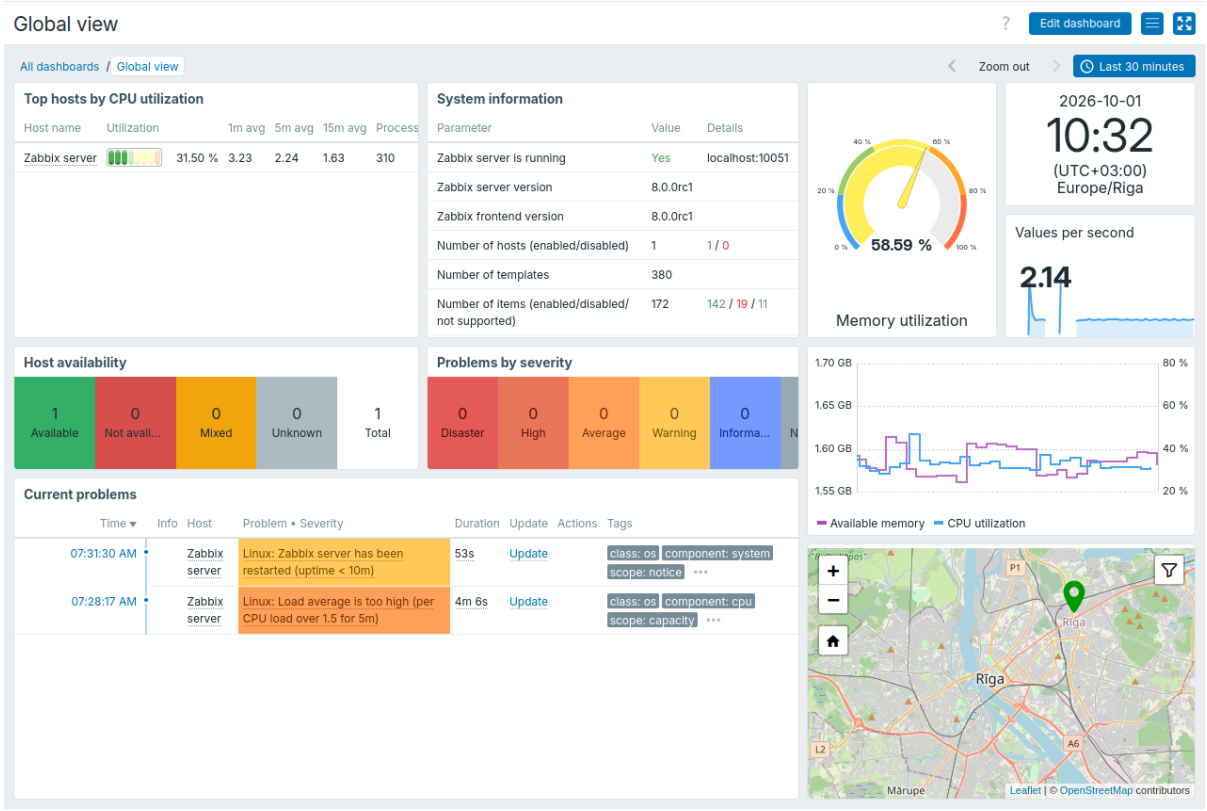
Dashboards have tags indicating their ownership and sharing settings:

- *My* - the dashboard is owned by you and visible only to you.
- *My + Shared* - the dashboard is owned by you but is public or shared with specific users or user groups.
- Dashboards without tags are owned by other users but are public or shared with you.

You can filter the list by dashboard name and owner.

View dashboard

Click a dashboard's name in the dashboard list to view the dashboard:



The Zabbix user interface remembers the last opened dashboard for the next session.

Note:
The data you can see on a dashboard depends on your access to hosts.

For widgets that show graphs, the time period selector above the widgets sets the time range. It shows the current time range as a label. Click the label to open or close the selector.

From here, you can manage dashboards as follows:

- Edit the dashboard.
- Perform the following actions from the dashboard *Actions* menu :
  - *Sharing* - edit the sharing settings for the dashboard.
  - *Create new* - create a new dashboard.
  - *Clone* - create a new dashboard by cloning the current one. You can edit the dashboard properties before saving.
  - *Delete* - delete the dashboard.
  - *Export* - export the dashboard.
  - *Create new report* - create a scheduled report. This option is disabled if you do not have permission to manage scheduled reports.
  - *View related reports* - view a list of existing reports for the current dashboard. This option is disabled if there are no related reports or you do not have permission to manage scheduled reports.
- Switch to kiosk mode.
- Start or stop the dashboard slideshow.

In dashboard view mode, the following options are available for widgets:

- Edit the widget.
- Perform the following actions from the widget *Actions* menu ...:
  - *Copy* - copy the widget.
  - *Download image* - download the widget as a PNG image. This option is available for widgets that support this feature (see individual widget pages).
  - *Refresh interval* - select the frequency of refreshing the widget contents.

For dashboards with multiple pages, the following actions are available in the page *Actions* menu in dashboard view mode:

- *Copy* - copy the page.
- *Properties* - customize the page properties, including the name and the slideshow interval.

Slideshow

If the *Start slideshow* checkbox is selected in the dashboard properties, the slideshow starts automatically when the dashboard has more than one page. Otherwise, you can start the slideshow by clicking the *Start slideshow* button.



The pages change according to the intervals specified in the dashboard and individual page properties. If the *Slideshow interval* property is defined for an individual page, the interval specified for that page is applied. Otherwise, the dashboard's *Default slideshow interval* is applied.

To stop the slideshow, click *Stop slideshow*.

In kiosk mode, use the following buttons:

-  to start the slideshow;
-  to stop the slideshow;
-  to navigate to the previous page;
-  to navigate to the next page.

Share dashboard

Dashboards can be public or private.

Public dashboards are visible to all users. Private dashboards are visible only to their owners and to the users or user groups they are shared with.

To change the sharing settings of a dashboard, click *Sharing* in the dashboard *Actions* menu while viewing the dashboard:

Parameter	Description
Type	Select the dashboard type: **Private** - the dashboard is visible only to selected user groups and users. **Public** - the dashboard is visible to all users.
List of user group shares	Select the user groups that can access the dashboard. You can grant read-only or read-write access.
List of user shares	Select the users that can access the dashboard. You can grant read-only or read-write access.

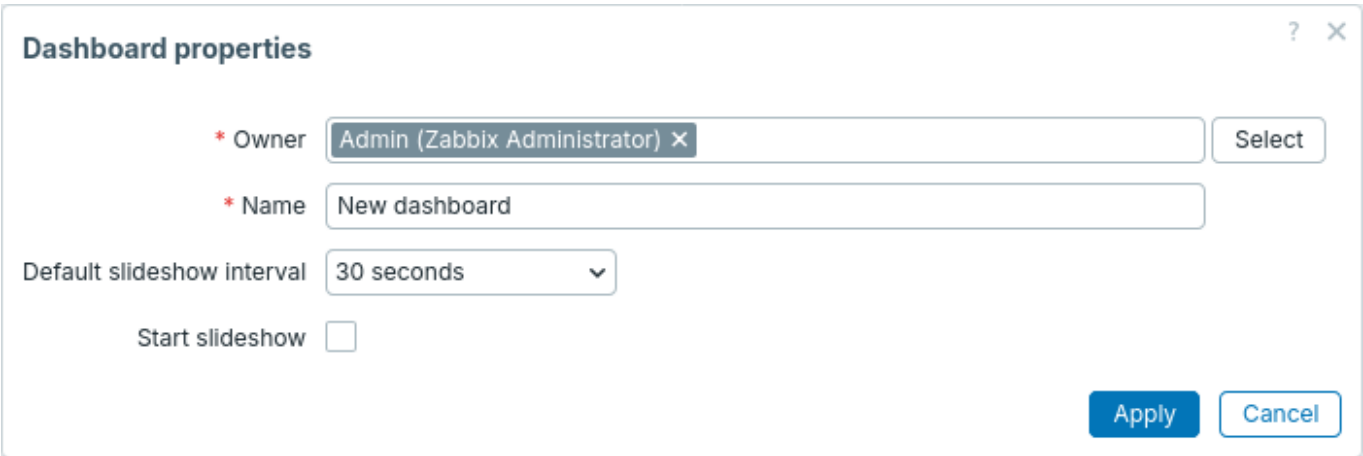
See Permissions to dashboards to find out how sharing options limit the actions available to users.

Create dashboard

You can create a dashboard in two ways:

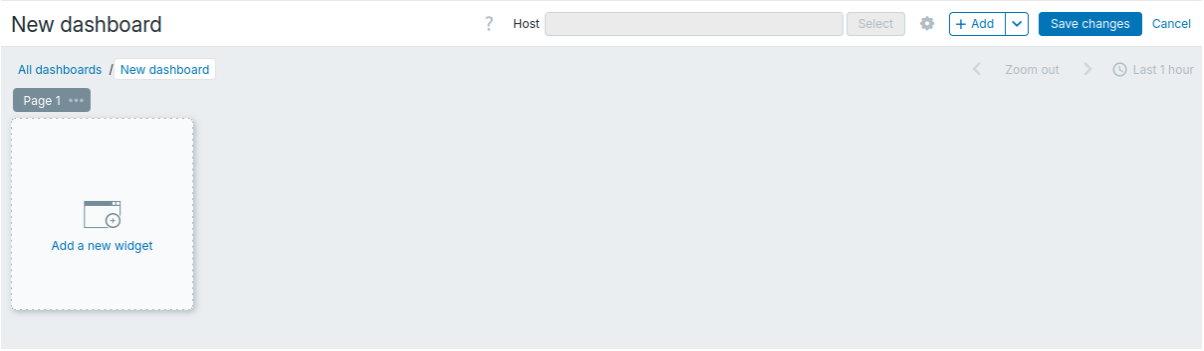
- In the dashboard list, click the *Create dashboard* button.
- In the dashboard view mode, click *Create new* in the dashboard *Actions* menu.

Then, specify the dashboard properties:



Parameter	Description
Owner	Select the user that will own the dashboard.
Name	Enter the dashboard name.
Default slideshow interval	Select how long each dashboard page is displayed before the slideshow moves to the next page.
Start slideshow	Mark this checkbox to start the slideshow automatically once the dashboard has more than one page.

Then, click *Apply*. The new dashboard opens in edit mode:



From here, you can add dashboard widgets and pages.

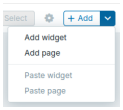
Note that the minimum width of a dashboard is 1200 pixels; in narrower windows, a horizontal scrollbar appears. A wider dashboard always fills the browser window; its widgets cannot extend beyond the window edges.

Click the *Save changes* button to save the dashboard, or *Cancel* to return to the previous screen.

Add widgets

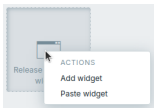
In dashboard edit mode, you can add widgets can be added in two ways:

- Click the *Add* button. Alternatively, click the drop-down arrow on the *Add* button and select *Add widget*.



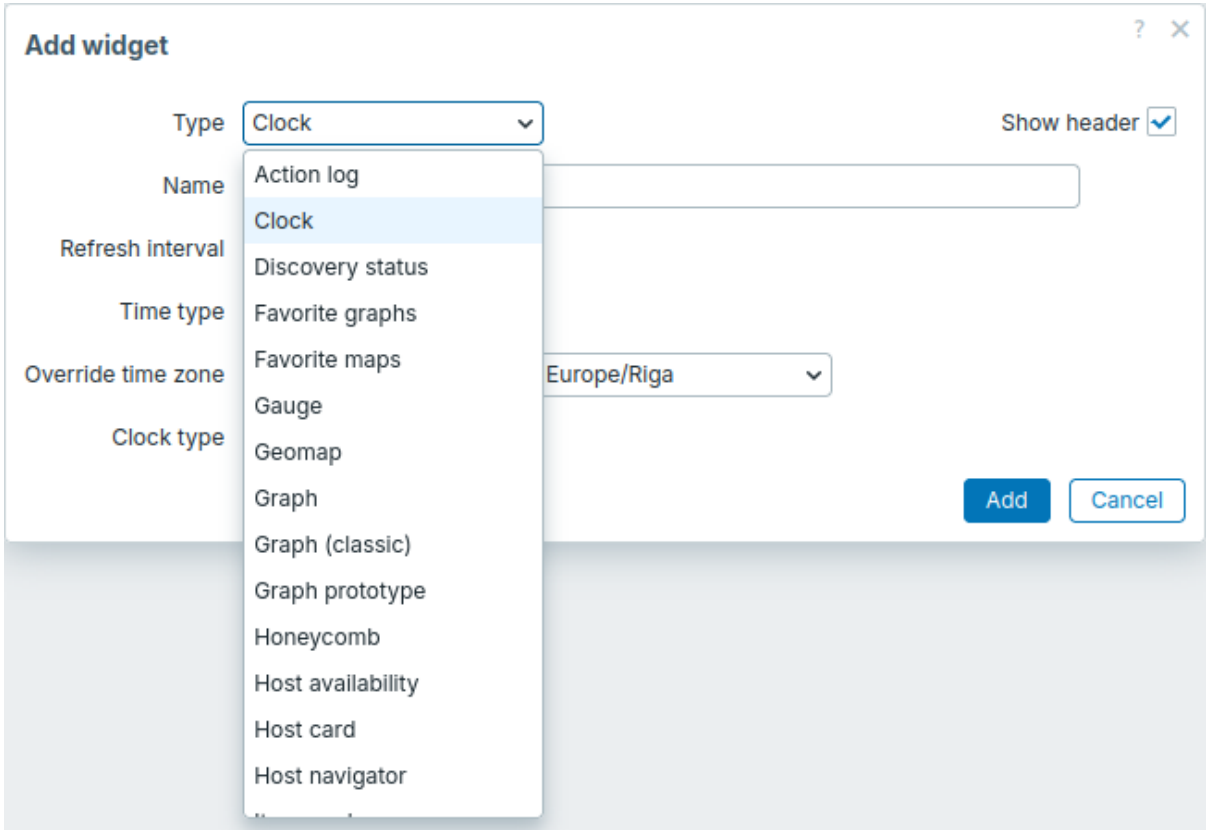
  The widget placeholder is added at its default size to the first available empty space on the dashboard, with the widget configuration form open.

- Move the mouse pointer over an empty area of the dashboard. Then click the widget placeholder to open the widget configuration form. Alternatively, click and drag the placeholder to resize it, then release to open the widget configuration form; if a widget is copied to the clipboard, you will first be prompted to choose *Add widget* or *Paste widget*.



Then, follow these steps:

1. In the widget configuration form, select the widget type.
2. Configure widget parameters.
3. Click *Add* to add the widget to the dashboard.



The widget is added at its default size or resized to fit the available space. You can resize and reposition widgets as described in the *Edit widgets* section.

You can also copy and paste widgets to create new widgets with the same configuration, either within the same dashboard or between dashboards opened in different browser tabs.

Copy/paste widgets

To copy and paste a widget:

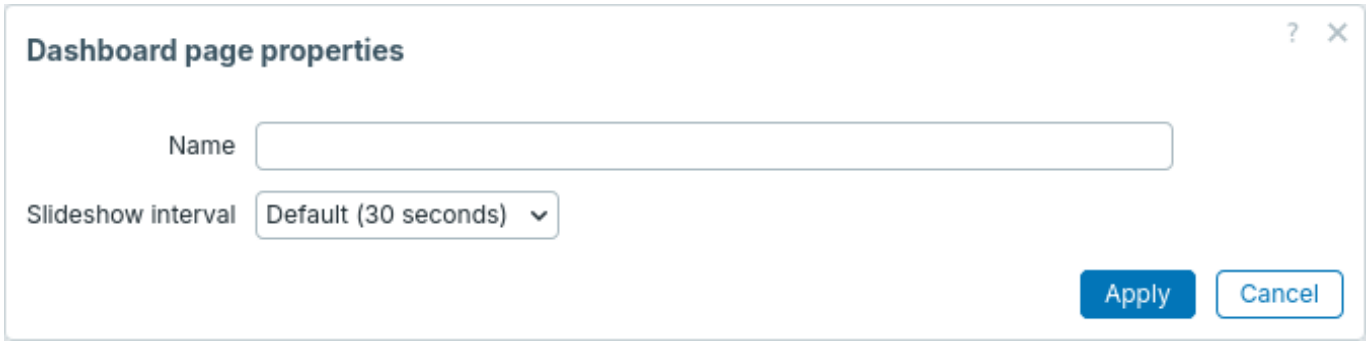
1. In the widget *Actions* menu ⋮, select *Copy*.
2. In dashboard edit mode, click the drop-down arrow on the *Add* button and select *Paste widget*. Alternatively, click an empty area in the dashboard and select *Paste widget*.

You can also paste a copied widget over an existing one using the *Paste* option in its *Actions* menu.

Add pages

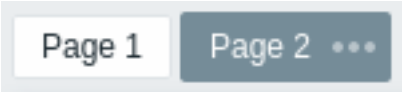
To add a new page to a dashboard:

1. In dashboard edit mode, click the drop-down arrow on the *Add* button and select *Add page*.
2. Fill in the page properties:
   - *Name* - enter the name of the page. If you leave the name empty, the page is added with a *Page N* name, where 'N' is the incremental page number.
   - *Slideshow interval* - select how long the page is displayed before the slideshow moves to the next page.



3. Click *Apply*.

A new tab (e.g., *Page 2*) is added for the new page.



You can also copy and paste pages to create new pages with the same properties, either within the same dashboard or between dashboards opened in different tabs.

Copy/paste pages

To copy and paste a page:

1. In the page *Actions* menu , select *Copy*.
2. In dashboard edit mode, click the drop-down arrow on the *Add* button and select *Paste page*.

Edit dashboard

To edit a dashboard, click the *Edit dashboard* button in the upper-right corner. Alternatively, click the *Edit* button for a widget. The dashboard opens in edit mode.

From here, you can manage dashboards as follows:

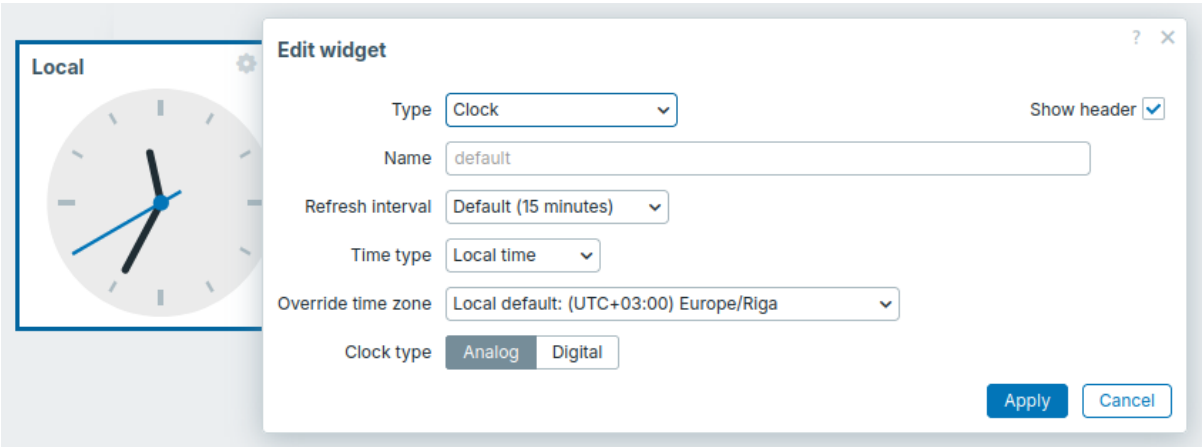
- Edit the dashboard properties by clicking the gear button .
- Edit dashboard widgets.
- Access the following dashboard actions by clicking the drop-down arrow on the *Add* button:
  - *Add widget* - add a widget to the dashboard.
  - *Add page* - add a page to the dashboard.
  - *Paste widget* - paste a copied widget. This option is grayed out if no widget has been copied. Only one widget can be copied at a time.
  - *Paste page* - paste a copied page. This option is grayed out if no page has been copied. Only one page can be copied at a time.

After editing the dashboard, click the *Save changes* button to save the dashboard, or *Cancel* to return to the previous screen.

Edit widgets

A widget can be edited in dashboard edit mode by clicking its *Edit* button . Alternatively, when viewing a dashboard, you can hover over a widget and click the *Edit* button.

While a widget is being edited, it is highlighted, and changes are previewed in real time.



You can reposition the widget and its configuration form by dragging their title bars and resize the widget by dragging its resize handles.



The following actions are available from the widget *Actions* menu ... when editing a widget:

- *Copy* - copy the widget.
- *Paste* - paste a copied widget over this widget. This option is grayed out if no widget has been copied.
- *Delete* - delete the widget.

If mandatory widget parameters are not configured, the widget enters the *Widget is not fully configured* state.

After editing, click *Apply* in the widget configuration form, then *Save changes* in the upper-right corner of the dashboard to apply your changes.

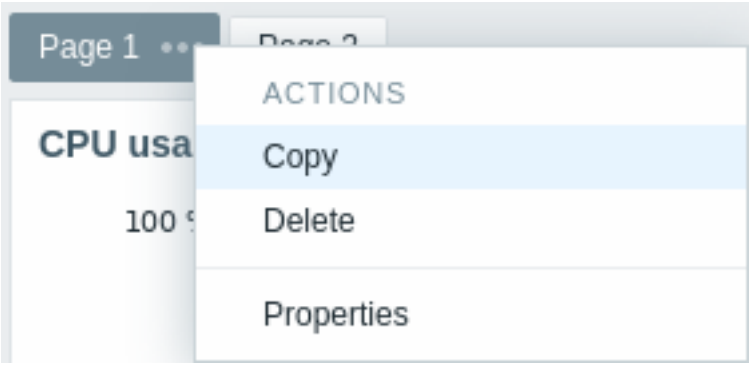
Edit pages

Pages can be edited in dashboard edit mode.

The pages can be reordered by dragging and dropping the page tabs. Reordering does not rename pages, except for pages with automatically assigned names (*Page N*), which are renumbered.

The following actions are available from the page *Actions* menu when editing a page:

- *Copy* - copy the page.
- *Delete* - delete the page.
- *Properties* - customize the page properties.



Permissions to dashboards

Permissions to dashboards for regular users and users of the *Admin* type are limited as described below, provided that access to dashboards and the *Create and edit dashboards* option are enabled for their assigned user role:

- They can see and clone a public dashboard even when no users or user groups are added to the dashboard's sharing preferences.
- They can see and clone a private dashboard if they have at least *Read* rights to it, as specified in the sharing preferences.
- They can edit and delete a dashboard only if they have *Read-write* rights to it, as specified in the sharing preferences.
- They cannot change the dashboard owner.

Dashboard widgets

Overview

Zabbix provides a variety of dashboard widgets that visualize and summarize monitoring data, giving you a clear overview of system performance and availability.

Below is an overview of all widgets, grouped by their primary purpose.

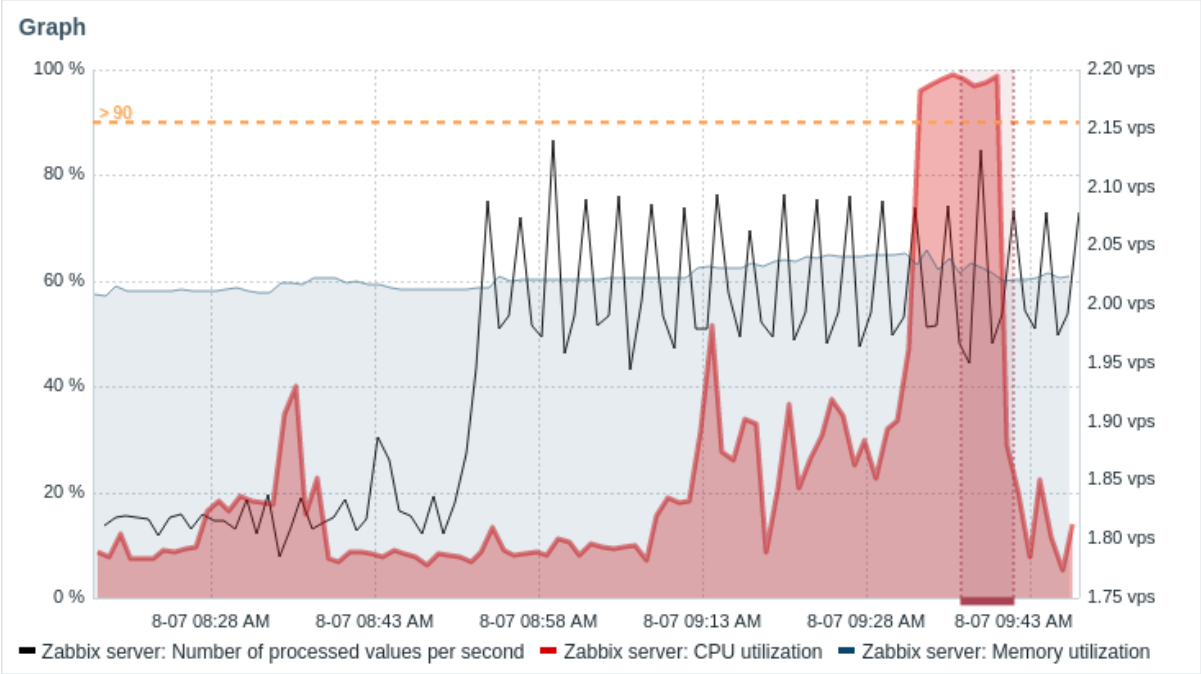
For widget configuration details, see Dashboard widget parameters.

Data visualization

These widgets present data in various visual formats, helping you monitor data and problems in different ways.

Graph

The Graph widget displays numeric item data as a vector-based graph. It can help you track metrics, spot issues, and compare values over time and across hosts.

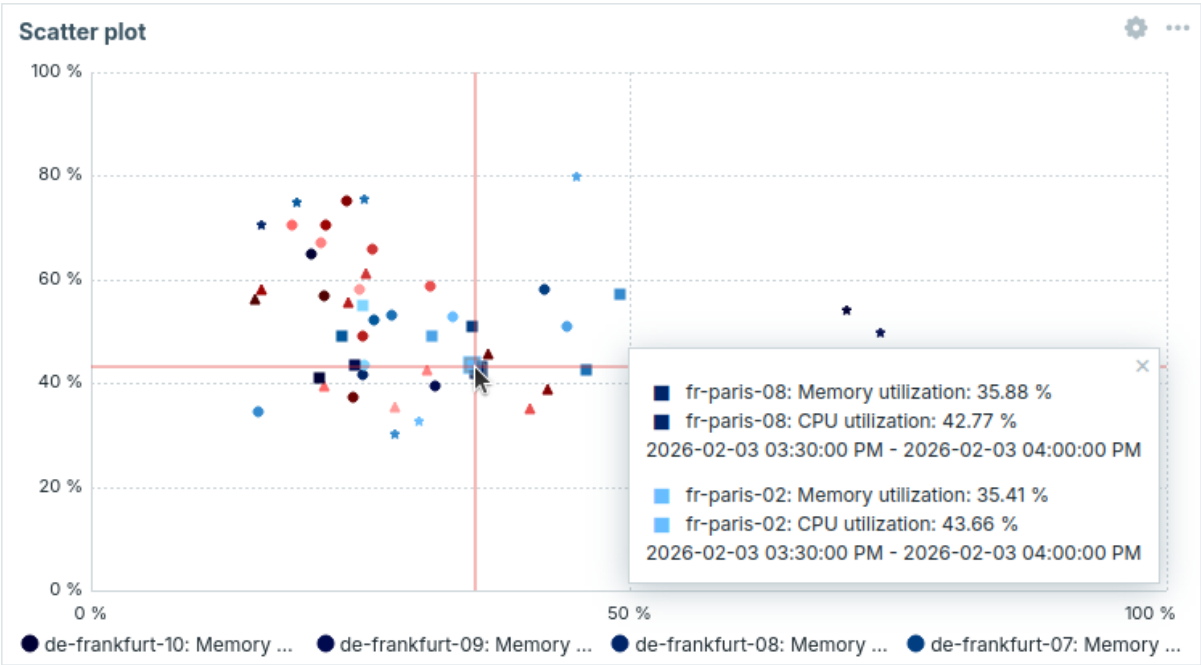


See also:

- Graph (classic) - the legacy Graph widget that displays numeric item data as an image-based graph.
- Graph prototype - displays a grid of automatically created custom graphs based on low-level discovery rule graph prototypes or item prototypes.
- Favorite graphs - displays an alphabetical list of shortcuts to graphs marked as favorites by the current user.

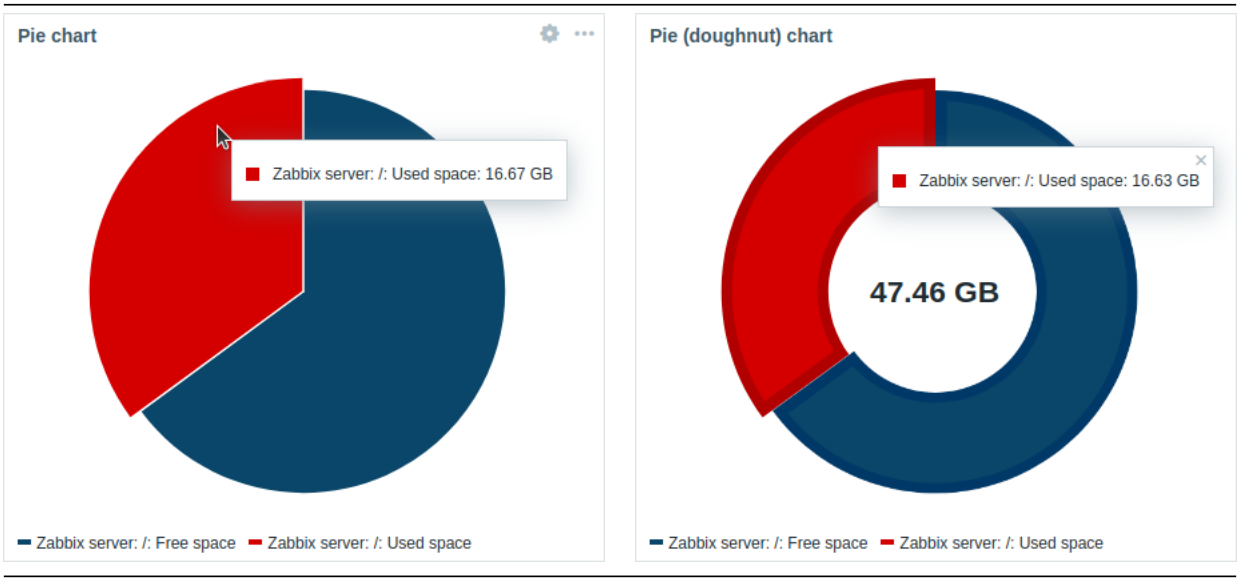
Scatter plot

The Scatter plot widget displays the relationship between two metrics by plotting individual data points along an X and Y axis. This helps reveal patterns, clusters, correlations, and outliers in the data set.



Pie chart

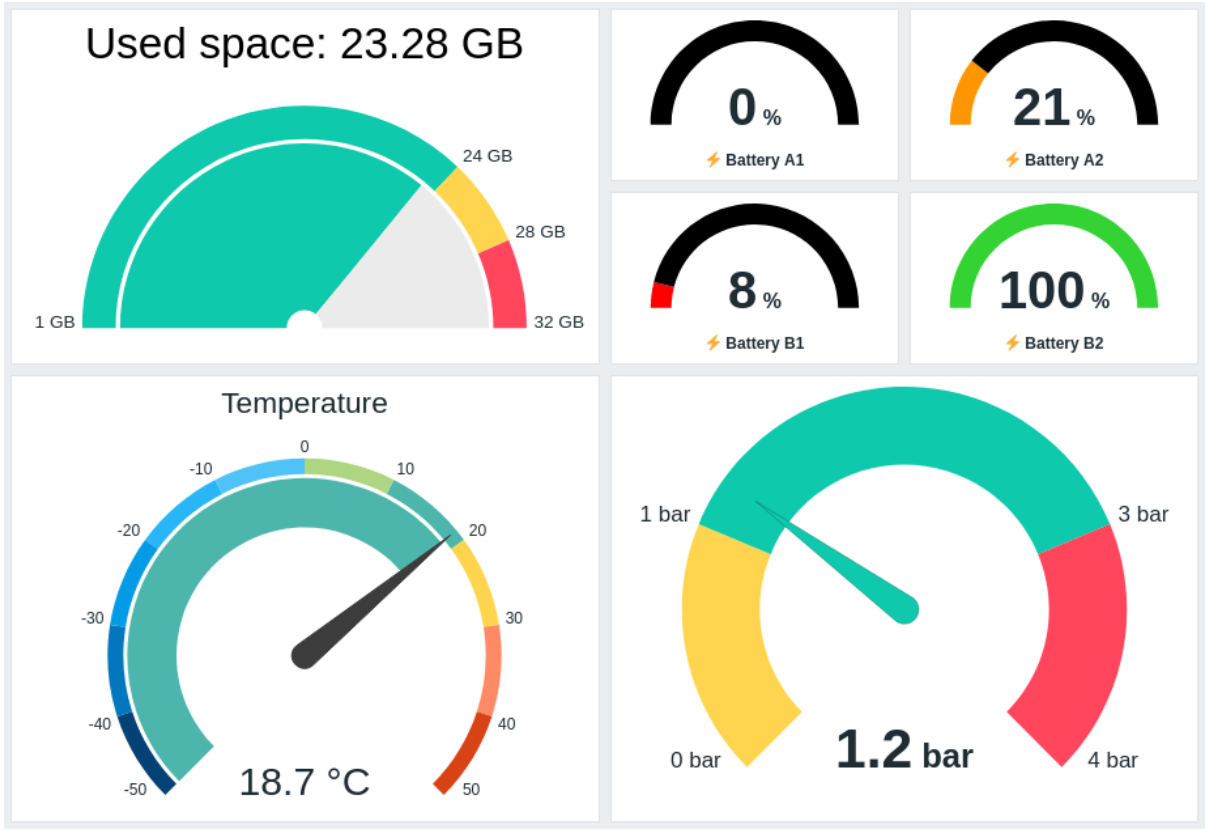
The Pie chart widget displays numeric item data as a vector-based pie or doughnut chart. This widget can help you visualize how items or hosts contribute to the overall dataset.



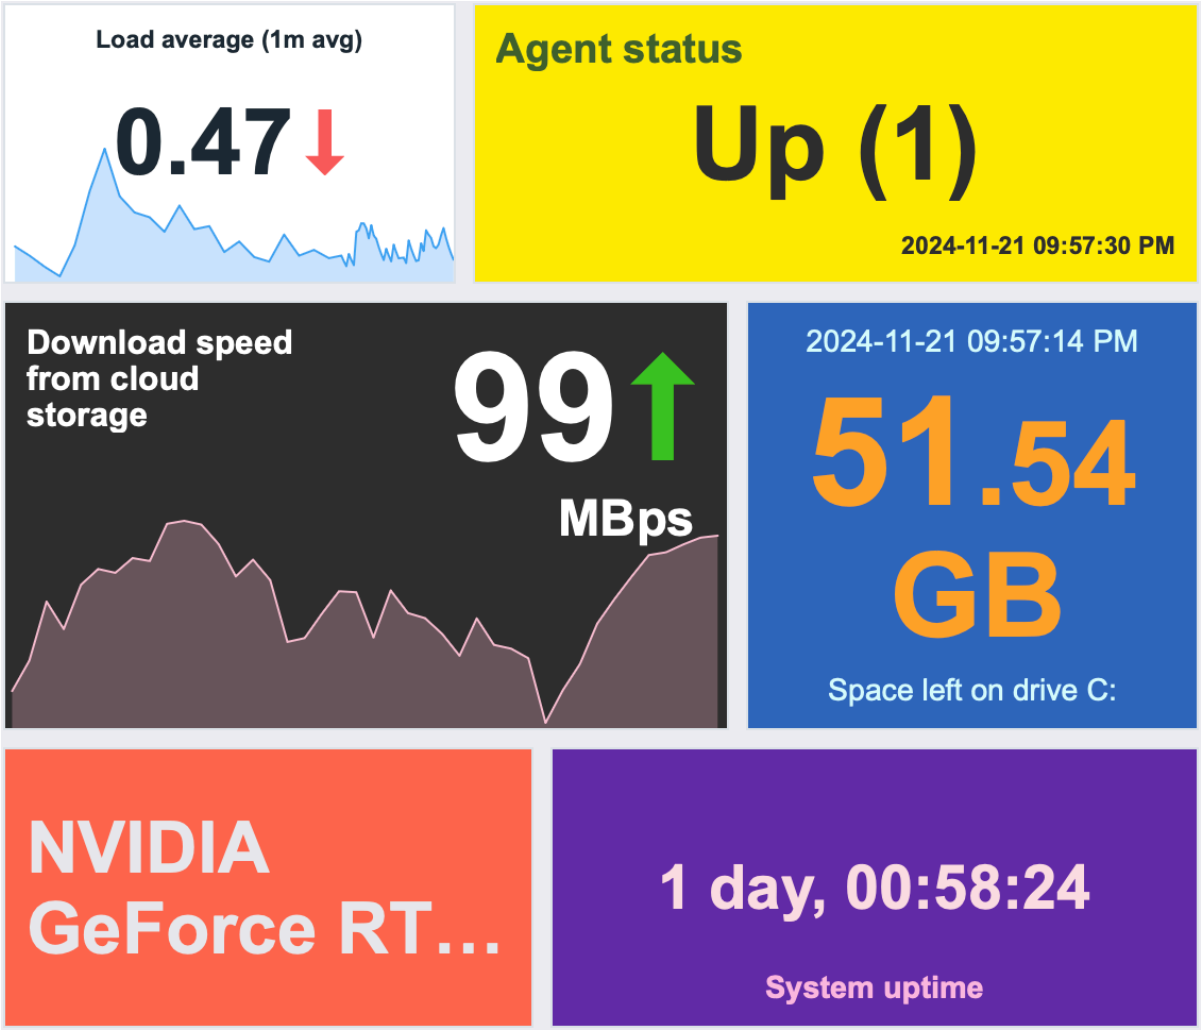
Gauge/Item value

The Gauge/Item value widgets display the value of a single item, useful for keeping an eye on important metrics, visualizing thresholds, and detecting sudden changes in the data.

Gauge widget:

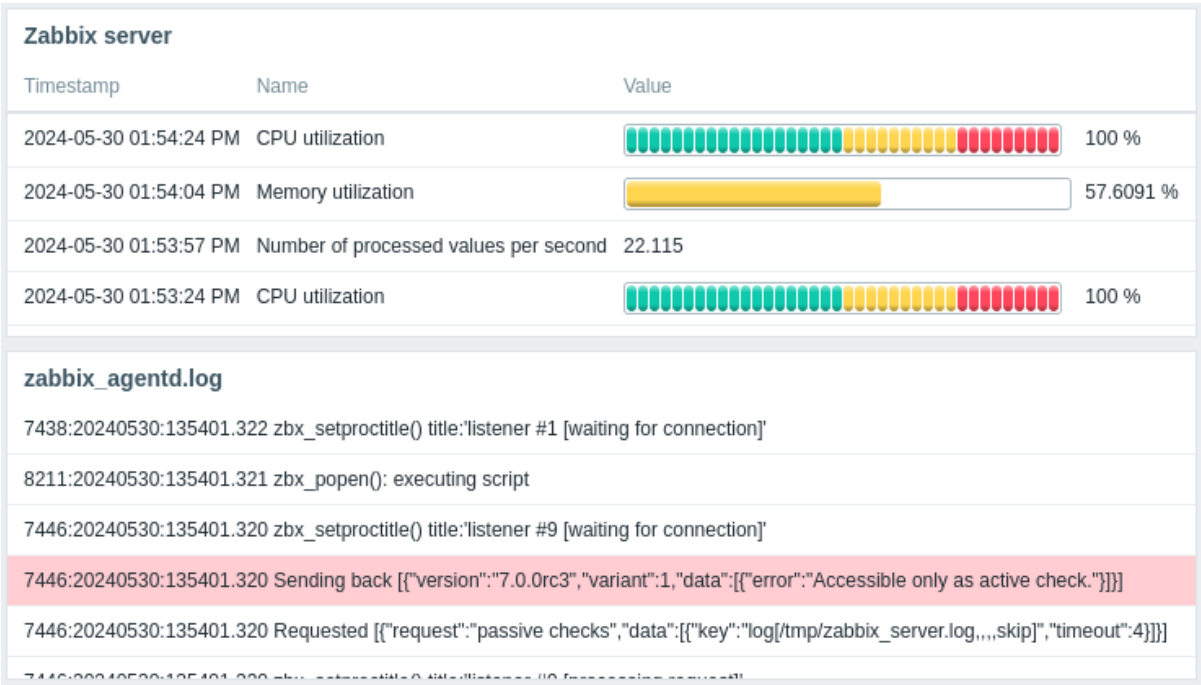


Item value widget:



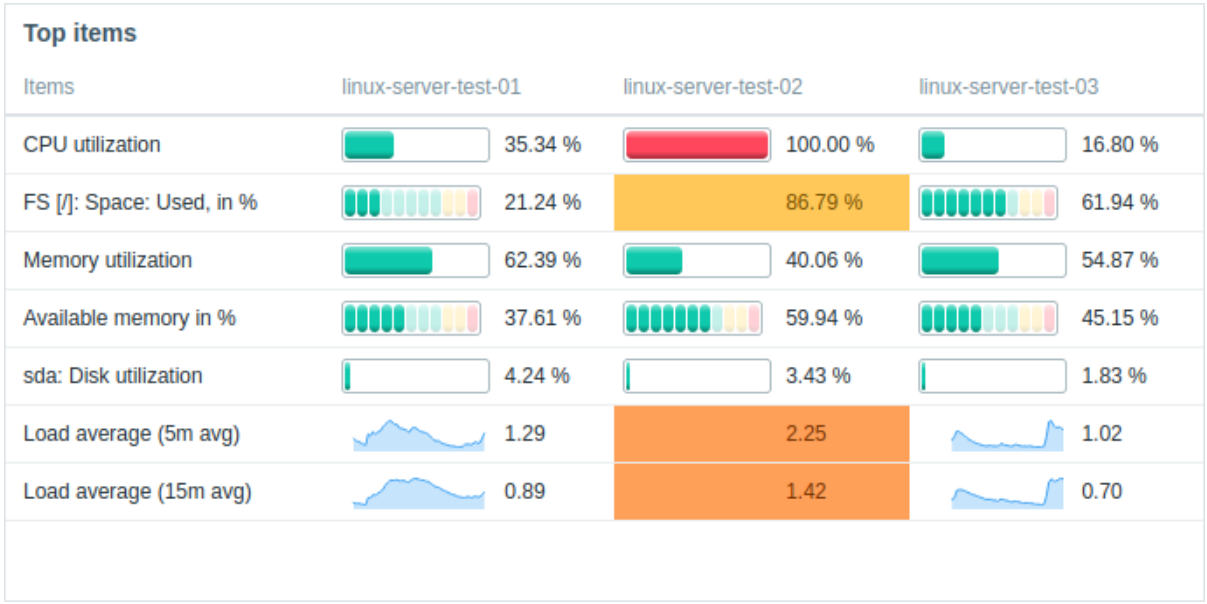
Item history

The Item history widget displays the latest data for all item types (numeric, text, etc.) in a table format. It can show progress bars and images (useful for browser items) and can highlight values (useful for log file monitoring).



Top items

The Top items widget displays the top/bottom values for selected items, providing a quick overview of their performance.

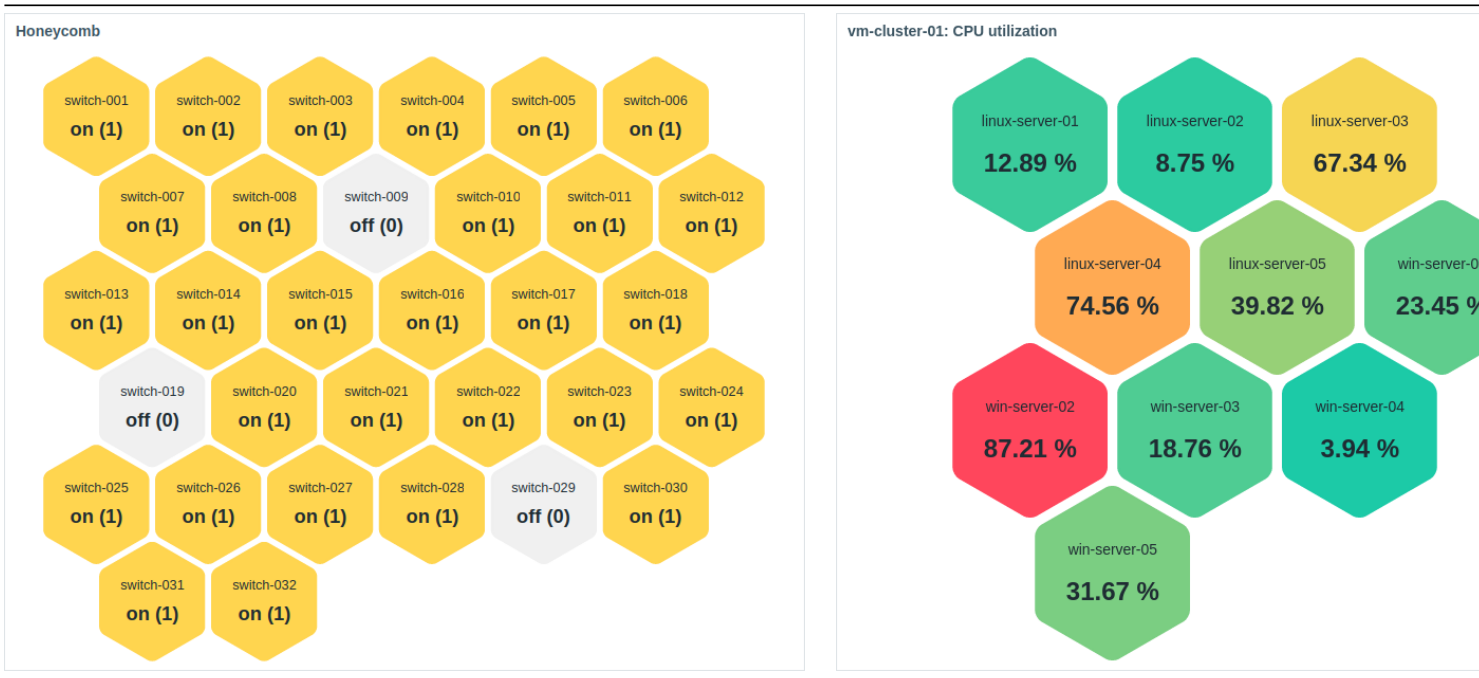


See also:

- Top hosts - displays the top/bottom values for selected items from multiple hosts.
- Top triggers - displays triggers with the highest number of problems.

Honeycomb

The Honeycomb widget displays hosts or items as hexagonal cells in a grid. This layout makes it easy to detect problem hosts/items, spot clusters of issues, compare groups at a glance, and more.

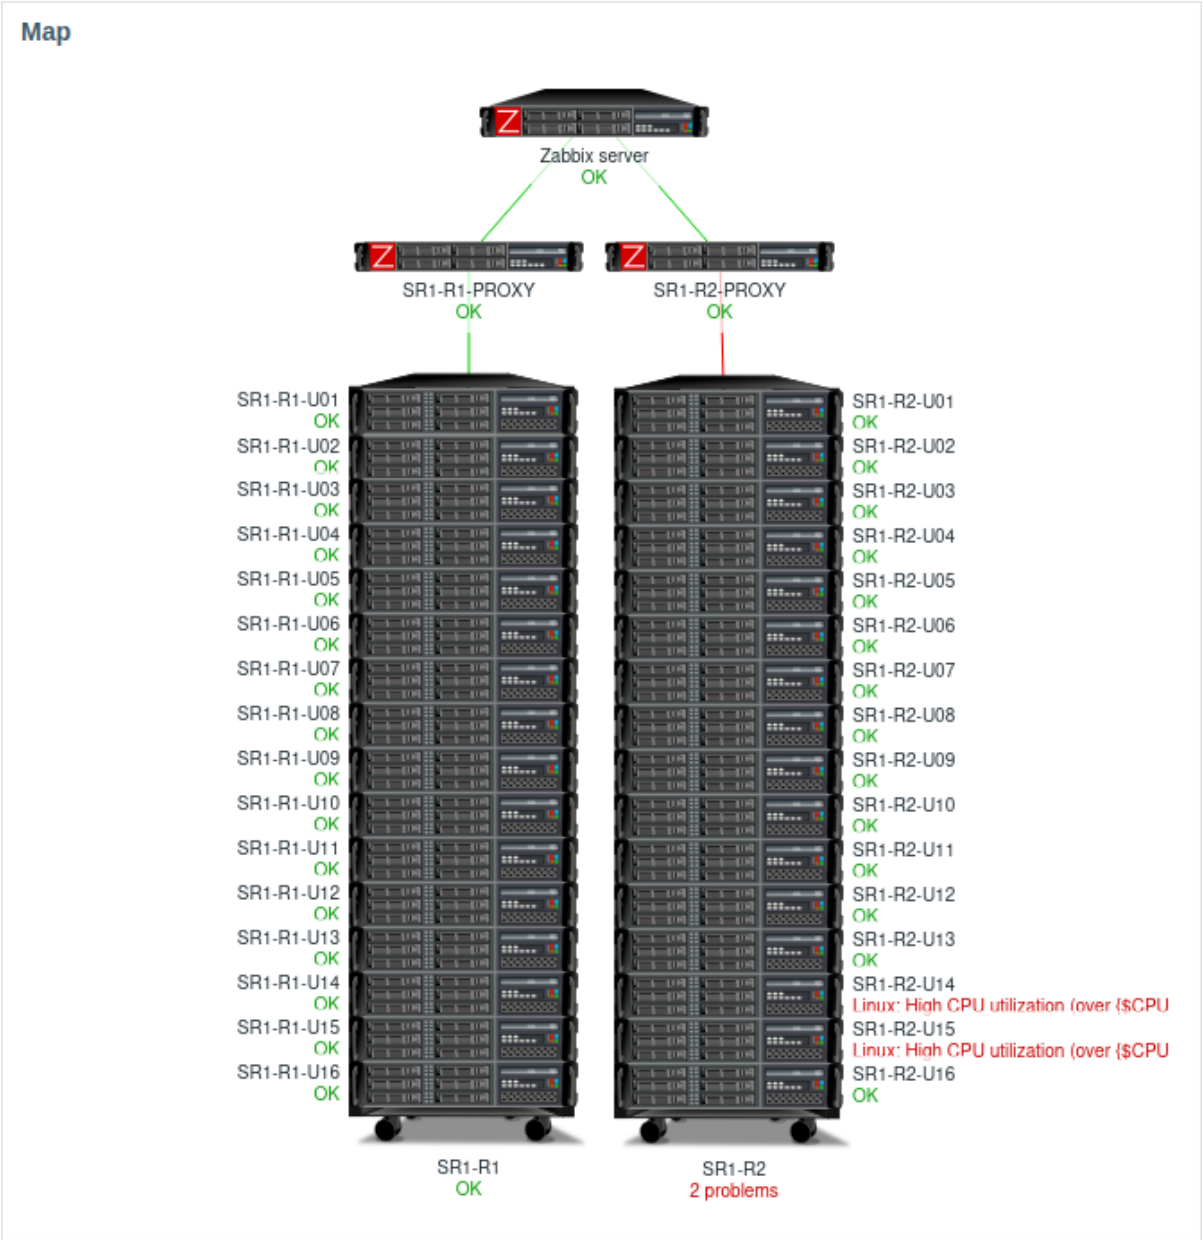


Maps

Map widgets visualize network topology or host locations, helping assess overall and location-specific health.

Map

The Map widget can display a network map or a similar visualization, providing you a dynamic overview of your network.

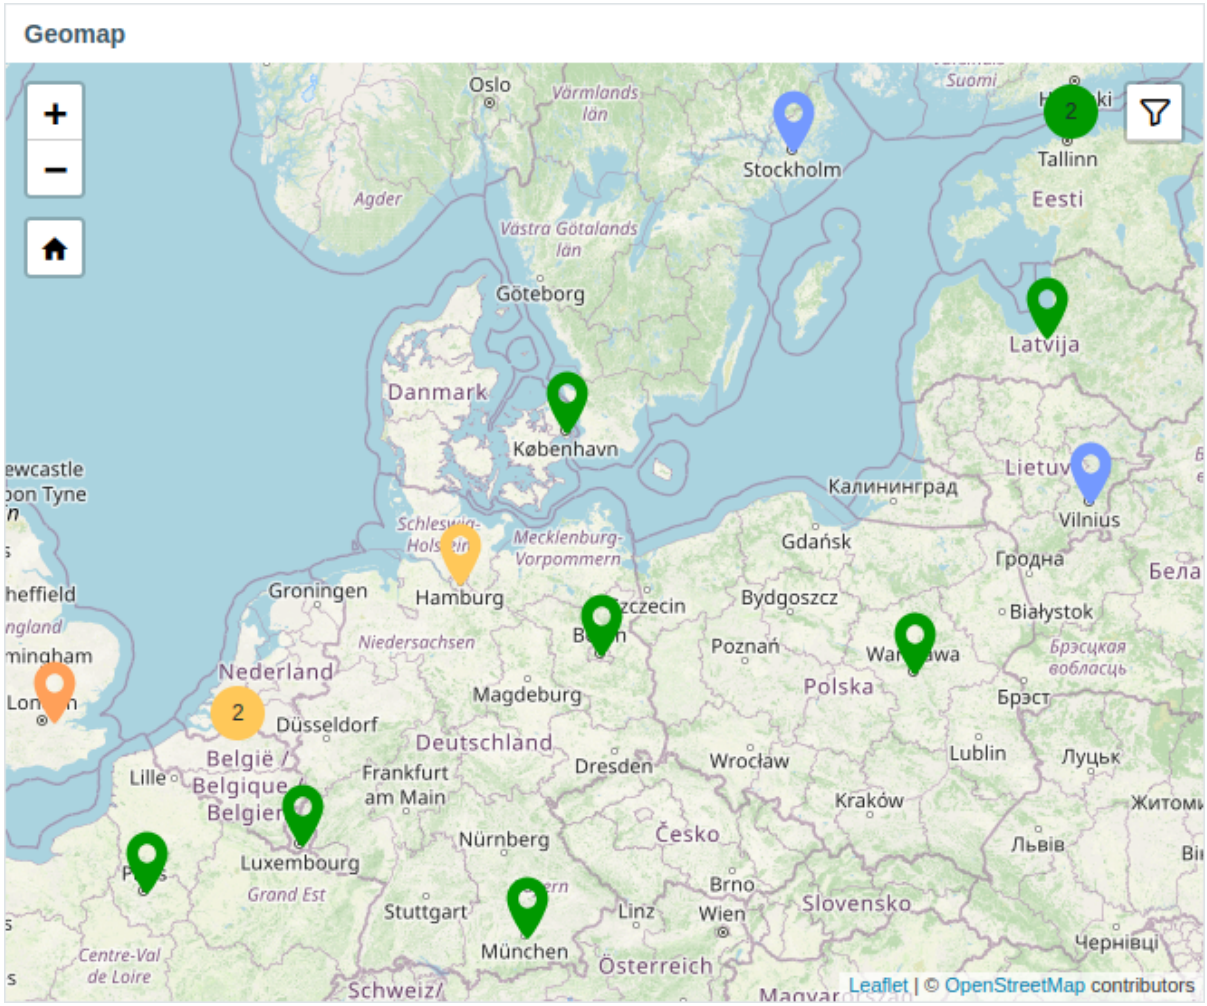


See also:

- Map navigation tree - displays a hierarchy of existing maps and the problem count for each map and map group.
- Favorite maps - displays a list of shortcuts to maps marked as favorites by the current user, sorted alphabetically.

Geomap

The Geomap widget displays hosts as markers on an interactive Leaflet-based map, with marker colors indicating each host's most severe problem.

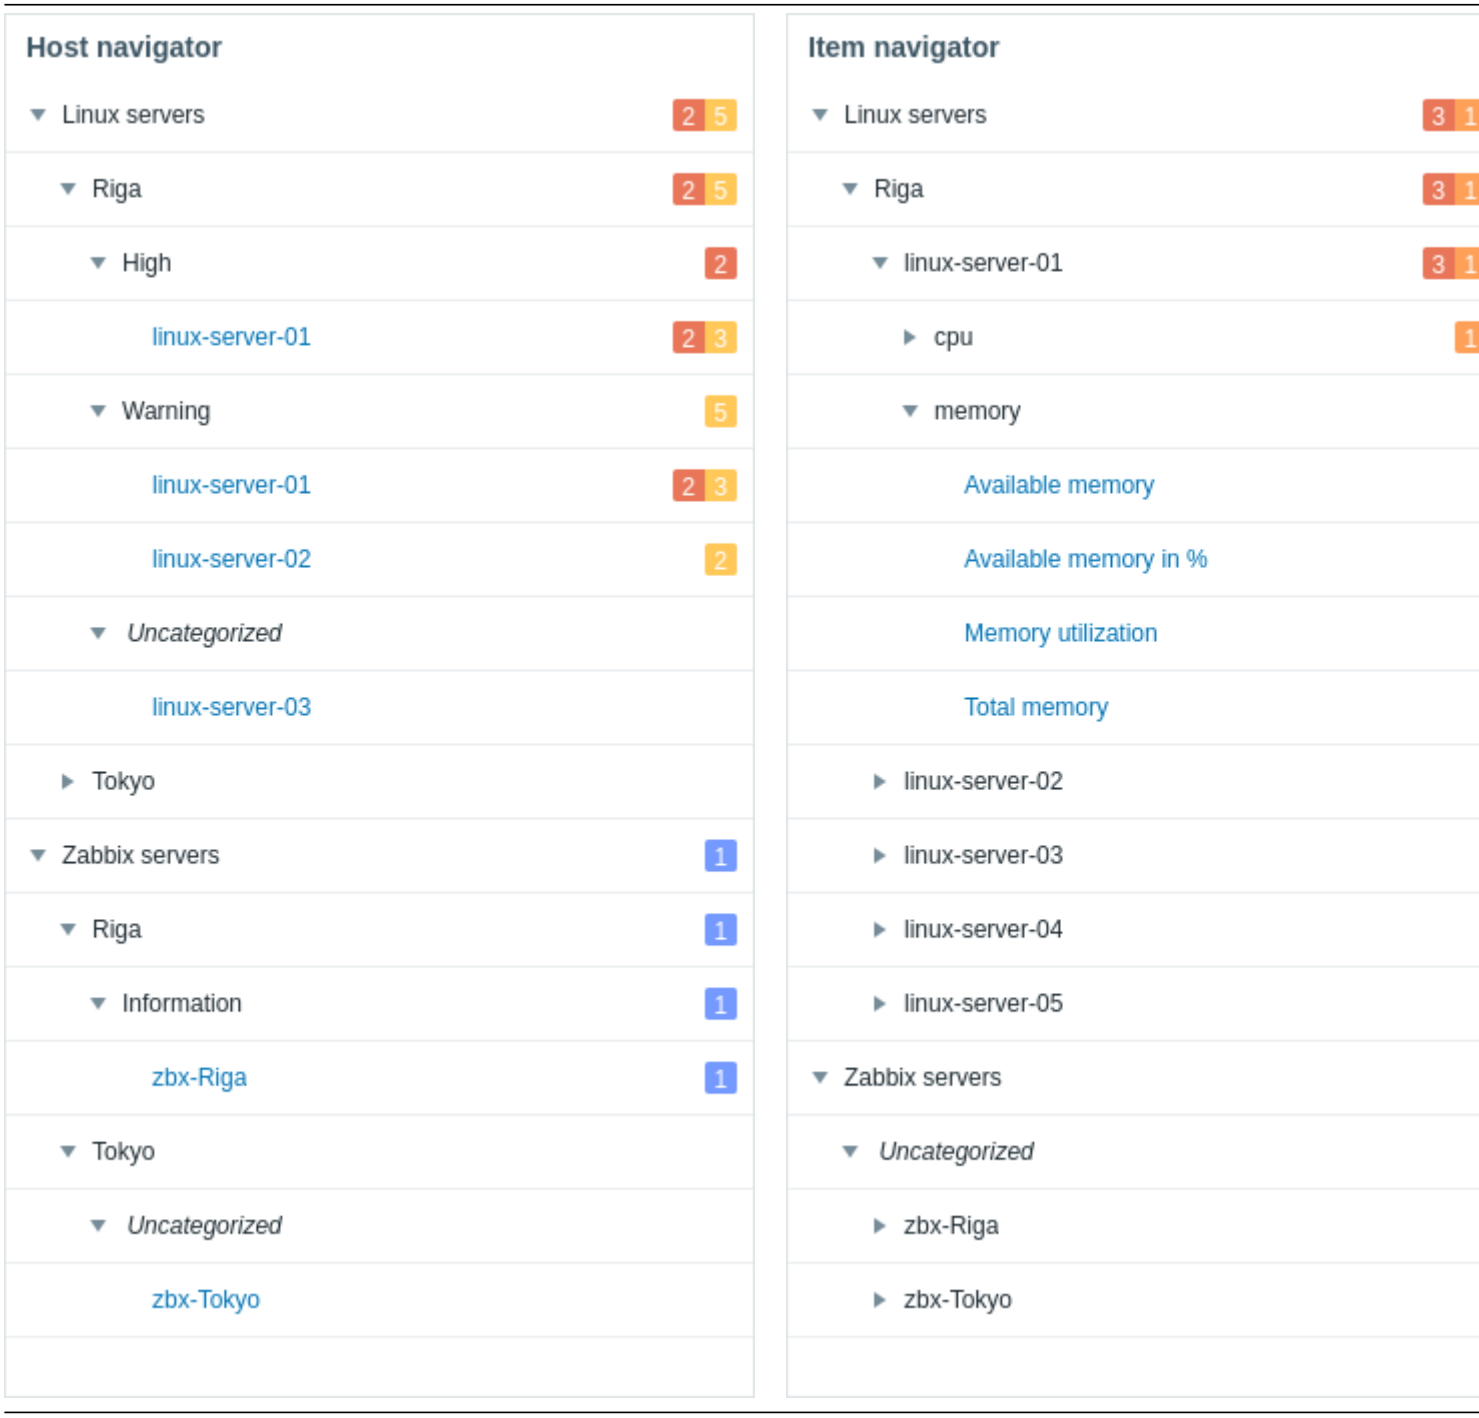


Host/Item details

These widgets provide lists and detailed views of hosts or items.

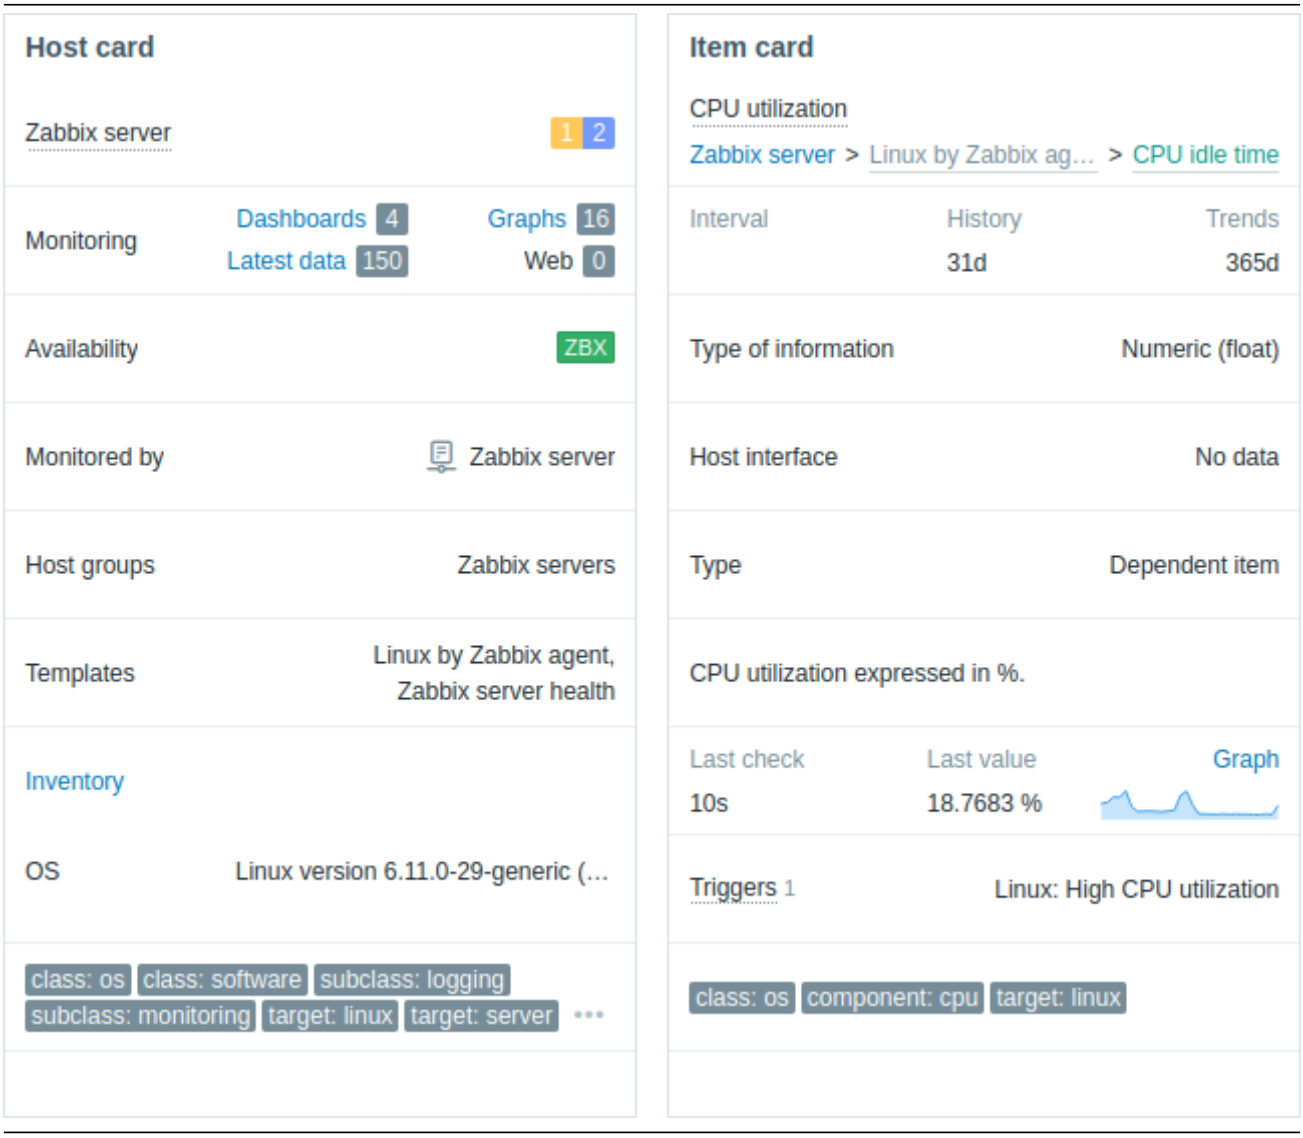
Host/Item navigator

The Host navigator/Item navigator widgets display a list of hosts/items based on various filtering and grouping options. These widgets are most useful for controlling what other widgets display based on the selected host/item.



Host/Item card

The Host card/Item card widgets display details of a single host/item, allowing you to quickly assess the status and configuration of a host or item (host availability, problems, inventory, item latest data, errors, description, etc.).



Status and reports

These widgets display status data (e.g., availability, triggers, problems) as well as reports.

Host availability

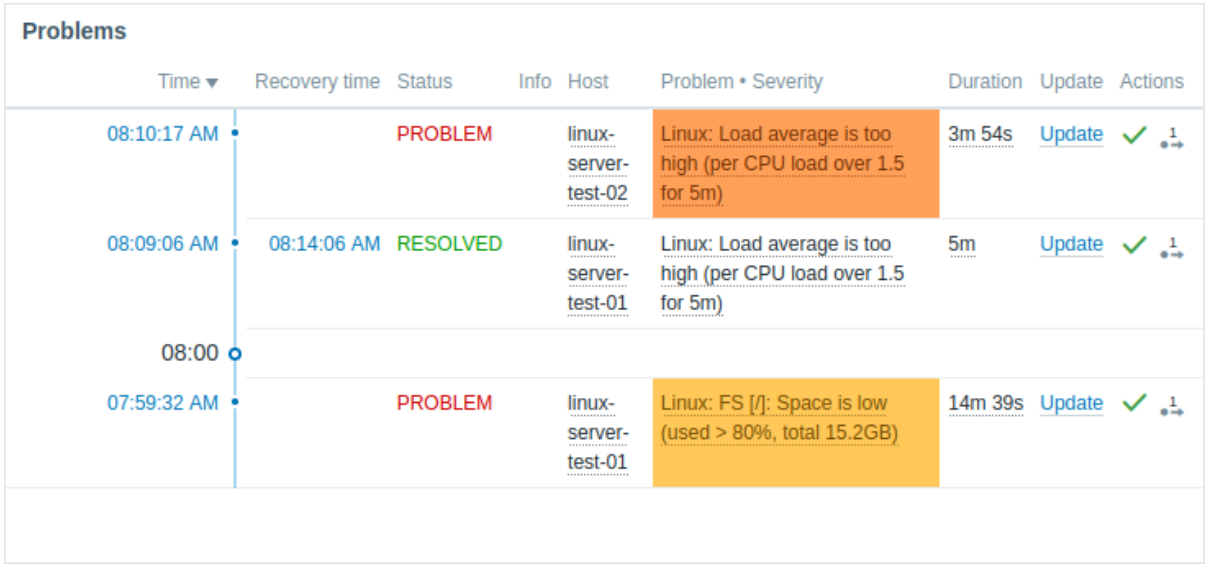
The Host availability widget displays host availability across selected host groups, allowing you to monitor the number of hosts that are up or down.

Host availability

	Total Hosts	Agent (active)	Agent (passive)	SNMP	JMX	IPMI
Available	10	3	8	2	0	0
Not available	2	0	3	0	0	0
Mixed	1	-	0	0	0	0
Unknown	4	0	4	0	0	0
Total	17	3	15	2	0	0

Problems

The Problems widget displays current problems filtered by different parameters (host groups, hosts, problem names, etc.), giving you a clear view of what needs attention.



See also:

- Problems by severity - displays the number of cause problems grouped by severity.
- Problem hosts - displays the number of cause problems per host group and shows the highest problem severity in each group.

Trigger overview

The Trigger overview widget displays the current states of host triggers as a table of colored blocks. It provides a quick visual snapshot of trigger health and activity on different hosts.



Action log

The Action log widget displays details of operations (notifications, remote commands) executed within an action. The detailed log helps with auditing, troubleshooting, and monitoring each action's execution.

Action log

Time ▾	Action	Media type	Recipient	Message	Status	Info
2025-08-15 07:58:30 AM	Report problems to Zabbix administrators	Email	Admin (Zabbix Administrator) *admin@example.com*	**Resolved in 1m 59s: Linux: Load average is too high (per CPU load over 1.5 for 5m)** Problem has been resolved at 10:58:30 on 2025.08.15 Problem name: Linux: Load average is too high (per CPU load over 1.5 for 5m) Problem duration: 1m 59s Host: Linux server Severity: Average Original problem ID: 45	Sent	
2025-08-15 07:56:30 AM	Report problems to Zabbix administrators	Email	Admin (Zabbix Administrator) *admin@example.com*	**Problem: Linux: Load average is too high (per CPU load over 1.5 for 5m)** Problem started at 10:56:30 on 2025.08.15 Problem name: Linux: Load average is too high (per CPU load over 1.5 for 5m) Host: Linux server Severity: Average Operational data: Load averages(1m 5m 15m): (6.256836 4.808105 3.398438), # of CPUs: 2 Original problem ID: 45	Sent	
2025-08-15 07:42:30 AM	Report problems to Zabbix administrators		Admin (Zabbix Administrator)		Failed	i

SLA report

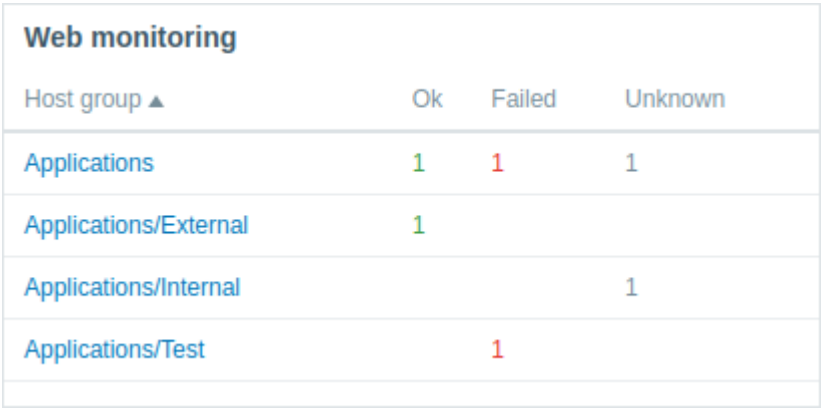
The SLA report widget displays SLA reports, helping you monitor service performance against defined targets.

SLA report

Day	SLO	SLI	Uptime	Downtime	Error budget	Excluded downtimes
2025-09-04	99.9%	99.8437	5h 40m 50s	28s	-8s	
2025-09-03	99.9%	99.9479	23h 59m 15s	45s	41s	
2025-09-02	99.9%	99.9190	23h 58m 50s	1m 10s	16s	
2025-09-01	99.9%	99.9251	22h 58m 58s	1h 1m 2s	21s	2025-09-01 01:00 AM Maintenance: 1h
2025-08-31	99.9%	99.6863	23h 55m 29s	4m 31s	-3m 5s	
2025-08-30	99.9%	99.9190	23h 58m 50s	1m 10s	16s	
2025-08-29	99.9%	99.7095	23h 55m 49s	4m 11s	-2m 45s	
2025-08-28	99.9%	99.9421	23h 59m 10s	50s	36s	
2025-08-27	99.9%	99.7188	23h 55m 57s	3m 15s	-2m 37s	
2025-08-26	99.9%	99.7755	23h 56m 46s	3m 14s	-1m 48s	
2025-08-25	99.9%	99.9094	22h 58m 45s	1h 1m 15s	8s	2025-08-25 01:00 AM Maintenance: 1h

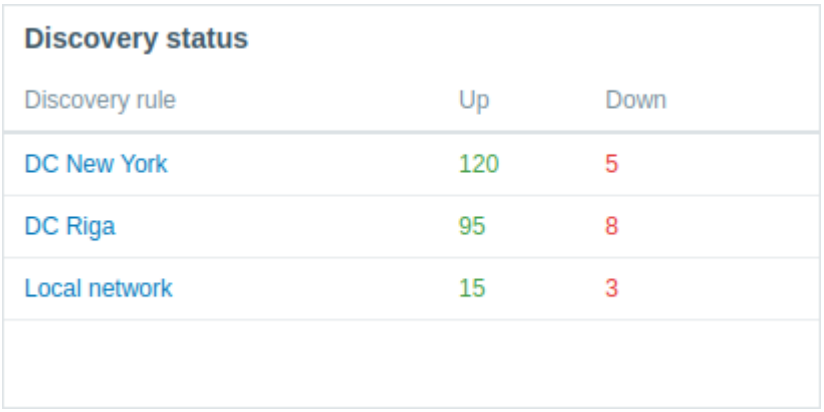
Web monitoring

The Web monitoring widget displays the status of active web scenarios.



Discovery status

The Discovery status widget displays the status of devices discovered by enabled network discovery rules. It helps you monitor whether network discovery is working correctly and the number of devices that are up or down.



Dashboard utilities and system information

These widgets provide supplementary information and tools to enhance your dashboards.

System information

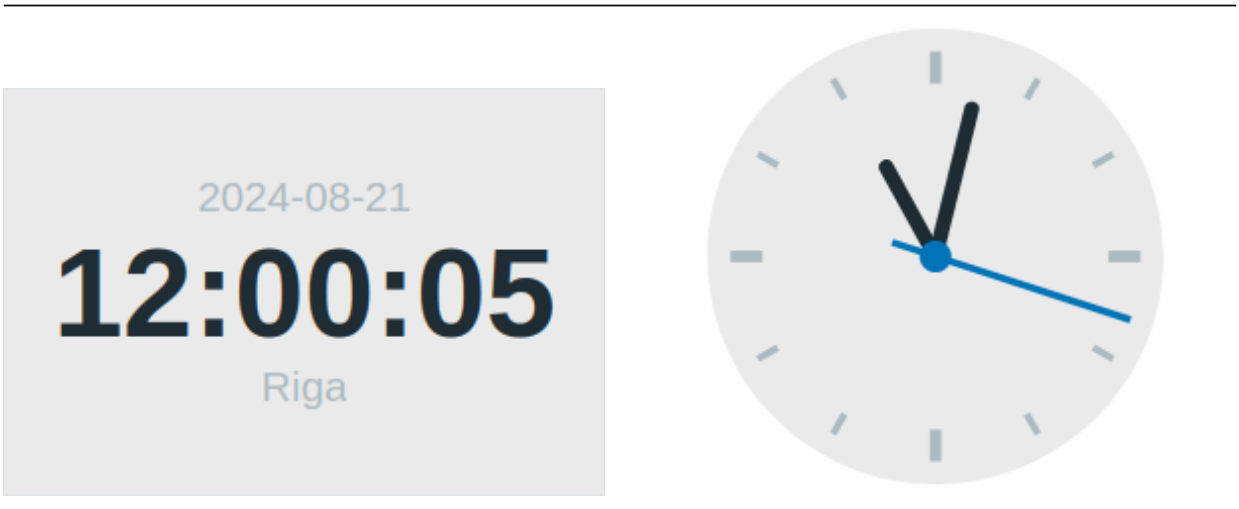
The System information widget displays a summary of key Zabbix server and system data or details of high availability nodes.

System information

Parameter	Value	Details
Zabbix server is running	Yes	localhost:10051
Zabbix server version	8.0.0	Up to date
Zabbix frontend version	8.0.0	Up to date
Software update last checked	2025-08-30	
Latest release	8.0.0	Release notes
Number of hosts (enabled/disabled)	2	2 / 0
Number of templates	312	
Number of items (enabled/disabled/not supported)	176	165 / 0 / 11
Number of triggers (enabled/disabled [problem/ok])	89	89 / 0 [4 / 85]
Number of users (online)	2	1
Required server performance, new values per second	2.68	
Global scripts on Zabbix server	Disabled	
High availability cluster	Disabled	

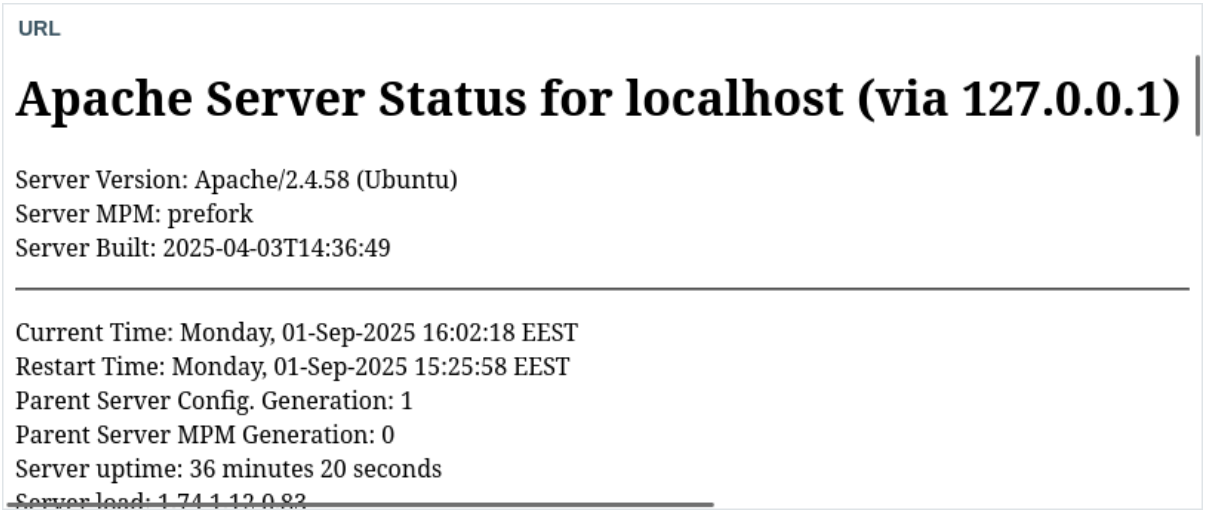
Clock

The Clock widget displays local, server, or host time in analog or digital format.



URL

The URL widget displays the content retrieved from a URL.



Dashboard widget parameters

Overview

This page describes parameters that are common to all dashboard widgets, as well as dynamic parameters that allow multiple widgets to exchange configuration data with each other or with the dashboard.

To see specific parameters for each widget, please use the sidebar to navigate to the corresponding widget page.

Common parameters

The following parameters are common for every single widget:

Name	Enter a widget name.
Refresh interval	Configure the default refresh interval. Default refresh intervals for widgets range from *No refresh* to *15 minutes* depending on the type of the widget. For example: - *No refresh* for *URL* widget; - *1 minute* for *Action log* widget; - *15 minutes* for *Clock* widget. Refresh intervals can be set to a default value for all users. Switch the dashboard to edit mode, click the edit a widget button and select the desired refresh interval from the dropdown list. Each user can also set their own widget refresh interval. In the dashboard view mode, click the three dots button on a widget and select the desired refresh interval from the dropdown list. Note that a user's unique refresh interval takes priority over the widget setting and is preserved even when the widget's setting is modified.
Show header	Mark the checkbox to permanently display the widget header. When unmarked, the header is hidden to save space and becomes visible only on widget mouseover (both in view and edit modes). The header is also semi-visible when dragging a widget to a new place.

Dynamic parameters

Multiple widget parameters (e.g., *Hosts*, *Override host*, *Time period*) support dynamic data sources. Instead of configuring static values, you can link these parameters to other widgets or the dashboard itself. Widgets are then automatically updated based on selections or broadcasts from other parts of the dashboard.

Widget compatibility

Some widgets can broadcast configuration data to other widgets, some can listen for data, and some can do both. For example:

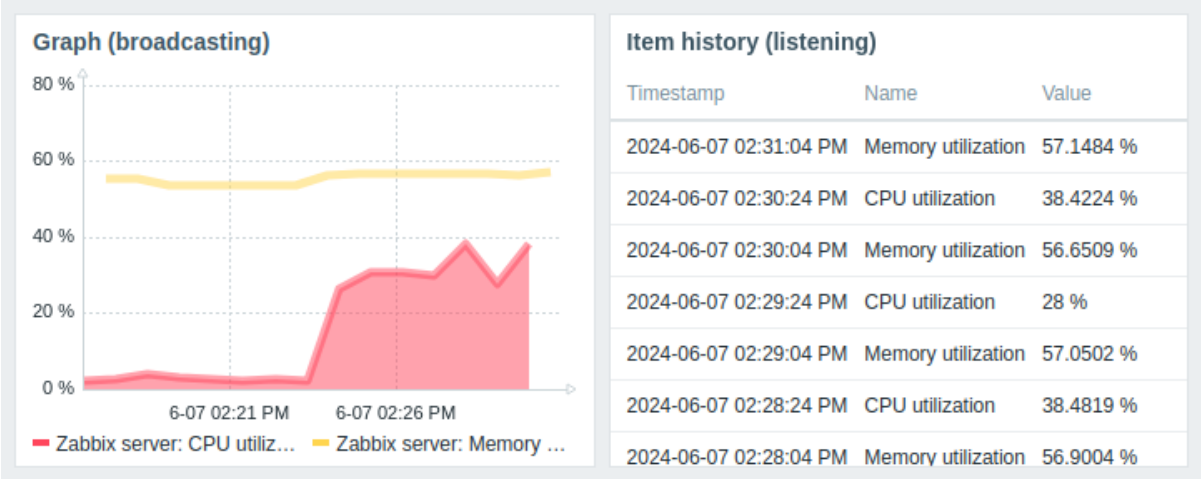
- The *Action log* widget can only retrieve time period data from *Graph*, *Graph (classic)*, and *Graph prototype* widgets.
- The *Geomap* widget can broadcast host data to widgets that listen for it (*Honeycomb*, *Top items*, etc.) and can also listen for host group and host data from widgets that broadcast it (*Honeycomb*, *Problem hosts*, etc.).
- The *Clock* widget cannot broadcast or listen for any data.

The following table outlines the broadcasting and listening capabilities of each widget.

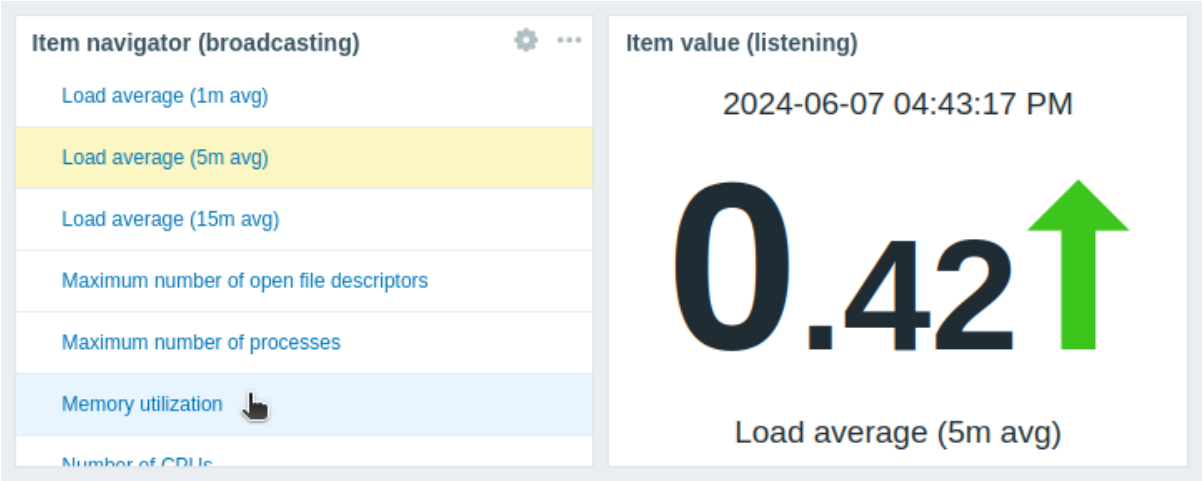
Widget	Broadcasts	Listens
Action log	-	Time period
Clock	-	-
Discovery status	-	-
Favorite graphs	-	-
Favorite maps	-	-
Gauge	-	Hosts, Items
Geomap	Hosts	Host groups, Hosts
Graph	Time period, Items	Hosts, Items, Time period
Graph (classic)	Time period	Hosts, Items, Graphs, Time period
Graph prototype	Time period	Hosts, Time period
Honeycomb	Hosts, Items	Host groups, Hosts
Host availability	-	Host groups
Host card	-	Hosts
Host navigator	Hosts	Host groups
Item card	-	Hosts, Items, Time period
Item history	Items	Hosts, Time period
Item navigator	Items	Host groups, Hosts
Item value	-	Hosts, Items, Time period
Map	Host groups, Hosts	Maps
Map navigation tree	Maps	-
Pie chart	-	Items, Time period
Problem hosts	Host groups	Host groups, Hosts
Problems	Events	Host groups, Hosts
Problems by severity	Host groups	Host groups, Hosts
Scatter plot	Time period, Items	Hosts, Items, Time period
SLA report	-	-
System information	-	-
Top hosts	Hosts	Host groups, Hosts
Top items	-	Host groups, Hosts
Top triggers	-	Time period
Trigger overview	-	Host groups, Hosts
URL	-	Hosts
Web monitoring	Host groups	Host groups, Hosts

Widgets differ in how they **broadcast** data to other widgets.

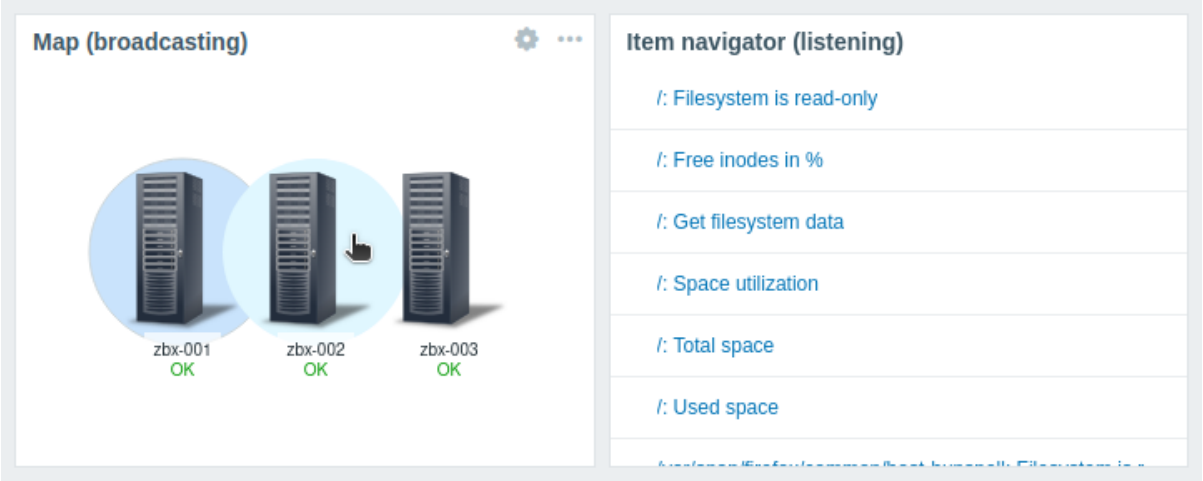
Widgets that support broadcasting start doing it automatically when created. For example, the *Graph* widget immediately broadcasts time period data to listening widgets.



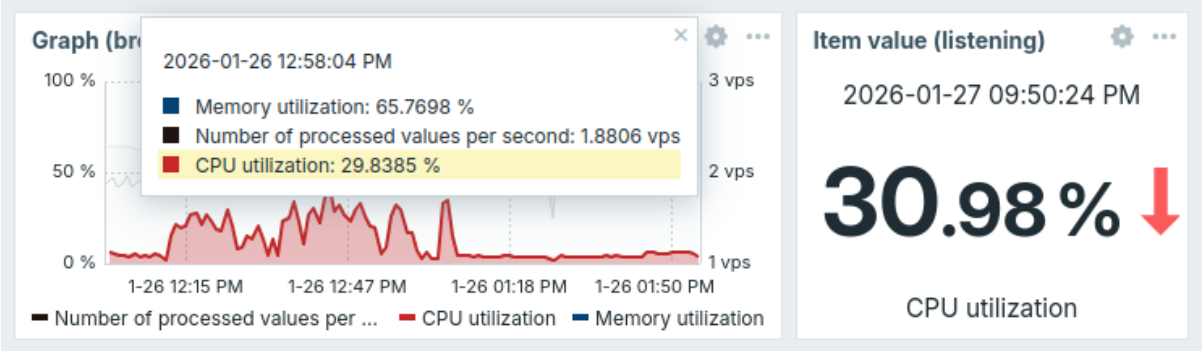
Widgets that support entity selection automatically broadcast data for the first available entity. For example, the *Item navigator* widget broadcasts data for the first item in its item list. When a different item is selected, it broadcasts data for that item. On mouseover, the item is highlighted in light blue; on selection, it is highlighted in yellow.



The *Map* widget behaves similarly, broadcasting data for the first element closest to the upper-left corner of the widget's visible area (for the *Geomap* widget, closest to the center). When a different element is selected, it broadcasts data for that element. On mouseover, the element is highlighted in light blue; on selection, it is highlighted in dark blue.

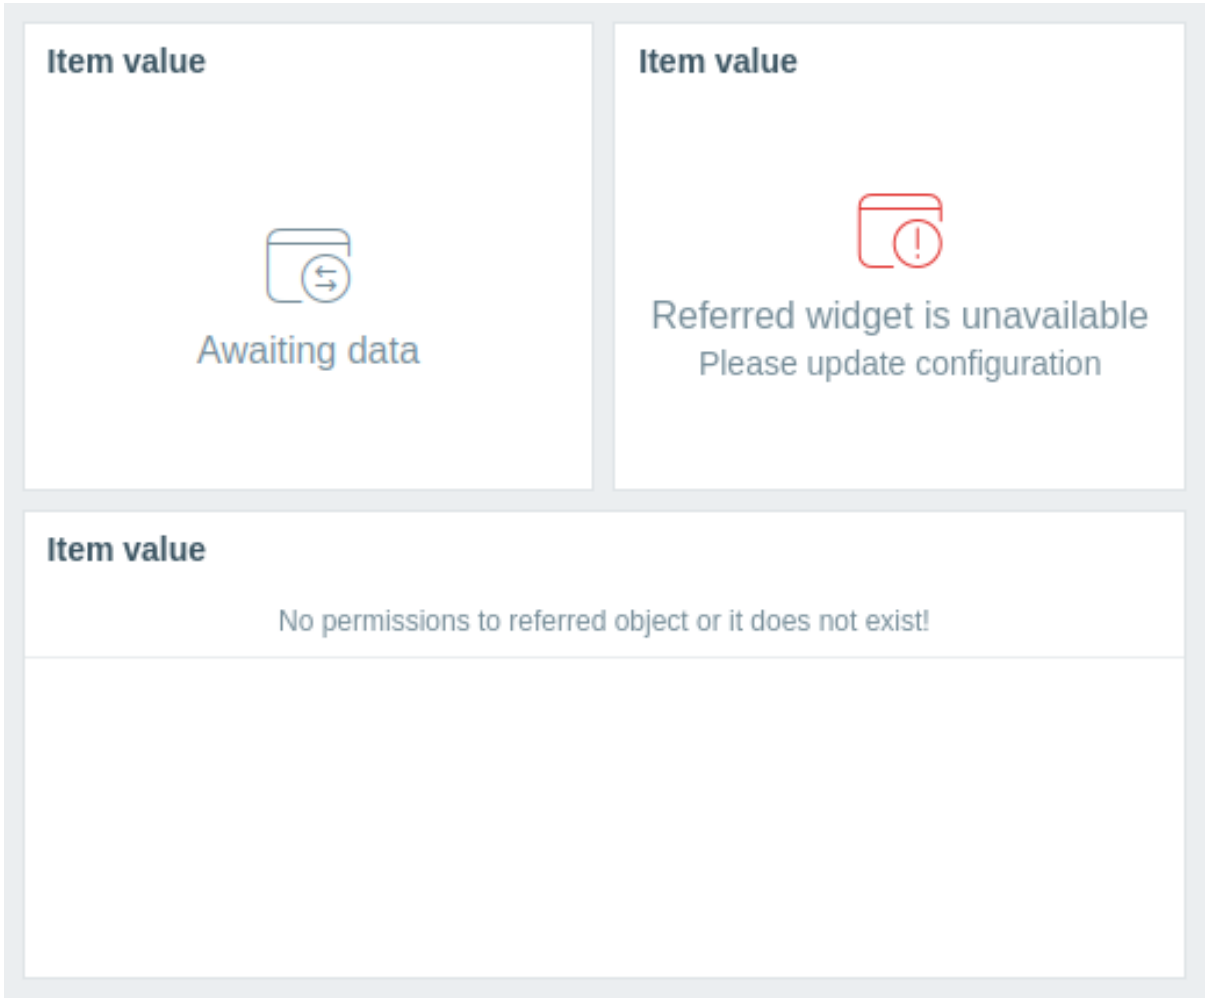


Widgets such as the *Graph* and *Scatter plot* can broadcast item data when an item is selected in their tooltips:



Widgets also have specific behaviors when **listening** for data from other widgets:

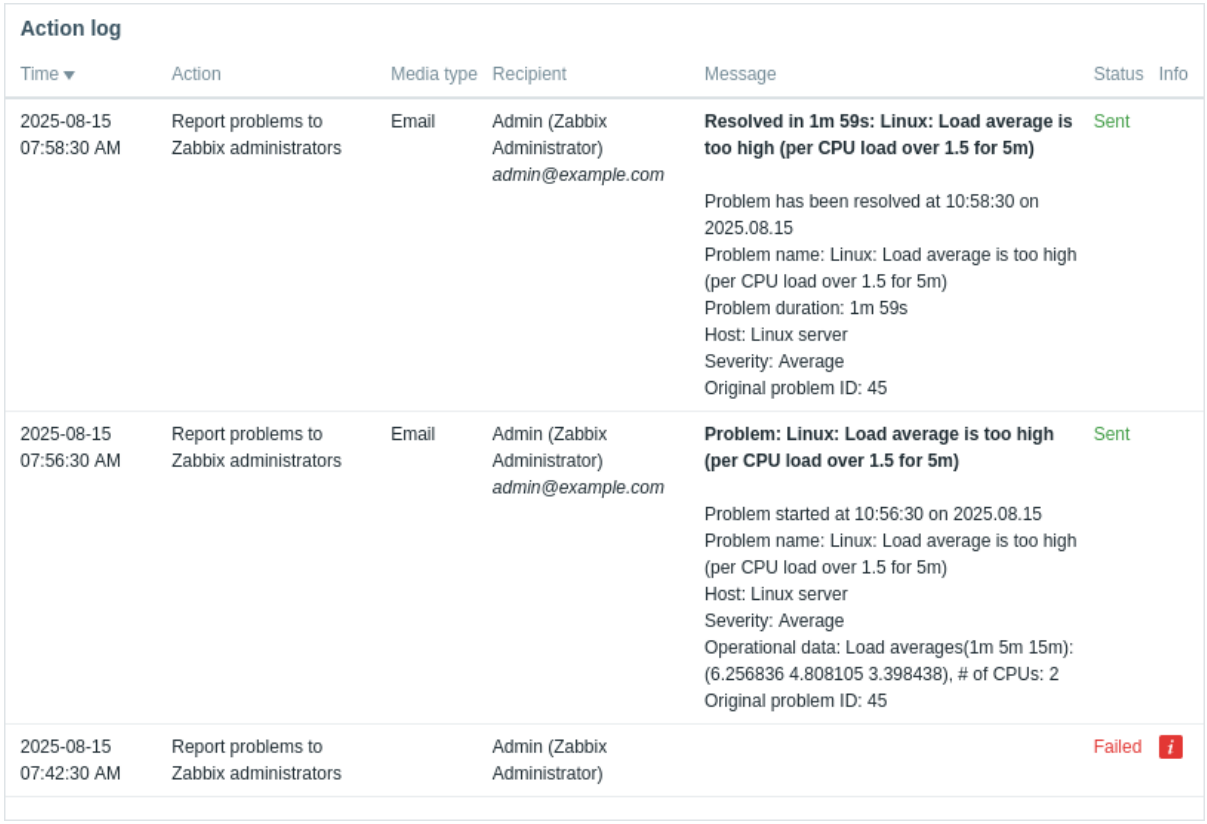
- If the data source widget is not broadcasting data, the listening widget enters the *Awaiting data* state.
- If the data source widget has been deleted, replaced with an incompatible widget, or moved to another dashboard page, the listening widget enters the *Referred widget is unavailable* state.
- If the specified host in the data source (widget or dashboard) lacks the entity configured in the listening widget (item, graph, map, etc.) or if the user lacks permissions to access the host, the listening widget displays the following message: *"No permissions to referred object or it does not exist!"*



Action log

Overview

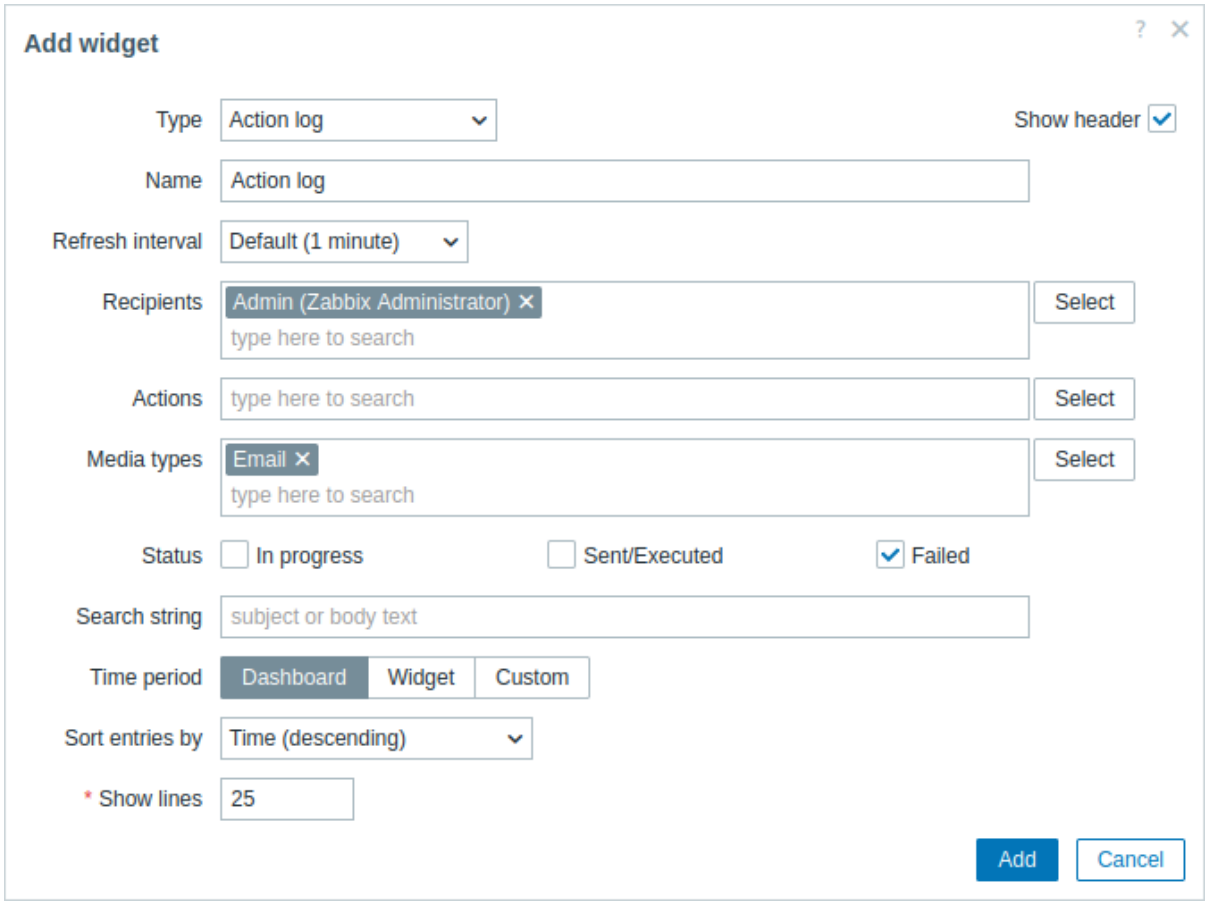
The *Action log* widget displays details of operations (notifications, remote commands) executed within an action. The detailed log helps with auditing, troubleshooting, and monitoring each action's execution.



It displays the same data as *Reports* > *Action log* and can display up to 1000 records.

Configuration

To configure, select *Action log* as type:



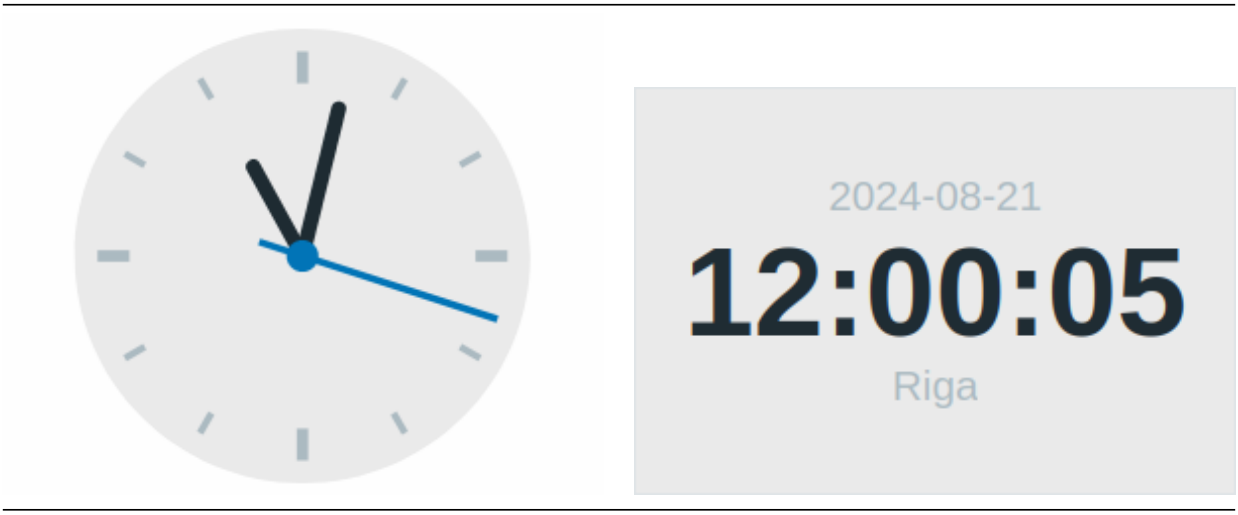
In addition to the parameters that are common for all widgets, you may set the following specific options:

Recipients	Filter entries by recipients. This field is auto-complete, so starting to type the name of a recipient will offer a dropdown of matching recipients. If no recipients are selected, details of action operations for all recipients will be displayed.
Actions	Filter entries by actions. This field is auto-complete, so starting to type the name of an action will offer a dropdown of matching actions. If no actions are selected, details of action operations for all actions will be displayed.
Media types	Filter entries by media types. This field is auto-complete, so starting to type the name of a media type will offer a dropdown of matching media types. If no media types are selected, details of action operations for all media types will be displayed.
Status	Mark the checkbox to filter entries by the respective status: **In progress** - action operations that are in progress are displayed; **Sent/Executed** - action operations that have sent a notification or have been executed are displayed; **Failed** - action operations that have failed are displayed.
Search string	Filter entries by the content of the message/remote command. If you enter a string here, only those action operations whose message/remote command contains the entered string will be displayed. Macros are not resolved.
Time period	Filter entries by time period. Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period.
Widget	Enter or select a compatible widget as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
Sort entries by	Sort entries by: **Time** (descending or ascending); **Type** (descending or ascending); **Status** (descending or ascending); **Recipient** (descending or ascending).
Show lines	Set how many action log lines will be displayed in the widget.

Clock

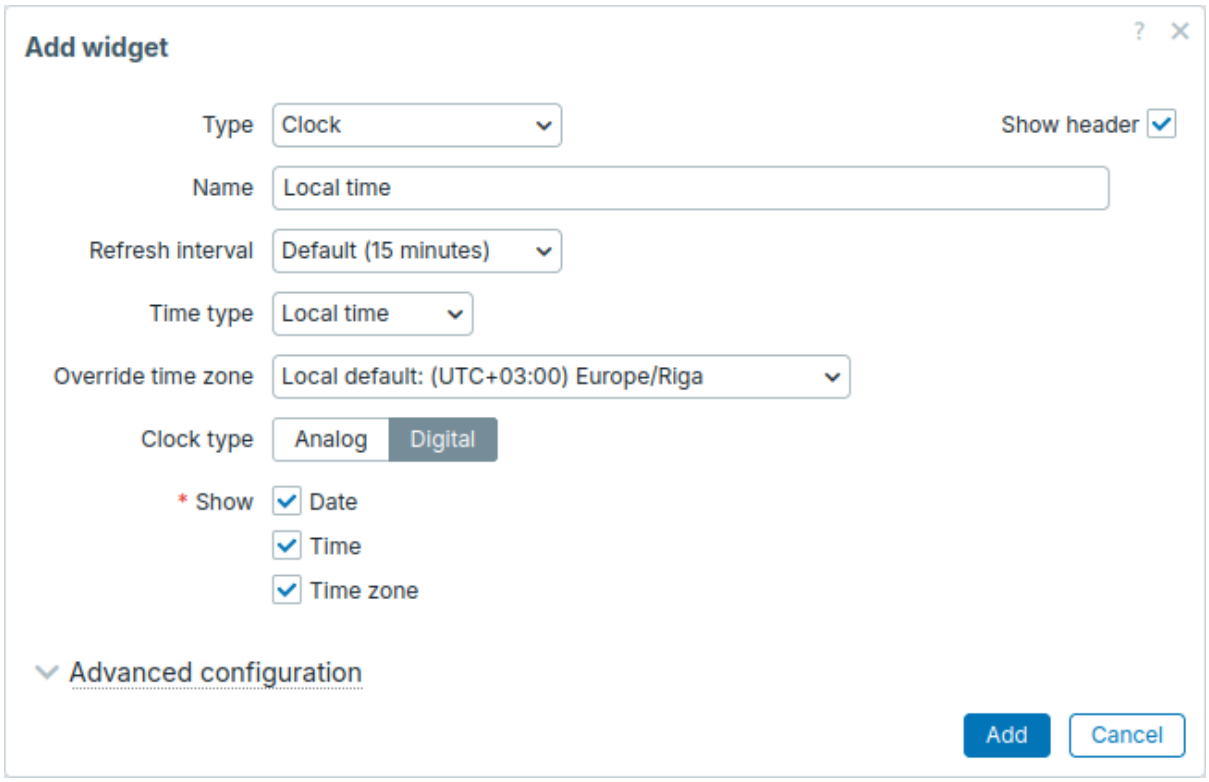
Overview

The *Clock* widget displays local, server, or host time in analog or digital format.



Configuration

To configure, select *Clock* as type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

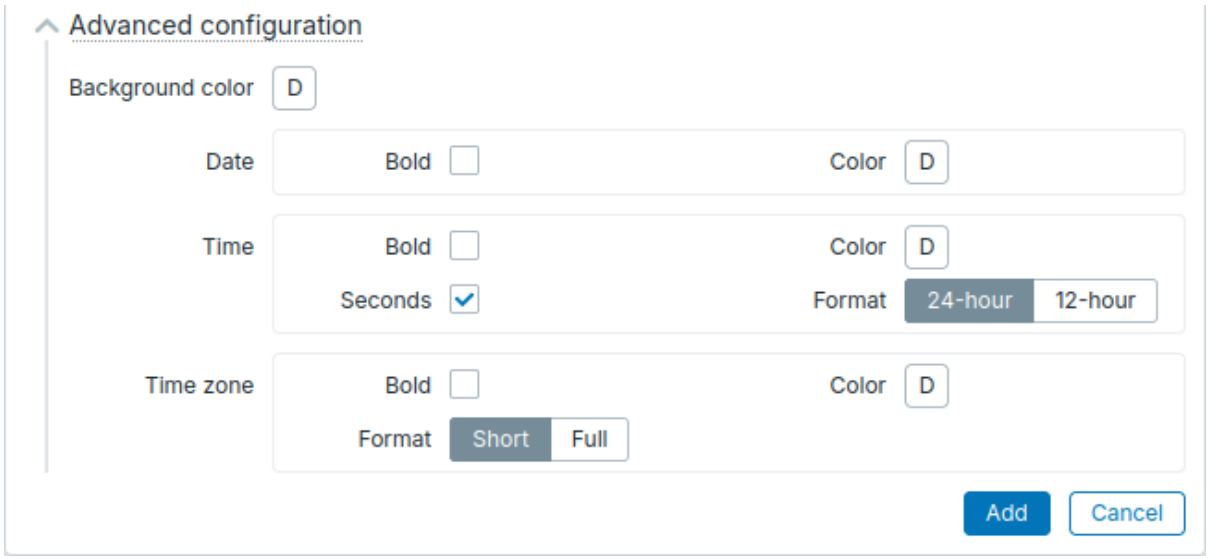
Time type	Select the time source: **Local time** - system time of the device running the browser; **Server time** - time zone set globally or for the Zabbix user; **Host time** - time retrieved from a host item.
Item	Select the item with the `system.localtime[local]` item key to display the host time. This parameter is available if *Time type* is set to "Host time".
Clock type	Select the clock type (analog or digital).
Override time zone	Select the time zone. This parameter is available for the *Local time* and *Server time* time types.
Show	Select one or more information units (date, time, time zone) to display in the digital clock. The unit size automatically adapts to the displayed units, widget size, and display scaling. This parameter is available if *Clock type* is set to "Digital".

Advanced configuration	Click the *Advanced configuration* header to display advanced configuration options for the digital clock. This section is available if *Clock type* is set to ”Digital”.

Advanced configuration

Advanced configuration options are available in the collapsible *Advanced configuration* section, and only for those elements that are selected in the *Show* field (see above).

Additionally, advanced configuration allows you to change the background color for the whole widget.

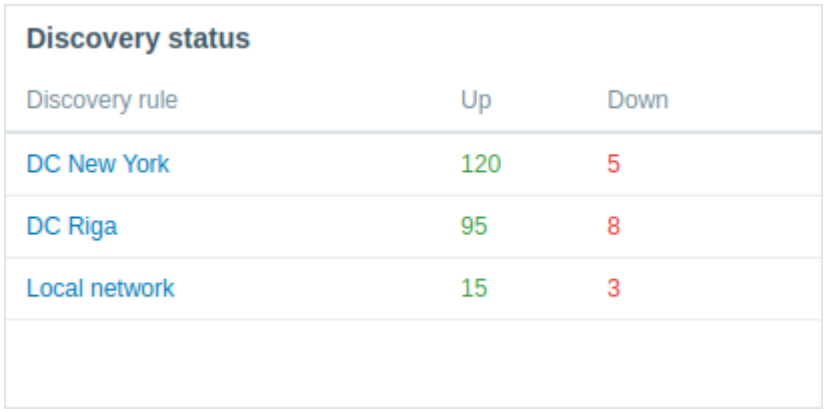


Background color		Select the background color from the color picker. D stands for default color (depends on the frontend theme). To return to the default value, click the *Use default* button in the color picker.
Date		This section is available when the *Date* checkbox is selected for the *Show* parameter.
	Bold	Mark the checkbox to display date in bold type.
	Color	Select the date color from the color picker. D stands for default color (depends on the frontend theme). To return to the default value, click the *Use default* button in the color picker.
Time		This section is available when the *Time* checkbox is selected for the *Show* parameter.
	Bold	Mark the checkbox to display time in bold type.
	Color	Select the time color from the color picker. D stands for default color (depends on the frontend theme). To return to the default value, click the *Use default* button in the color picker.
	Seconds	Mark the checkbox to display seconds. Otherwise, only hours and minutes will be displayed.
	Format	Select to display a 24-hour or 12-hour time.
Time zone		This section is available when the *Time zone* checkbox is selected for the *Show* parameter.
	Bold	Mark the checkbox to display time zone in bold type.
	Color	Select the time zone color from the color picker. D stands for default color (depends on the frontend theme). To return to the default value, click the *Use default* button in the color picker.
	Format	Select to display time zone in short format (e.g., `New York`) or full format (e.g.,`(UTC-04:00) America/New York`).

Discovery status

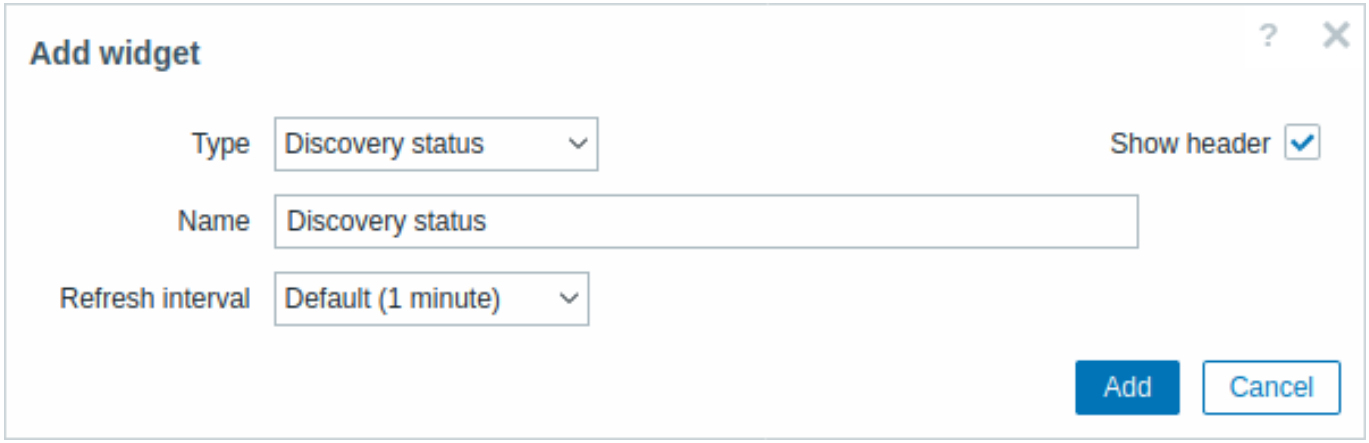
Overview

The *Discovery status* widget displays the status of devices discovered by enabled network discovery rules. It helps monitor whether network discovery is working correctly and the number of devices that are up or down.



Configuration

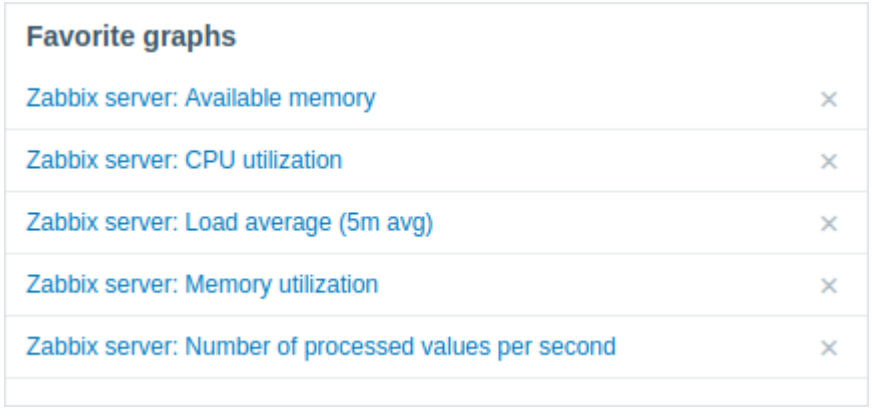
All configuration parameters for this widget are common to all widgets.



Favorite graphs

Overview

The *Favorite graphs* widget displays an alphabetical list of shortcuts to graphs marked as favorites by the current user.



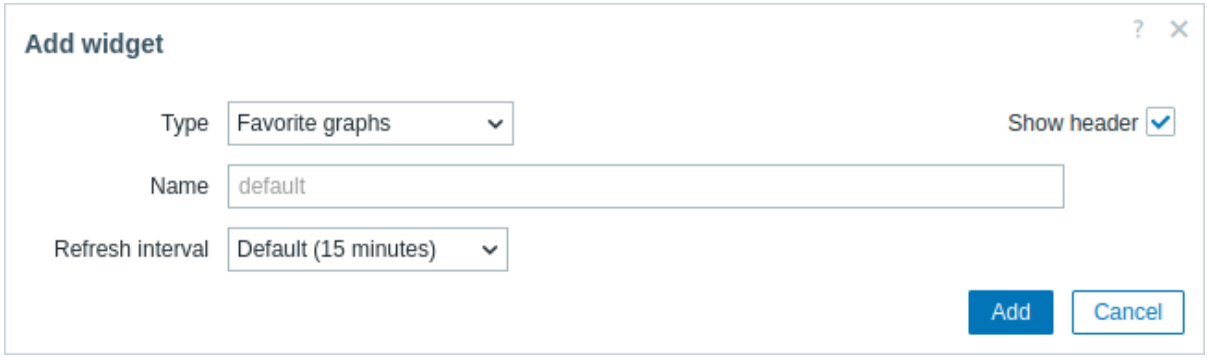
Clicking x removes the graph from the list and from favorites.

Graphs are marked as favorites by clicking the  *Add to favorites* button when viewing item graphs in *Monitoring* > *Latest data*.



Configuration

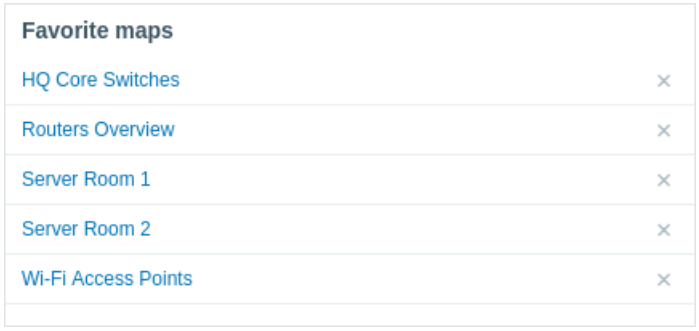
All configuration parameters for this widget are common to all widgets.



Favorite maps

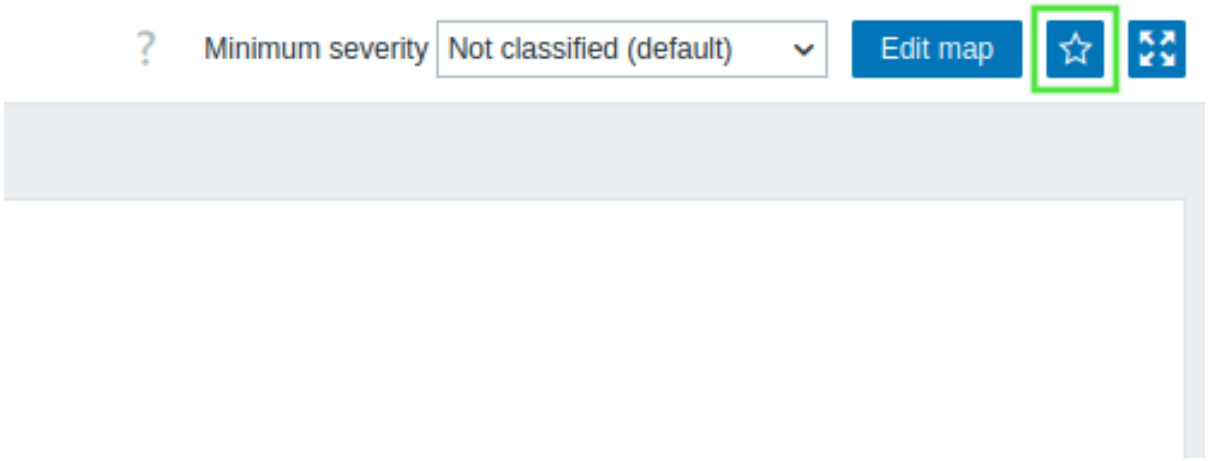
Overview

The *Favorite maps* widget displays a list of shortcuts to maps marked as favorites by the current user, sorted alphabetically.



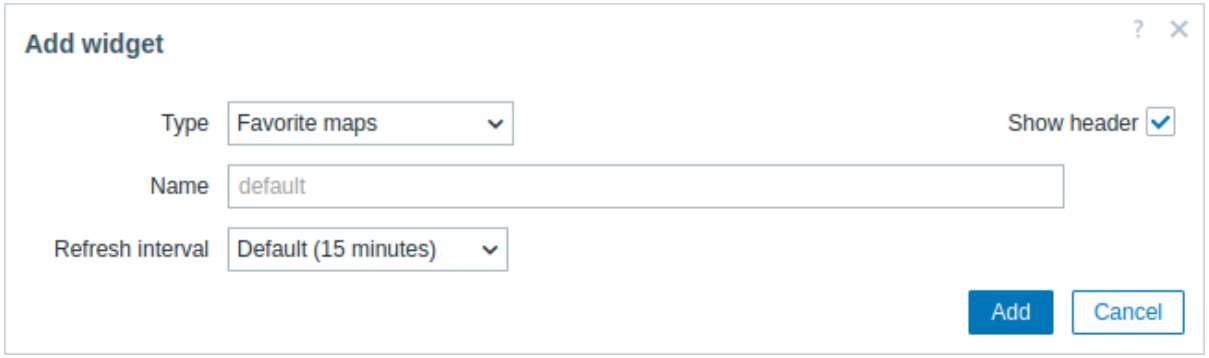
Clicking x removes the map from favorites.

Maps are marked as favorites by clicking the *Add to favorites* button when viewing maps in *Monitoring* > *Maps*.



Configuration

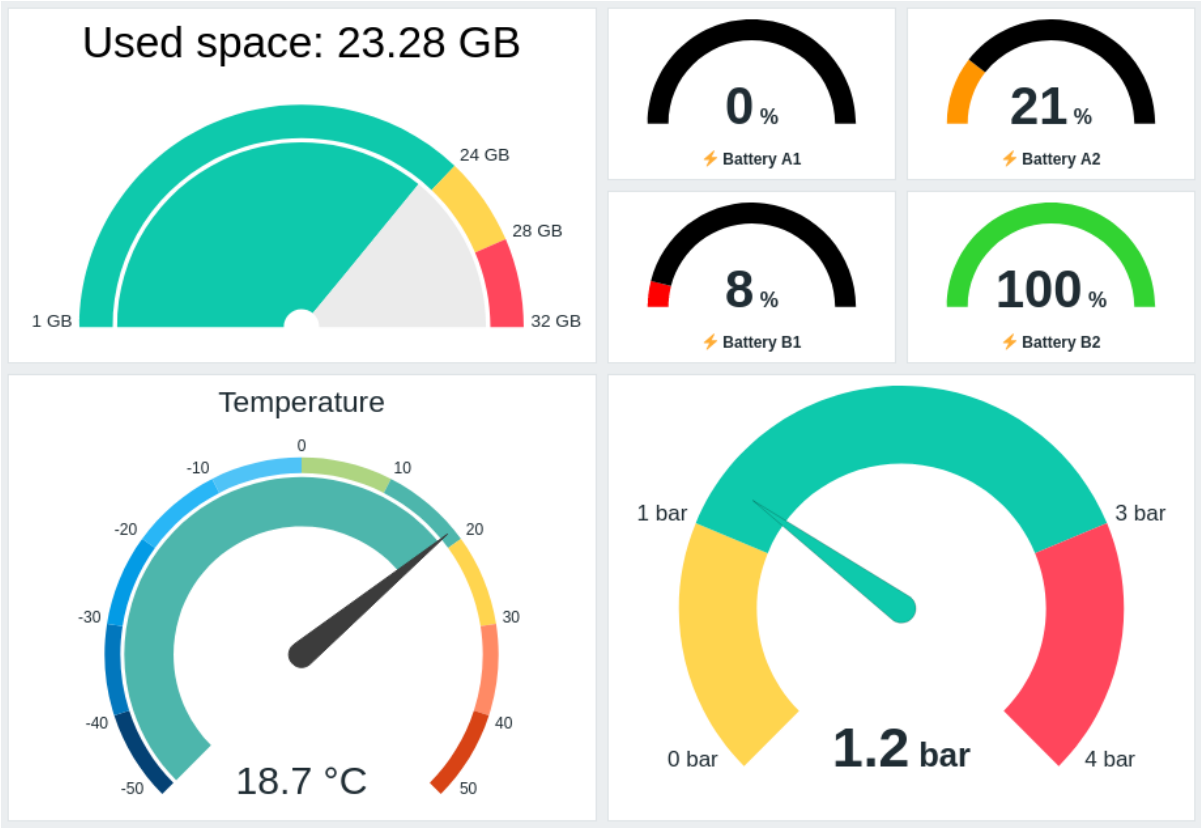
All configuration parameters for this widget are common to all widgets.



Gauge

Overview

The *Gauge* widget displays the value of a single numeric item as a gauge. It is useful for keeping an eye on important metrics, visualizing thresholds, and detecting sudden changes.



You can configure the widget to display the following:

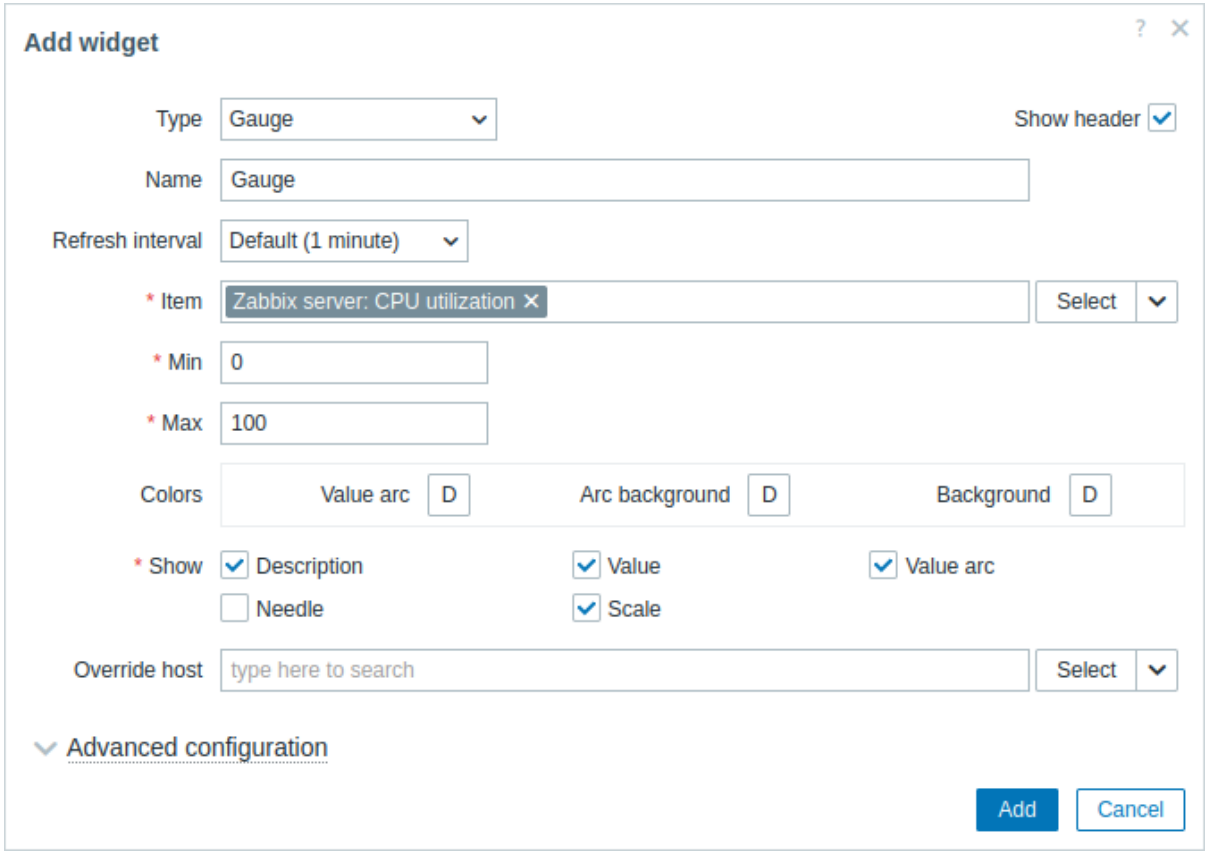
- Item description (Used space: {ITEM.LASTVALUE}, Temperature)
- Item value and units (21 %, 18.7 °C, 1.2 bar)
- Gauge scale (1GB/24GB/28GB/32GB, -50/-40/-30/etc.)
- Gauge arc (gauge value arc and gauge thresholds arc)
- Gauge needle

Clicking the widget opens a simple graph for the item.

The information displayed in the *Gauge* widget can be downloaded as a PNG image by selecting the *Download image* option in the widget *Actions* menu ... that appears when you hover over the widget. Depending on your browser settings, the image is either automatically saved to the `Downloads` folder, or a *Save* dialog opens, allowing you to choose the location.

Configuration

To configure, select *Gauge* as type:

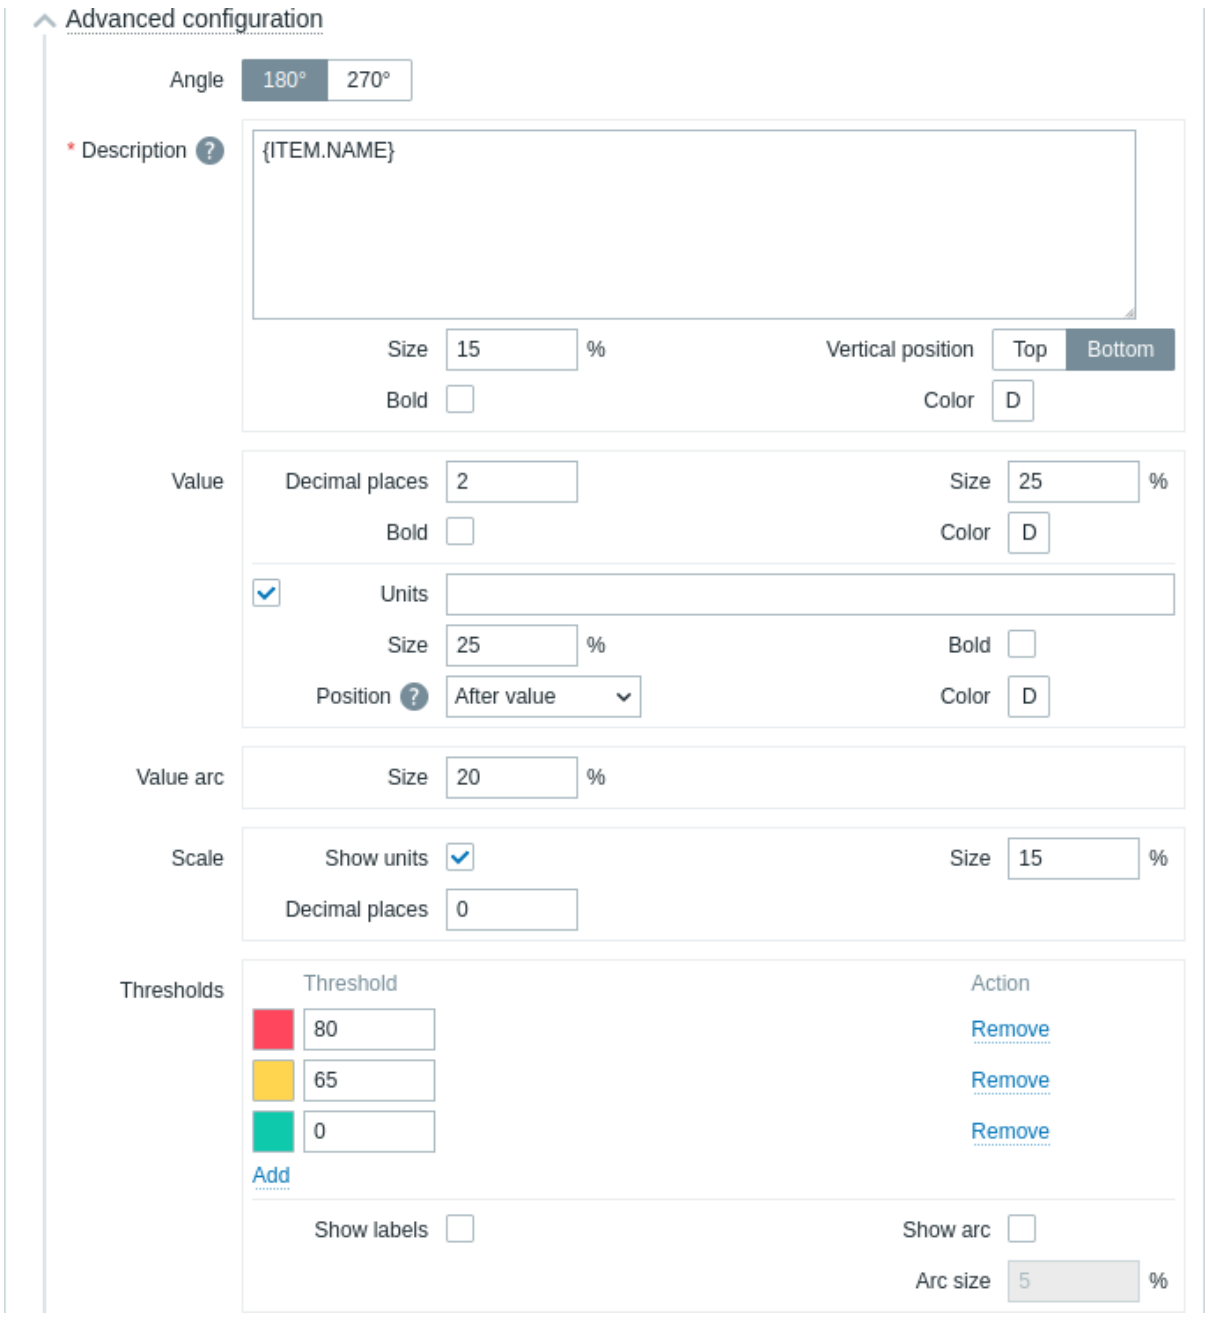


In addition to the parameters that are common for all widgets, you may set the following specific options:

Item	Select the item. Alternatively, select a compatible widget as the data source for items. This field is auto-complete, so starting to type the name of an item will offer a dropdown of matching items. Note that you can select only items that return numeric data (except binary data).
Min	Enter the minimum value of the gauge. Suffixes (for example, ”1d”, ”2w”, ”4K”, ”8G”) are supported. Value mappings are supported.
Max	Enter the maximum value of the gauge. Suffixes (for example, ”1d”, ”2w”, ”4K”, ”8G”) are supported. Value mappings are supported.
Colors	Select the color from the color picker: **Value arc** - select the gauge value arc color; **Arc background** - select the gauge value arc and gauge thresholds arc background color; **Background** - select the widget background color. ”D” stands for the default color, which depends on the frontend theme. If *Thresholds* are set, the default color for **Value arc** depends on the threshold color. To return to the default color, click the *Use default* button in the color picker.
Show	Mark the checkbox to display the respective gauge element - description, value, value arc, needle, scale (the minimum and maximum value of the gauge at the beginning and end of the gauge arc). Unmark to hide. At least one element must be selected. Note that the gauge needle and scale can be displayed if the gauge value arc or gauge thresholds arc (see advanced configuration options) is displayed. Also note that if the gauge needle is displayed, the value is placed under the needle; if the needle is hidden, the value is aligned with the bottom of the gauge arc.
Override host	Select a compatible widget or the dashboard host selector as the data source for hosts. This parameter is not available when configuring the widget on a template dashboard.
Advanced configuration	Click the *Advanced configuration* header to display advanced configuration options. This is also where you’ll be able to adjust the gauge elements selected in the *Show* field.

Advanced configuration

Advanced configuration options are available in the collapsible *Advanced configuration* section:



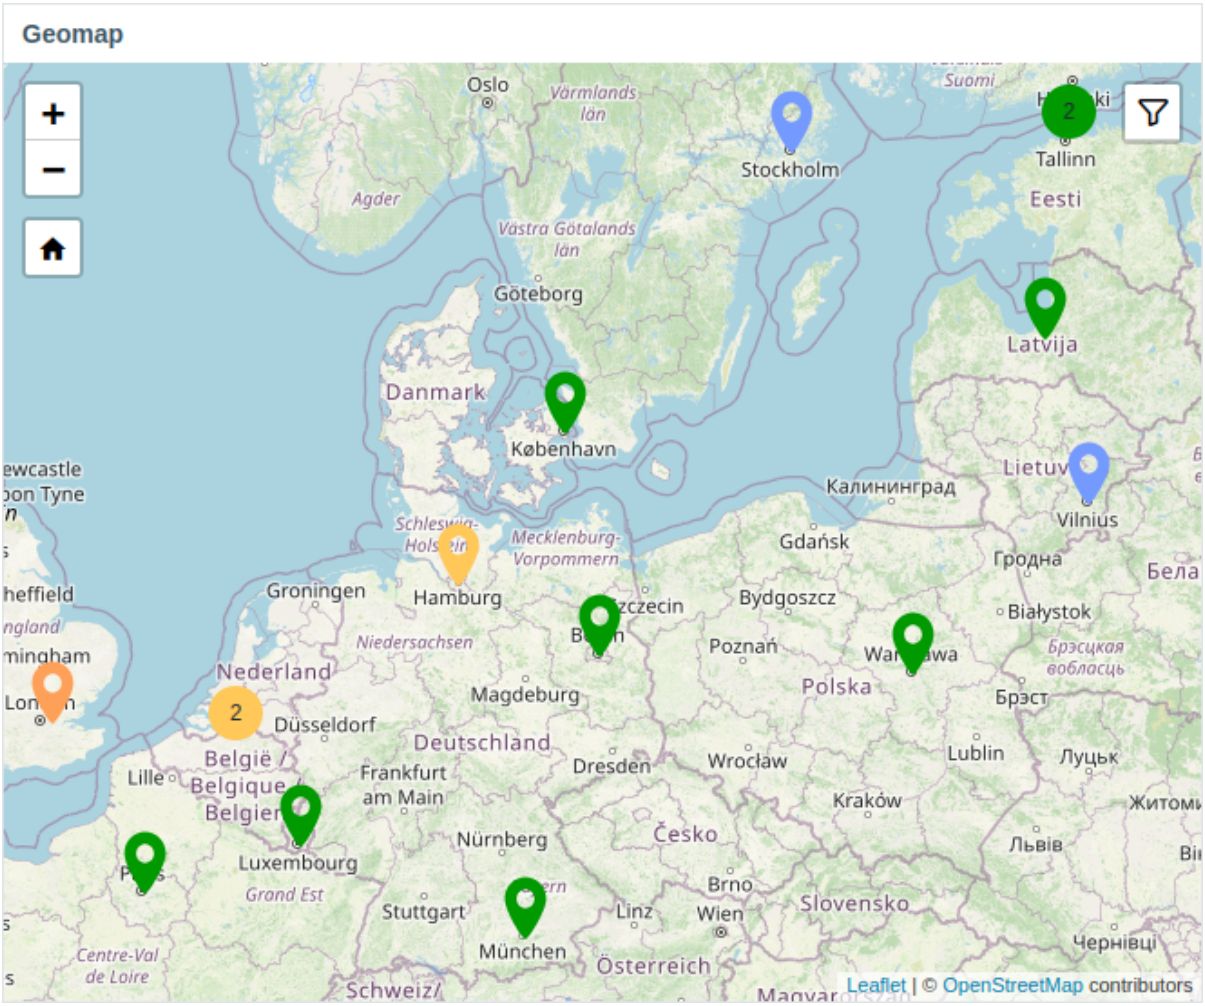
Angle	Select the gauge angle (180° or 270°).
Description	
Description	Enter the item description. This description may override the default item name. Multiline descriptions are supported. A combination of text and supported macros is possible. {HOST.*}, {ITEM.*}, {INVENTORY.*} and user macros are supported.
Size	Enter the font size height for the item description (in percent, relative to the total widget height).
Vertical position	Select the vertical position of the item description (top or bottom, relative to the gauge arc).
Bold	Mark the checkbox to display the item description in bold.
Color	Select the item description color from the color picker. "D" stands for the default color, which depends on the frontend theme. To return to the default color, click the *Use default* button in the color picker.
Value	
Decimal places	Enter the number of decimal places to display with the value. This option affects only items that return numeric (float) data.
Size	Enter the font size height for the value (in percent, relative to the gauge arc height).
Bold	Mark the checkbox to display the value in bold.
Color	Select the value color from the color picker. "D" stands for the default color, which depends on the frontend theme. To return to the default color, click the *Use default* button in the color picker.
Units	

Units	Mark the checkbox to display units with the item value. If you enter a unit name, it will override the units set in the item configuration.
Size	Enter the font size height for the item units (in percent, relative to the gauge arc height).
Bold	Mark the checkbox to display item units in bold.
Position	Select the position of the item units (above, below, before or after, relative to the item value). This option is ignored for the following time-related units: unixtime, uptime, s.
Color	Select the item units color from the color picker. "D" stands for the default color, which depends on the frontend theme. To return to the default color, click the *Use default* button in the color picker.
Value arc	
Arc size	Enter the size height of the gauge value arc (in percent, relative to the gauge arc radius).
Needle	
Color	Select the gauge needle color from the color picker. "D" stands for the default color, which depends on the frontend theme. If *Thresholds* are set, the default color for the needle depends on the threshold color. To return to the default color, click the *Use default* button in the color picker.
Scale	
Show units	Mark the checkbox to display units with the minimum and maximum value of the gauge.
Size	Enter the font size height for the minimum and maximum value of the gauge (in percent, relative to the gauge arc height).
Decimal places	Enter the number of decimal places to display with the minimum and maximum value of the gauge. This option affects only items that return numeric (float) data.
Thresholds	
Thresholds	Click *Add* to add a threshold, select a threshold color from the color picker, and specify a numeric value. The thresholds list will be sorted in ascending order when saved. Note that the colors configured as thresholds will be displayed correctly only for numeric items. Suffixes (for example, "1d", "2w", "4K", "8G") are supported. Value mappings are supported.
Show labels	Mark the checkbox to display threshold values as labels on the gauge scale.
Show arc	Mark the checkbox to display the gauge thresholds arc.
Arc size	Enter the size height of the gauge thresholds arc (in percent, relative to the gauge arc radius).

Geomap

Overview

The *Geomap* widget displays hosts as markers on an interactive Leaflet-based map, with marker colors indicating each host's most severe problem.

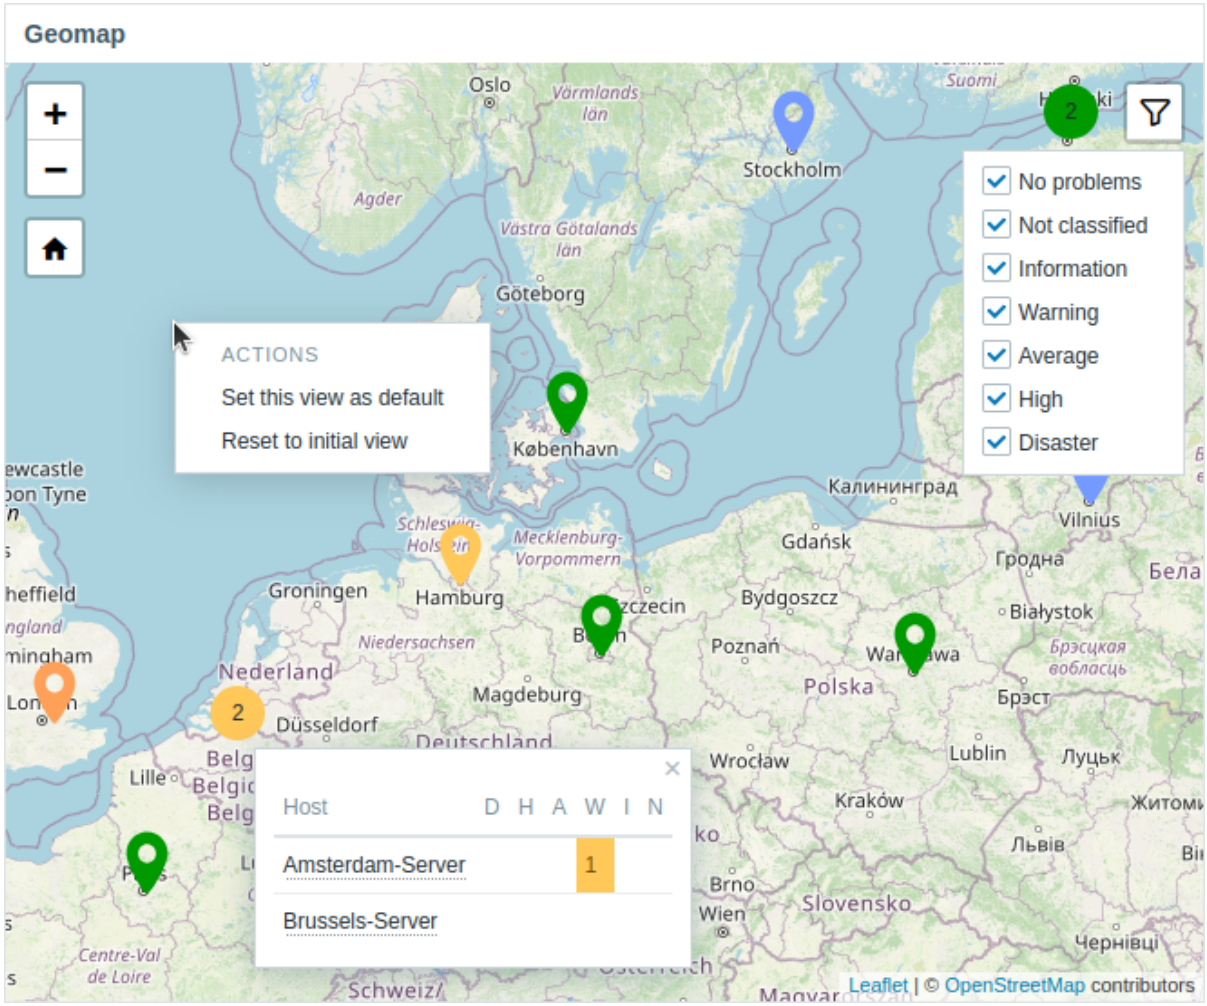


Note:
Zabbix offers multiple predefined map tile providers and lets you add custom ones in *Administration* > *General* > *Geographical maps*.

By default, the widget displays all enabled hosts that have valid coordinates (latitude -90 to 90, longitude -180 to 180) in their inventory.

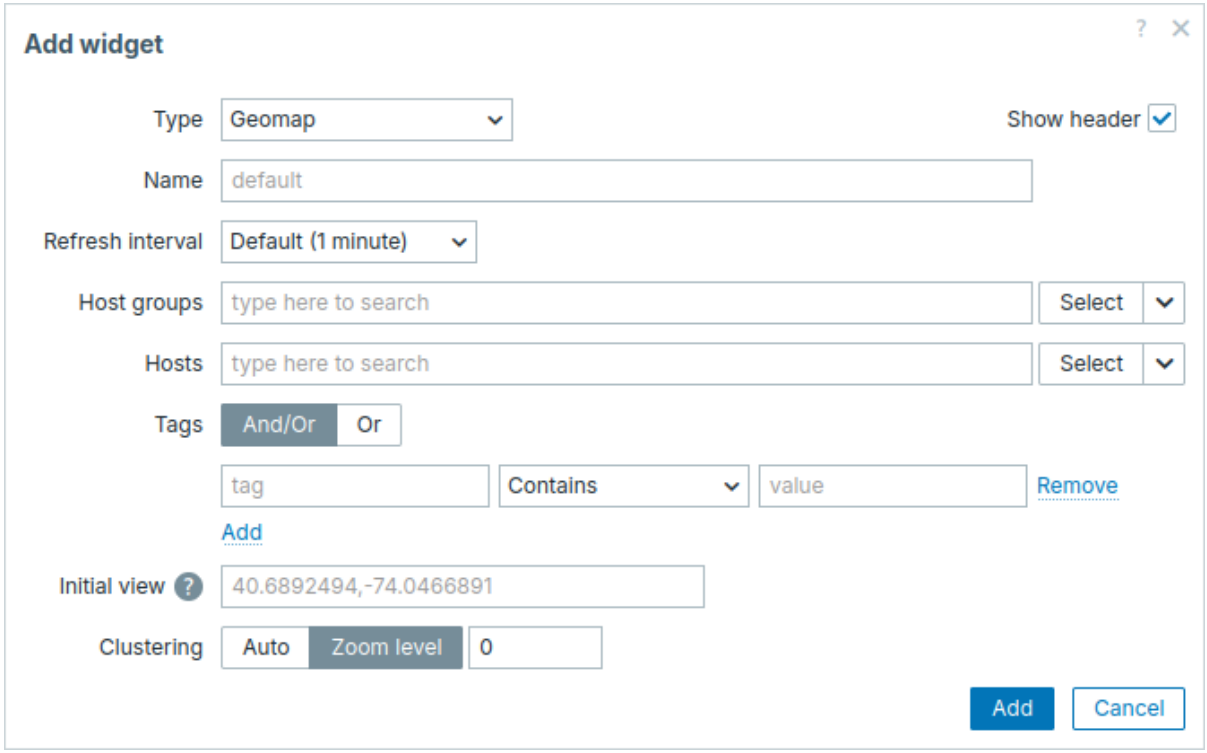
Use the following options to interact with the widget:

- Clicking a marker shows the host name and unresolved problems; clicking the host name opens the host menu.
- Filter hosts by problem severity using the filter icon in the upper-right corner.
- Right-click the map to set or reset the default view; use the home icon to return to the default view.



Configuration

To add the widget, select *Geomap* as type.



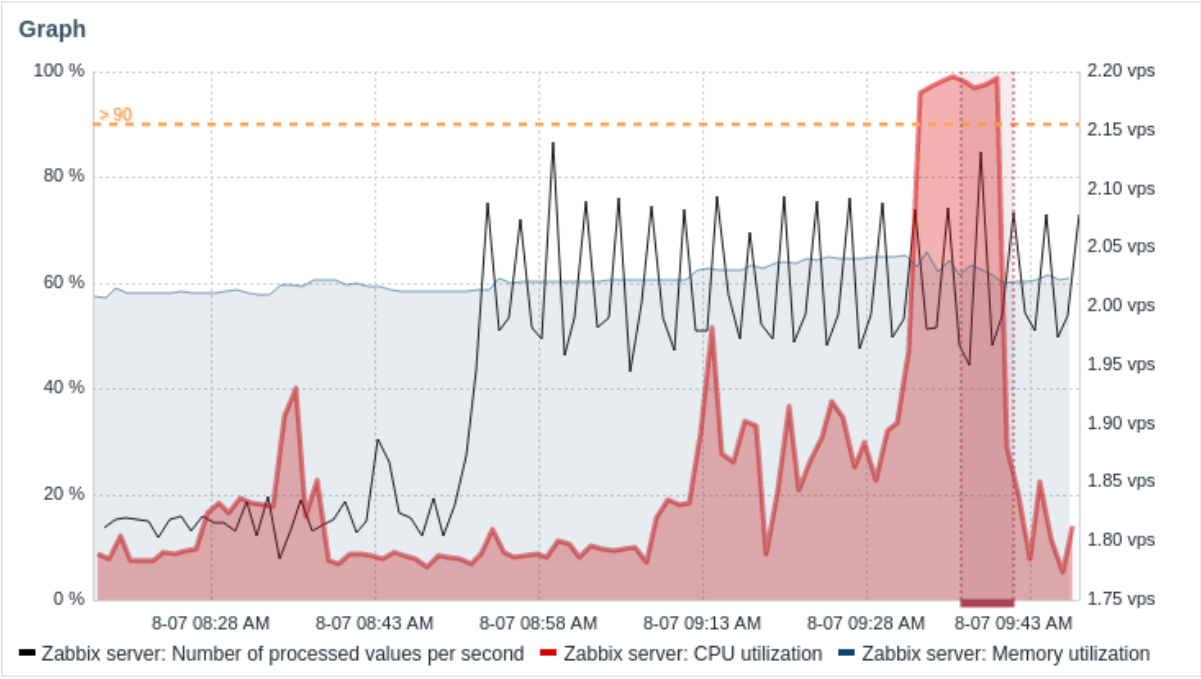
In addition to the parameters that are common for all widgets, you may set the following specific options:

Host groups	Select host groups to be displayed on the map. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. If nothing is selected in both *Host groups* and *Hosts* fields, all hosts with valid coordinates will be displayed. This parameter is not available when configuring the widget on a template dashboard.
Hosts	Select hosts to be displayed on the map. Alternatively, select a compatible widget or the dashboard as the data source for hosts. This field is auto-complete, so starting to type the name of a host will offer a dropdown of matching hosts. If nothing is selected in both *Host groups* and *Hosts* fields, all hosts with valid coordinates will be displayed. This parameter is not available when configuring the widget on a template dashboard.
Tags	Specify tags to limit the number of hosts displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the *Or* condition; **Or** - enough if one condition is met. This parameter is not available when configuring the widget on a template dashboard.
Initial view	Comma-separated center coordinates and an optional zoom level to display when the widget is initially loaded in the format `<latitude>,<longitude>,<zoom>` If initial zoom is specified, the Geomap widget is loaded at the given zoom level. Otherwise, initial zoom is calculated as half of the max zoom for the particular tile provider. The initial view is ignored if the default view is set (see below). Examples: 40.6892494,-74.0466891 40.6892494,-122.0466891
Clustering	Specify how nearby host markers are combined into a single marker with a count: **Auto** - markers are combined automatically; **Zoom level** - markers are combined based on the specified zoom level (0 - maximum zoom-out, 30 - maximum zoom-in). For example, if the parameter is set to 0, markers are not combined at any zoom level; if set to 15, markers are combined starting from maximum zoom-out down to zoom level 15; if set to 30, markers are combined at all zoom levels. Note that the maximum zoom-in level is determined by the map tile provider.

Graph

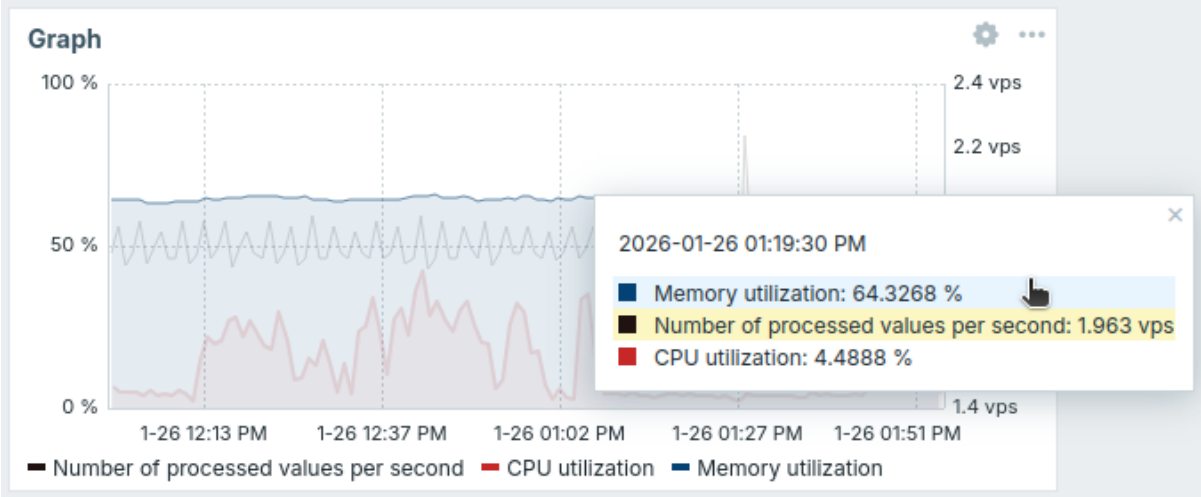
Overview

The *Graph* widget displays numeric item data as a vector-based graph. It can help you track trends, spot problems, and compare values over time and across hosts.

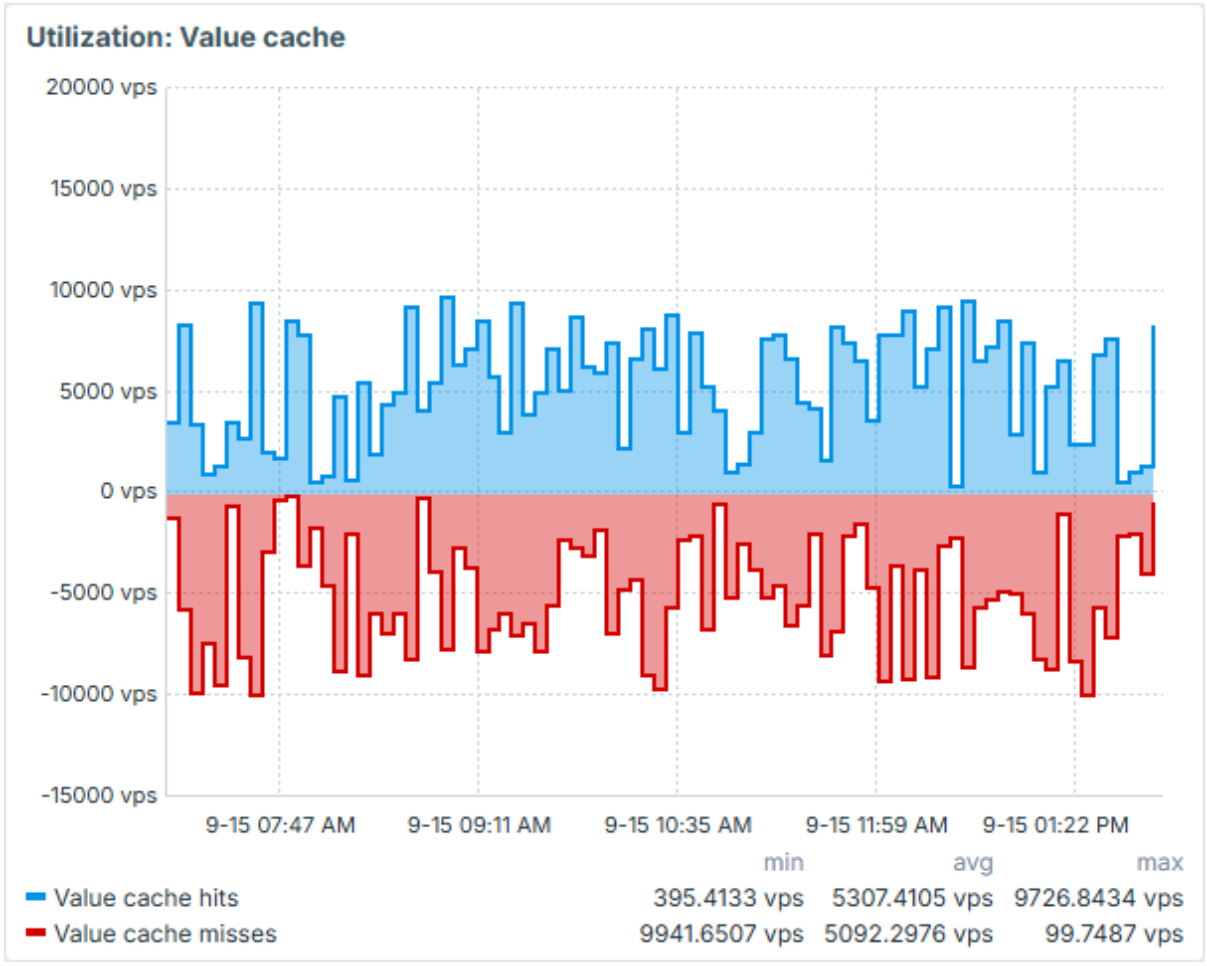


Hovering over the graph displays a tooltip with a list of items sorted by value in descending order.

Hovering over an item in the tooltip highlights its graph (others are dimmed); selecting the item broadcasts its data to other widgets.



The graph can be displayed with inverted Y-axis values; this only affects the visual representation and does not alter the original values.

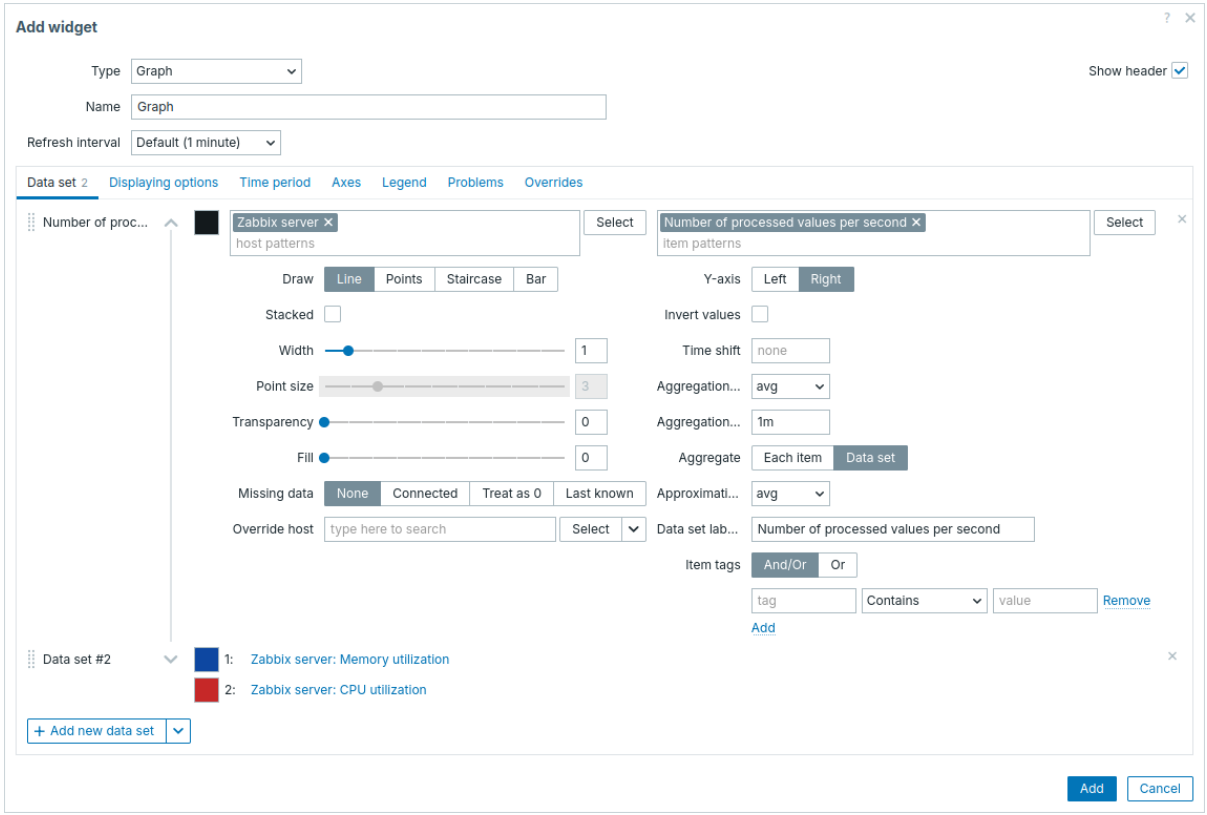


The legacy Graph (classic) widget is also available.

The information displayed in the *Graph* widget can be downloaded as a PNG image by selecting the *Download image* option in the widget *Actions* menu ⋯ that appears when you hover over the widget. Depending on your browser settings, the image is either automatically saved to the `Downloads` folder, or a *Save* dialog opens, allowing you to choose the location.

Configuration

To configure, select *Graph* as type:



Data set

The **Data set** tab allows selecting data for the graph by adding data sets. Two types of data sets can be added:

- *Item patterns* - data from matching items is displayed. You can pick a single base color or select a palette row to assign distinct colors to each matched item.
- *Item list* - data from selected items is displayed. You can choose each item's color individually from the picker.

By default, an *Item patterns* data set is added.

Data set	For **Item patterns** data set: Select or enter host and item patterns; data of items that match these patterns will be displayed on the graph; up to 50 items may be displayed. Wildcard patterns may be used for selection (for example, ∗ will return results that match zero or more characters). To specify a wildcard pattern, enter the string manually and press *Enter*. The wildcard symbol is always interpreted, so it is not possible to add, for example, an item named *item** individually if there are other matching items (for example, *item2*, *item3*). Specifying host and item patterns is mandatory for "Item patterns" data sets. See also: Data set configuration details. For **Item list** data set: Select items for the graph by clicking the *Add item* button. You may also select compatible widgets as the data source for items by clicking the *Add widget* button. Specifying items or widgets is mandatory for "Item list" data sets. See also: Data set configuration details. Note that only numeric item types are allowed. When configuring the widget on a template dashboard, the parameter for specifying host patterns is not available, and the parameter for specifying an item list allows only the selection of items configured on the template.
Draw	Choose the draw type of the metric. Possible draw types: *Line* (set by default), *Points*, *Staircase*, and *Bar*. Note that if there is only one data point in the line/staircase graph, it is drawn as a point regardless of the draw type. The point size is calculated from the line width, but it cannot be smaller than 3 pixels, even if the line width is less.

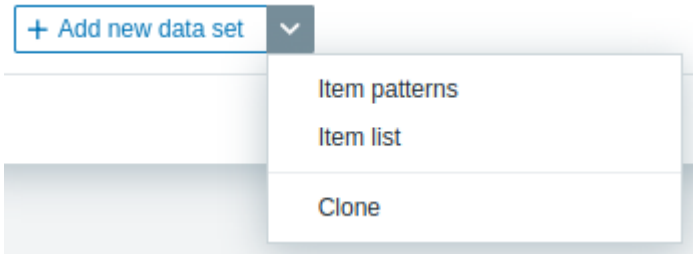
Stacked	Mark the checkbox to display data as stacked (filled areas displayed). This option is disabled when *Points* draw type is selected.
Width	Set the line width. This option is available when *Line* or *Staircase* draw type is selected.
Point size	Set the point size. This option is available when *Points* draw type is selected.
Transparency	Set the transparency level.
Fill	Set the fill level. This option is available when *Line* or *Staircase* draw type is selected.
Missing data	Select the option for displaying missing data: **None** - the gap is left empty; **Connected** - two border values are connected; **Treat as 0** - the missing data is displayed as 0 values; **Last known** - the missing data is displayed with the same value as the last known value; not applicable for the *Points* and *Bar* draw type.
Override host	Select a compatible widget or the dashboard host selector as the data source for hosts. This parameter is not available when configuring the widget on a template dashboard.
Y-axis	Select the side of the graph where the Y-axis will be displayed.
Invert values	Mark the checkbox to multiply all Y-axis values by -1 before displaying them on the graph. This affects only visual representation; the original data remains unchanged. This option does not affect the tooltip, legend, or percentile line. If the *Simple triggers* option is selected for the axis, the trigger thresholds will be displayed inverted.
Time shift	Specify time shift if required. You may use time suffixes in this field. Negative values are allowed.
Aggregation function	Specify which aggregation function to use for each item or for an entire data set within the selected *Aggregation interval*: **min** - display the smallest value; **max** - display the largest value; **avg** - display the average value; **sum** - display the sum of values; **count** - display the count of values; **first** - display the first value; **last** - display the last value. If **not used** is selected, all values are displayed (no aggregation). Only numeric data can be displayed for *min*, *max*, *avg* and *sum*. For *count*, non-numeric data will be changed to numeric. See also: Aggregation in graphs.
Aggregation interval	Specify the interval for aggregating values. You may use time suffixes in this field. A numeric value without a suffix will be regarded as seconds. Note that if the widget is configured to display historical data based on trends (*History data selection* is set to *Trends* or *Auto*), it is recommended to use an aggregation interval that is a multiple of 1 hour (e.g., 3600, 60m, 1h, 3h, etc.). Trends store hourly aggregated values, so using an aggregation interval that is not a multiple of 1 hour (e.g., 100s, 7min, 15min, 90min, etc.) may lead to results that are hard to interpret.
Aggregate	Specify whether to aggregate: **Each item** - each item in the dataset will be aggregated and displayed separately; **Data set** - all dataset items will be aggregated and displayed as one value.
Approximation	Specify what value to display when more than one value exists per vertical graph pixel: **all** - display the smallest, the largest and the average values; **min** - display the smallest value; **max** - display the largest value; **avg** - display the average value. This setting is useful when displaying a graph for a large time period with frequent update interval (such as one year of values collected every 10 minutes).

Data set label	Specify the data set label that is displayed in graph *Data set* configuration and in graph *Legend* (for aggregated data sets). All data sets are numbered including those with a specified *Data set label*. If no label is specified, the data set will be labeled automatically according to its number (e.g., ”Data set #2”, ”Data set #3”, etc.). Data set numbering is recalculated after reordering/dragging data sets. Data set labels that are too long will be shortened to fit where displayed (e.g., ”Number of proc...”).
Item tags	Specify tags to filter the items displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met.

Data set configuration details

Existing data sets are displayed in a list. You may:

- Click the move icon and drag a data set to a new place in the list.
- Click the expand icon to expand data set details. When expanded, this icon turns into a collapse icon.
- Click the color icon to open the picker. You can enter a hex code, choose a solid swatch, or switch to the *Palette* tab and select a row of predefined colors. The chosen color applies directly for *Item list* data sets or as the base for generated shades in *Item patterns*. Use Tab to move between dialog controls, arrow keys to navigate swatches or palette rows, Enter to select, and Esc to cancel.
- Click the *Add new data set* button to add an empty data set that allows you to select host and item patterns. If you click the downward pointing icon next to the *Add new data set* button, a drop-down menu appears, allowing you to add a new *Item patterns* or *Item list* data set or allowing you to *Clone* the currently open data set. If all data sets are collapsed, the *Clone* option is not available.

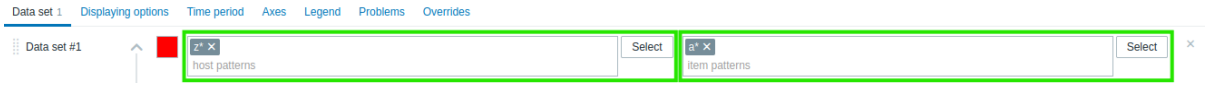


The **Item patterns** data set contains *Host patterns* and *Item patterns* fields that both recognize full names or patterns containing a wildcard symbol (*). This functionality allows you to select all the host names and item names containing the selected pattern. While typing the item name or item pattern in the *Item patterns* field, only items belonging to the selected host name(s) are displayed in the dropdown list.

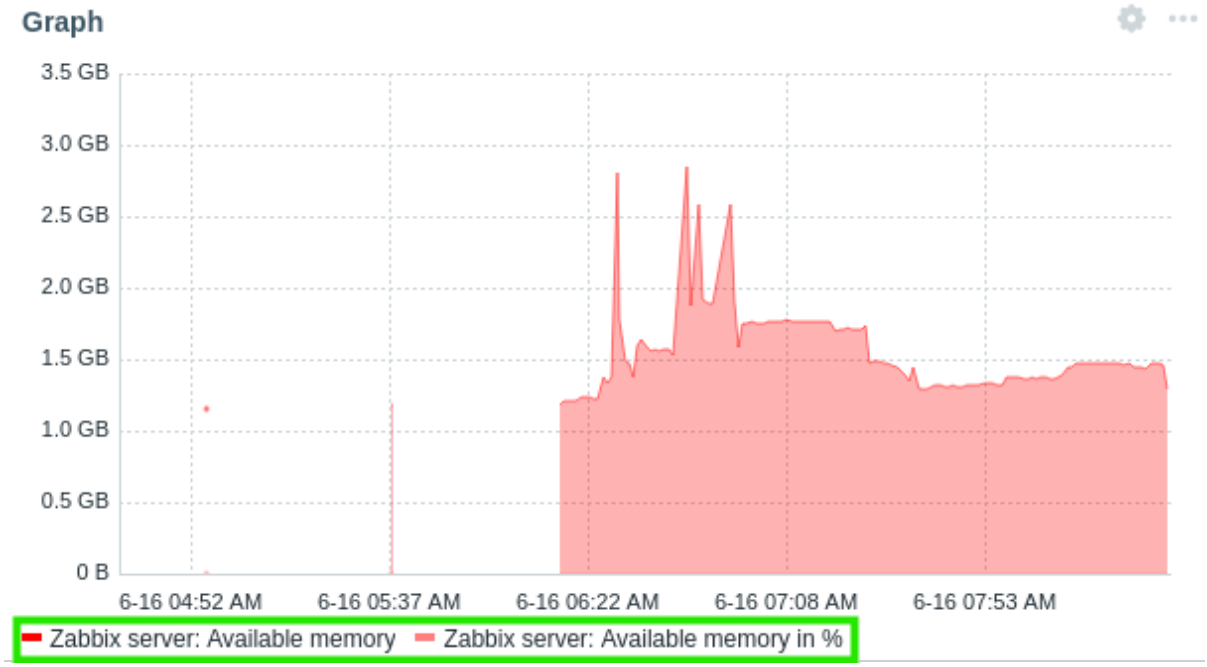
For example, having typed a pattern **z*** in the *Host patterns* field, the dropdown list displays all host names containing this pattern: **z*, Zabbix server, Zabbix proxy**. After pressing *Enter*, this pattern is accepted and is displayed as **z***. Similarly, having typed the pattern **a*** in the *Item patterns* field, the dropdown list displays all item names containing this pattern: **a*, Available memory, Available memory in %**.



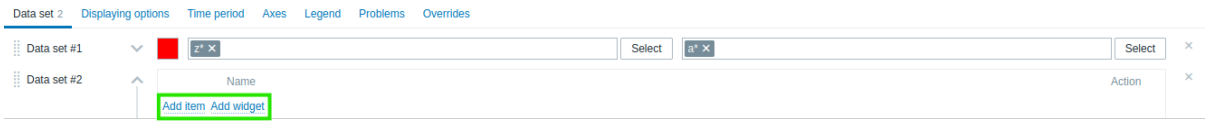
After pressing *Enter*, the pattern is accepted and shown as **a***.



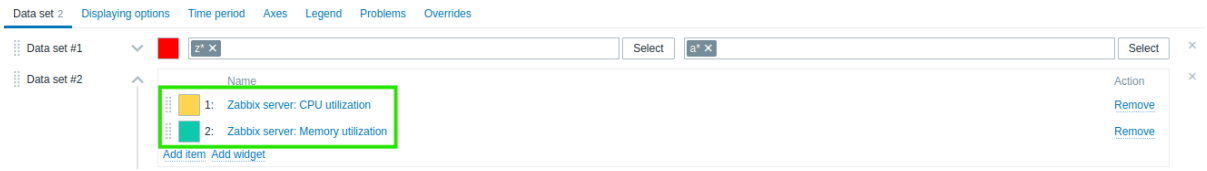
The graph then displays all items belonging to the selected host name(s).



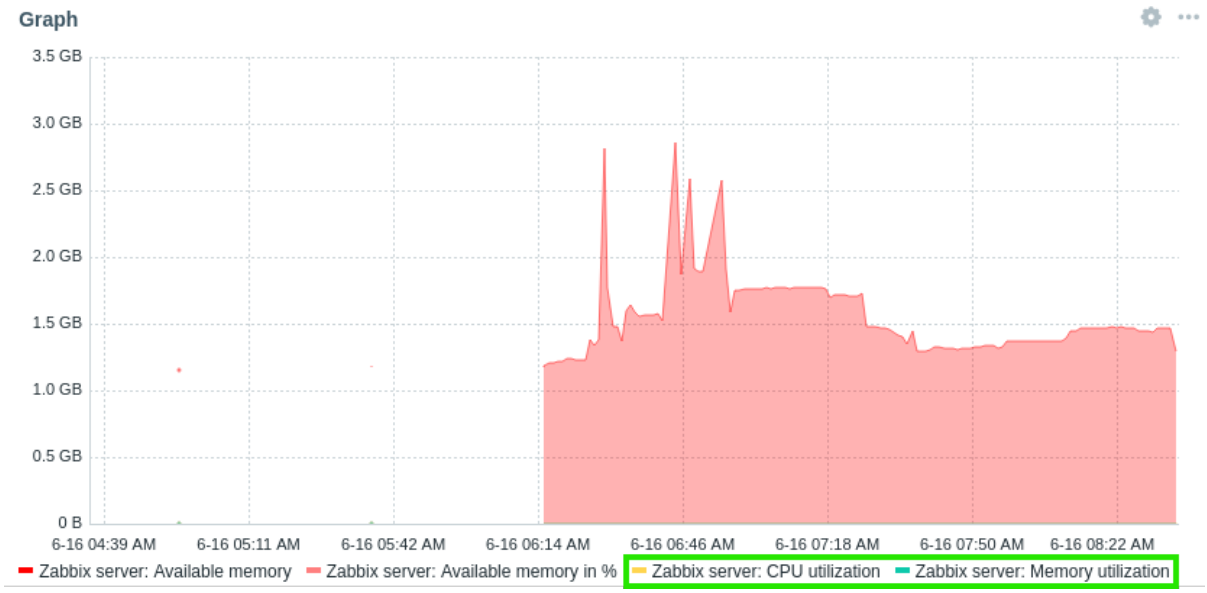
The **Item list** data set contains the *Add item* button that allows you to add items to be displayed on the graph. You can also add compatible widgets as the data source for items by clicking the *Add widget* button.



For example, clicking the *Add item* button opens a pop-up window containing a *Host* parameter. Having selected a host, all its items that are available for selection are displayed in a list.



After selecting one or more items, they will be displayed in the data set item list and in the graph.

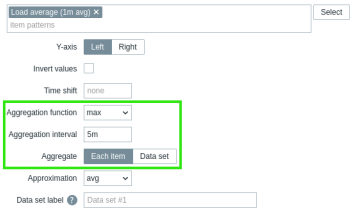


Aggregation in graphs

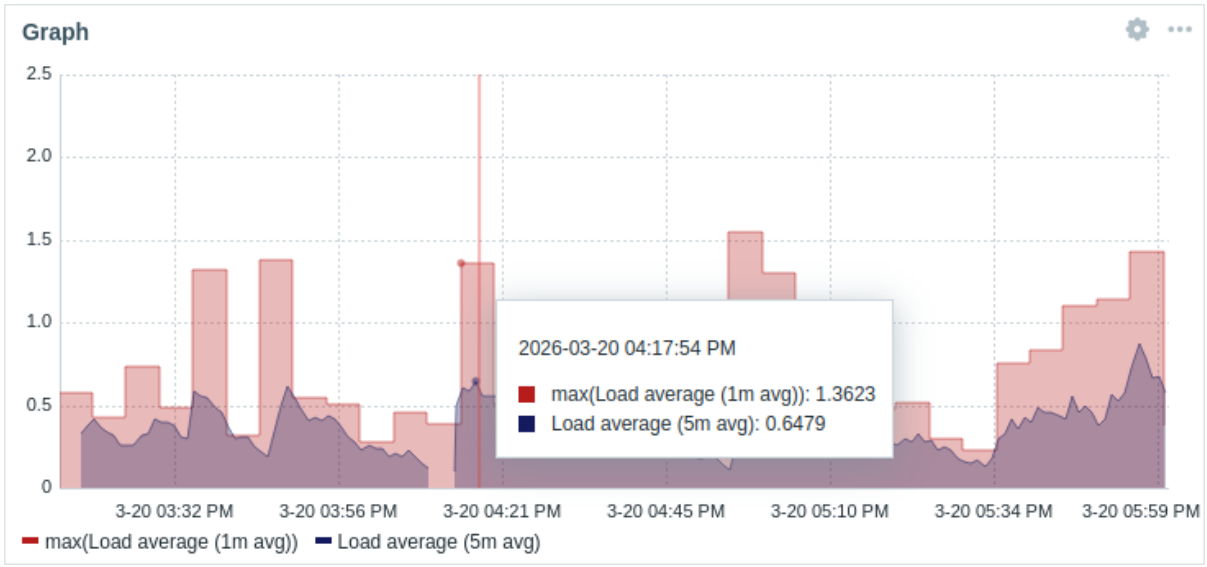
Aggregation allows an aggregated value (minimum, maximum, average, etc.) to be displayed for the specified interval (5 minutes, an hour, a day, etc.) instead of all individual values. Within a data set, aggregation can be applied either to each item separately or to the entire data set.

Some use cases for the aggregation function can be as follows: - The average request count to an NGINX server per day. - The lowest disk space among clusters per week.

In the following example, the ”max” aggregation function is applied to the item with an aggregation interval of five minutes:



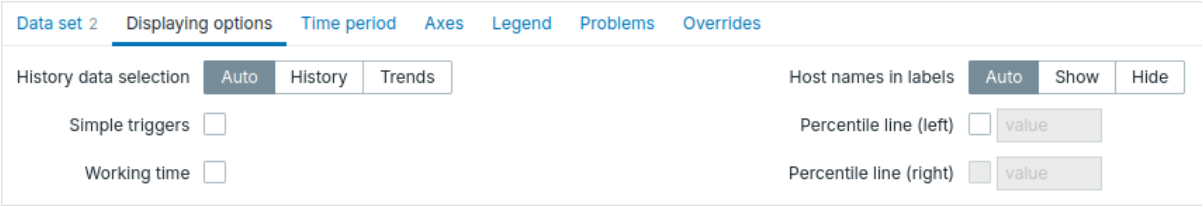
The example demonstrates the maximum CPU load over five minutes (aggregated from one-minute average values), compared to the average CPU load over five minutes.



When *Show aggregation function* is selected for the legend, items are displayed in parentheses, prefixed by the aggregation function in the legend and the tooltip. If *Data set label* is specified for a graph, the label is displayed in parentheses, prefixed by the aggregation function.

Displaying options

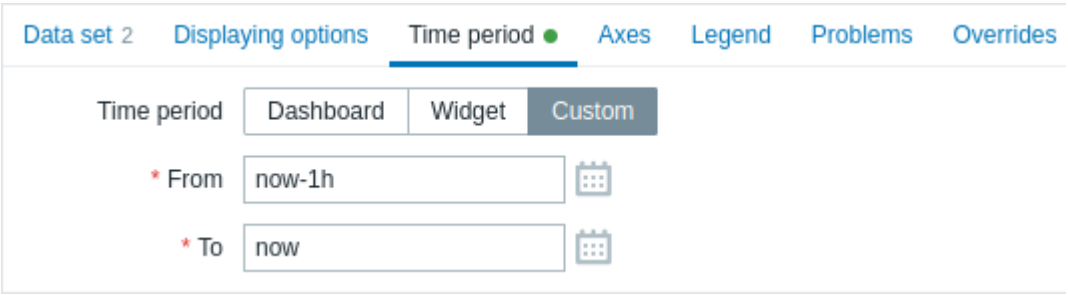
The **Displaying options** tab allows you to define history data selection:



History data selection	Set the source of graph data: **Auto** - data is sourced according to the classic graph algorithm (default); **History** - data from history; **Trends** - data from trends.
Simple triggers	Mark the checkbox to show the trigger thresholds for simple triggers. The thresholds will be drawn as dashed lines using the trigger severity color. A simple trigger is a trigger with one function (only `last`, `max`, `min`, `avg`) for one item in the expression. A maximum of three triggers can be drawn. Note that the trigger has to be within the drawn range to be visible.
Working time	Mark the checkbox to show working time on the graph. Working time (working days) is displayed in graphs as a white background, while non-working time is displayed in gray (with the *Original blue* default frontend theme).
Host names in labels	Choose whether to display host names in the graph legend and in the tooltip that appears when hovering over the graph data: **Auto** - host names are displayed only if more than one host is present across the data sets (default). **Show** - host names are displayed. **Hide** - host names are hidden. This parameter is not available when configuring the widget on a template dashboard.
Percentile line (left)	Mark the checkbox and enter the percentile value to show the specified percentile as a line on the left Y-axis of the graph. If, for example, a 95% percentile is set, then the percentile line will be at the level where 95 percent of the values fall under.
Percentile line (right)	Mark the checkbox and enter the percentile value to show the specified percentile as a line on the right Y-axis of the graph. If, for example, a 95% percentile is set, then the percentile line will be at the level where 95 percent of the values fall under.

Time period

The **Time period** tab allows you to set a time period for which to display data in the graph:

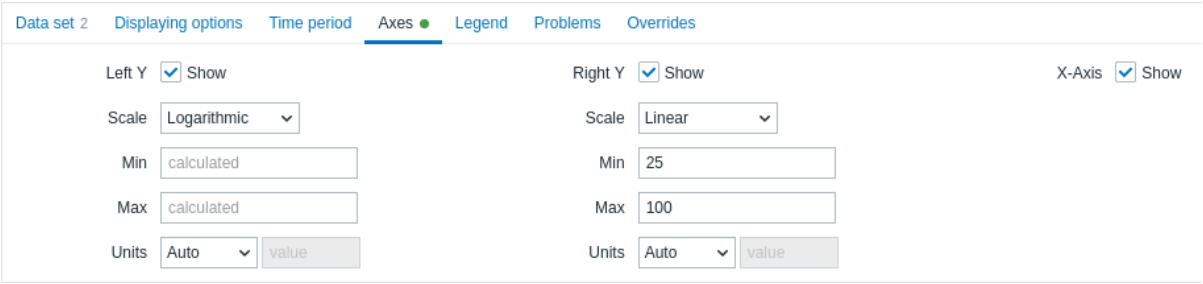


Time period	Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period.
Widget	Enter or select a compatible widget as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".

To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".

Axes

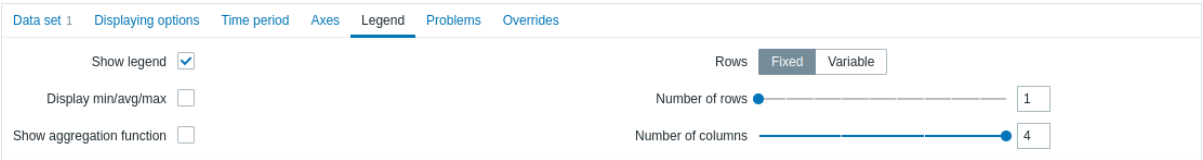
The **Axes** tab allows you to customize how axes are displayed:



Left Y	Mark this checkbox to make left Y-axis visible. The checkbox may be disabled if unselected either in *Data set* or in *Overrides* tab.
Right Y	Mark this checkbox to make right Y-axis visible. The checkbox may be disabled if unselected either in *Data set* or in *Overrides* tab.
X-Axis	Unmark this checkbox to hide X-axis (marked by default).
Scale	Choose the scale for the graph axis values from the dropdown: **Linear** - axis values increase by a fixed amount (e.g., 10, 20, 30), suitable for data that changes steadily or covers a small to moderate range; **Logarithmic** - axis values increase exponentially (e.g., 10, 100, 1000), suitable for data that changes rapidly or covers a large range.
Min	Set the minimum value of the corresponding axis. Visible range minimum value of Y-axis is specified.
Max	Set the maximum value of the corresponding axis. Visible range maximum value of Y-axis is specified.
Units	Choose the unit for the graph axis values from the dropdown: **Auto** - axis values are displayed using the unit of the first item in the data set; **Static** - axis values are displayed using the unit specified in the *value* input field; if the field is left blank, only numeric values are displayed.

Legend

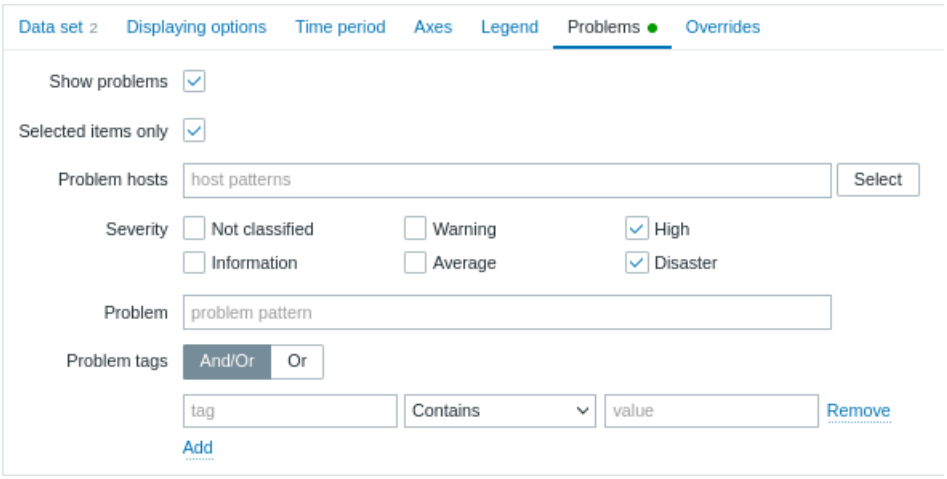
The **Legend** tab allows you to customize the graph legend:



Show legend	Unmark this checkbox to hide the legend on the graph (marked by default).
Display min/avg/max	Mark this checkbox to display the minimum, average, and maximum values of the item in the legend.
Show aggregation function	Mark this checkbox to show the aggregation function in the legend.
Rows	Select the display mode for legend rows: **Fixed** - the number of rows displayed is determined by the *Number of rows* parameter value; **Variable** - the number of rows displayed is determined by the amount of configured items while not exceeding the *Maximum number of rows* parameter value.
Number of rows/ Maximum number of rows	If *Rows* is set to "Fixed", set the number of legend rows to be displayed (1-10). If *Rows* is set to "Variable", set the maximum number of legend rows to be displayed (1-10).
Number of columns	Set the number of legend columns to be displayed (1-4). This parameter is available if *Display min/avg/max* is unmarked.

Problems

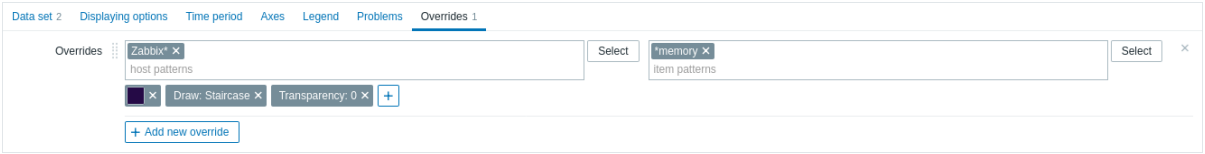
The **Problems** tab allows you to customize the problem display:



Show problems	Mark this checkbox to enable problem displaying on the graph (unmarked, i.e., disabled by default).
Selected items only	Mark this checkbox to include problems for the selected items only to be displayed on the graph.
Problem hosts	Select the problem hosts to be displayed on the graph. Wildcard patterns may be used (for example, * will return results that match zero or more characters). To specify a wildcard pattern, just enter the string manually and press *Enter*. While you are typing, note how all matching hosts are displayed in the dropdown. This parameter is not available when configuring the widget on a template dashboard.
Severity	Mark problem severities to filter problems to be displayed on the graph. If no severities are marked, all problems will be displayed.
Problem	Specify the problem's name to be displayed on the graph.
Problem tags	Specify problem tags to limit the number of problems displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the *Or* condition; **Or** - enough if one condition is met.

Overrides

The **Overrides** tab allows you to add custom overrides for data sets:



Overrides are useful when several items are selected for a data set using the * wildcard and you want to change how the items are displayed by default (e.g., default base color or any other property).

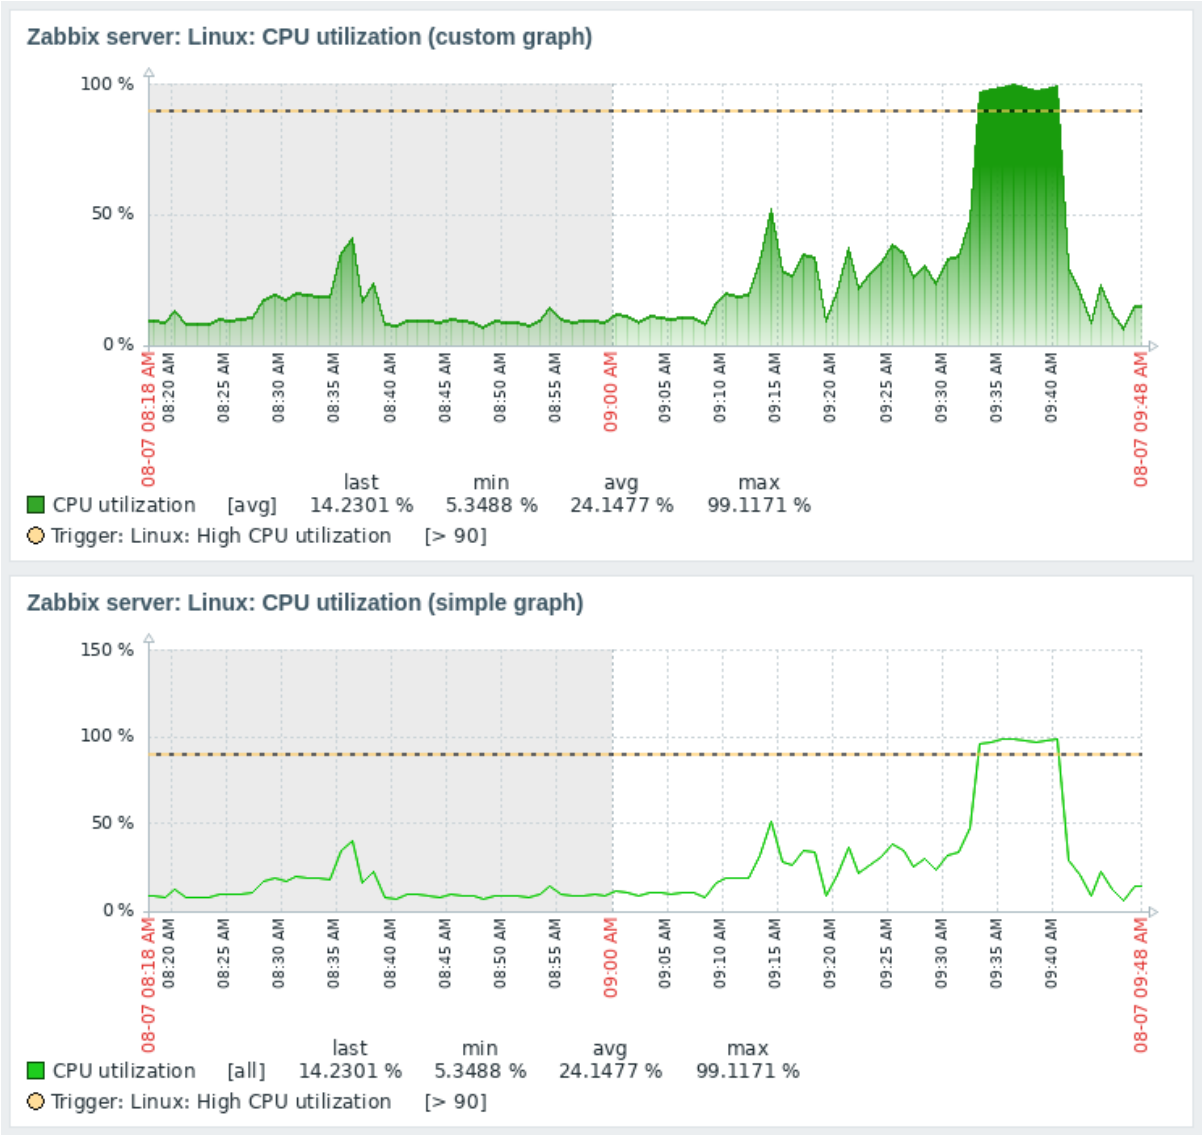
Existing overrides (if any) are displayed in a list. To add a new override:

- Click the + Add new override button
- Select hosts and items for the override. Alternatively, you may enter host and item patterns. Wildcard patterns may be used (for example, * will return results that match zero or more characters). To specify a wildcard pattern, just enter the string manually and press *Enter*. While you are typing, note how all matching hosts are displayed in the dropdown. The wildcard symbol is always interpreted, therefore it is not possible to add, for example, an item named "item*" individually if there are other matching items (e.g., item2, item3). Host pattern and item pattern parameters are mandatory. The parameter for specifying host patterns is not available when configuring the widget on a template dashboard. The parameter for specifying an item list allows only the selection of items configured on the template when configuring the widget on a template dashboard.
- Click +, to select override parameters. At least one override parameter should be selected. For parameter descriptions, see the *Data set* tab above.

Graph (classic)

Overview

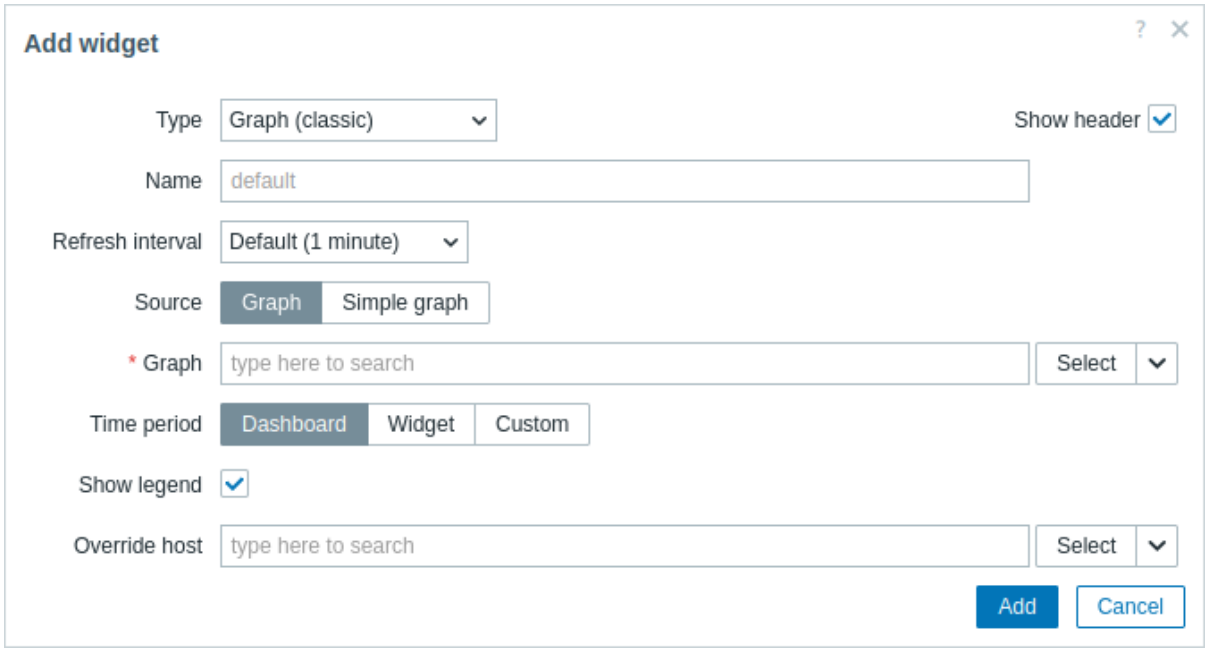
The *Graph (classic)* widget displays numeric item data as an image-based custom graph or simple graph. It supports time period selection, trigger lines, and optional legend display.



The information displayed in the *Graph (classic)* widget can be downloaded as a PNG image by selecting the *Download image* option in the widget *Actions* menu ... that appears when you hover over the widget. Depending on your browser settings, the image is either automatically saved to the `Downloads` folder, or a *Save* dialog opens, allowing you to choose the location.

Configuration

To configure, select *Graph (classic)* as type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

Source	Select the graph type: **Graph** - custom graph; **Simple graph** - simple graph.
Graph	Select the custom graph to display. Alternatively, select a compatible widget as the data source for graphs. This parameter is available if *Source* is set to "Graph".
Item	Select the item to display in a simple graph. Alternatively, select a compatible widget as the data source for items. This parameter is available if *Source* is set to "Simple graph".
Time period	Set a time period for which to display data in the graph. Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period.
Widget	Enter or select a compatible widget as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
Show legend	Unmark this checkbox to hide the legend on the graph (marked by default).
Override host	Select a compatible widget or the dashboard host selector as the data source for hosts. This parameter is not available when configuring the widget on a template dashboard.

The legend of the graph consists of three sections:

- Items and their aggregated values
- Percentiles (if configured)
- Triggers (if any are associated with the displayed items)

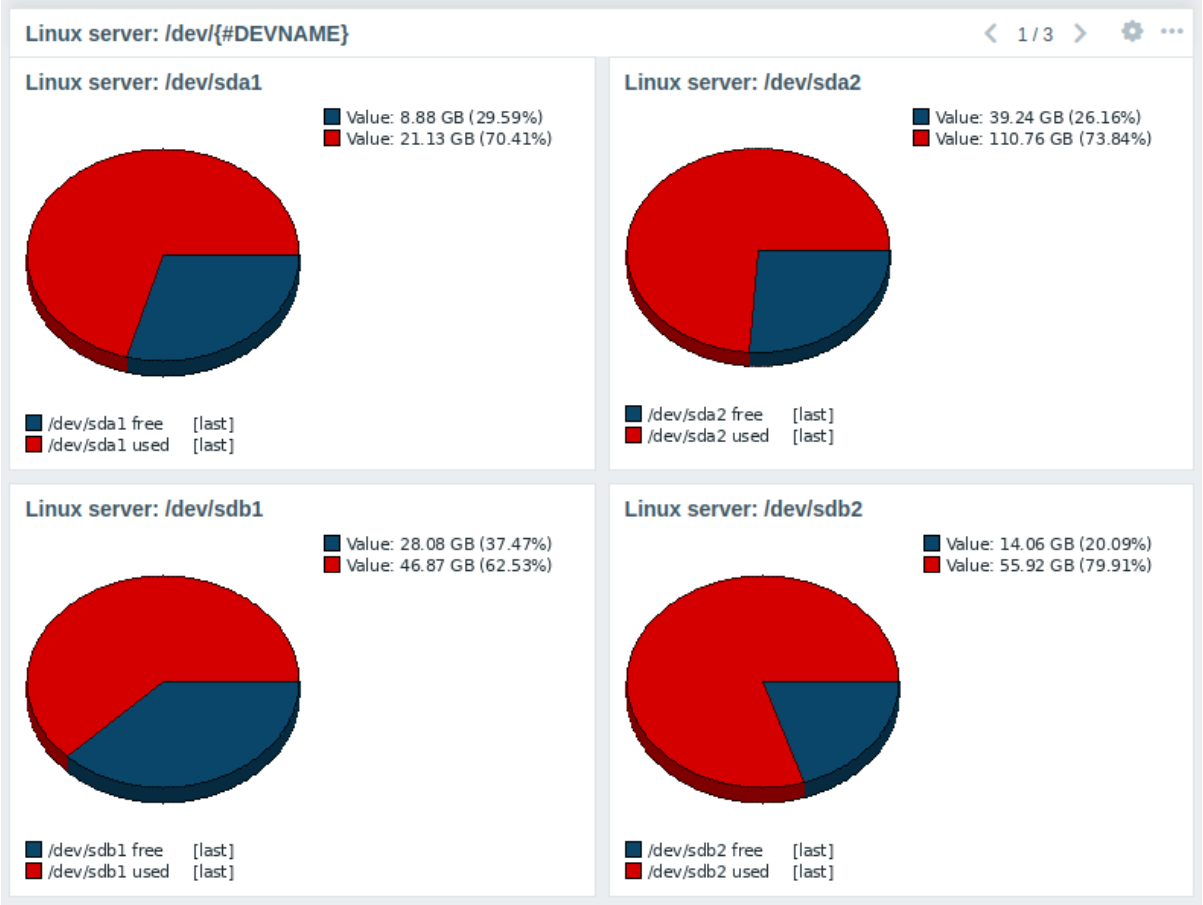
If the height of the graph within the widget is insufficient, the legend may not be displayed or may be displayed only partially. Triggers and percentiles are hidden first, followed by the item legend. To display the full legend, increase the widget's vertical size.

Note:
No more than 3 trigger lines can be displayed. If there are more triggers then the triggers with lower severity are prioritized for display.

Graph prototype

Overview

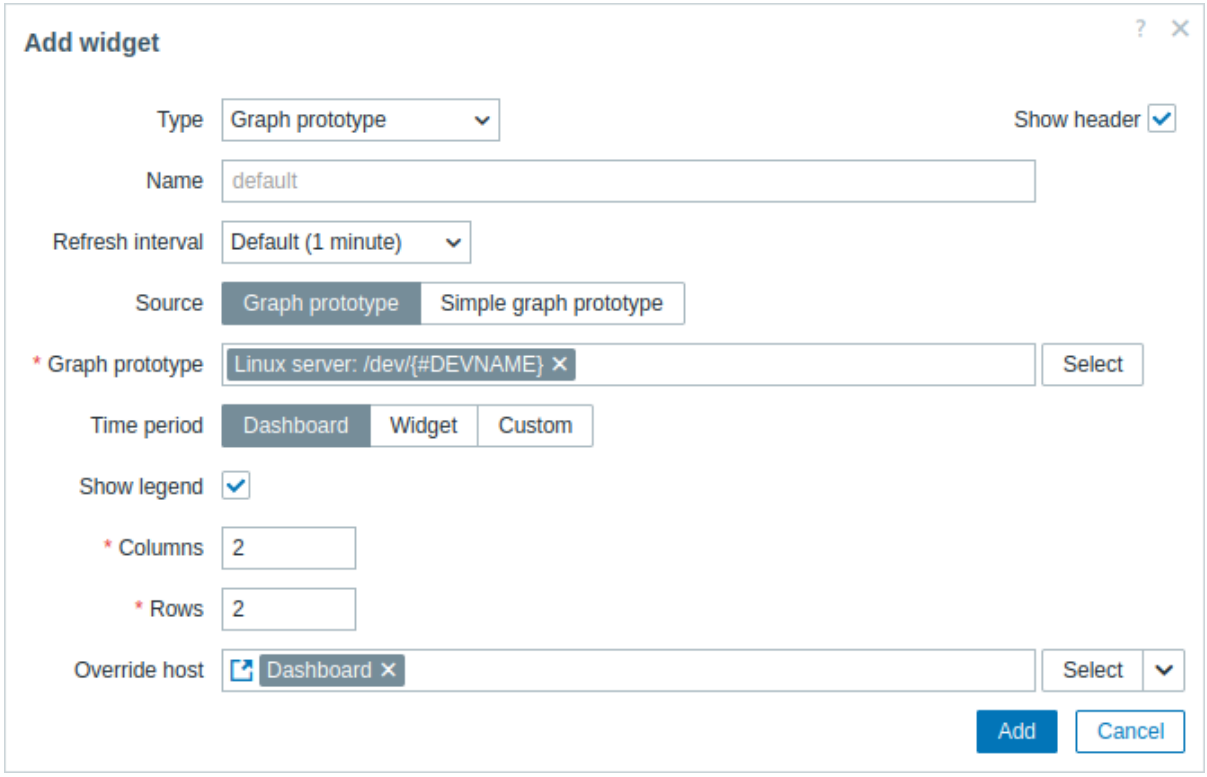
The *Graph prototype* widget displays a grid of automatically created custom graphs based on low-level discovery rule graph prototypes or item prototypes.



Graphs are arranged in a configurable grid layout; if the number of graphs exceeds the layout, paging is enabled in the widget.

Configuration

To configure, select *Graph prototype* as widget type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

Source	Select the source of the graphs: **Graph prototype** or **Simple graph prototype**.
Graph prototype	Select a graph prototype to display graphs discovered by the graph prototype. This parameter is available if *Source* is set to "Graph prototype".
Item prototype	Select an item prototype to display simple graphs for items discovered by the item prototype. This parameter is available if *Source* is set to "Simple graph prototype".
Time period	Set a time period for which to display data in the graphs. Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period.
Widget	Enter or select a compatible widget as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
Show legend	Unmark this checkbox to hide the legend on the graphs (marked by default).
Override host	Select a compatible widget or the dashboard host selector as the data source for hosts. This parameter is not available when configuring the widget on a template dashboard.
Columns	Enter the number of columns of graphs to display within a graph prototype widget.
Rows	Enter the number of rows of graphs to display within a graph prototype widget.

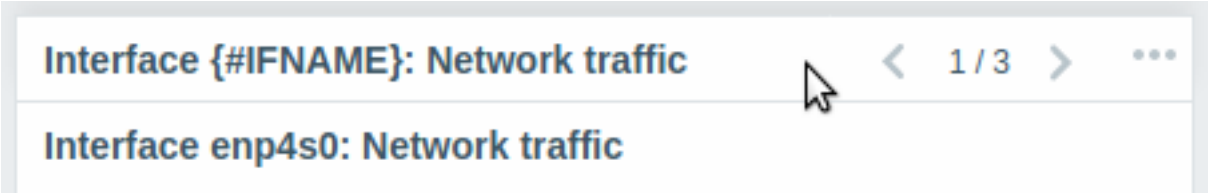
The legend of the graph consists of three sections:

- Items and their aggregated values
- Percentiles (if configured)
- Triggers (if any are associated with the displayed items)

If the height of the graph within the widget is insufficient, the legend may not be displayed or may be displayed only partially. Triggers and percentiles are hidden first, followed by the item legend. To display the full legend, increase the widget's vertical size.

Note:
No more than 3 trigger lines can be displayed. If there are more triggers then the triggers with lower severity are prioritized for display.

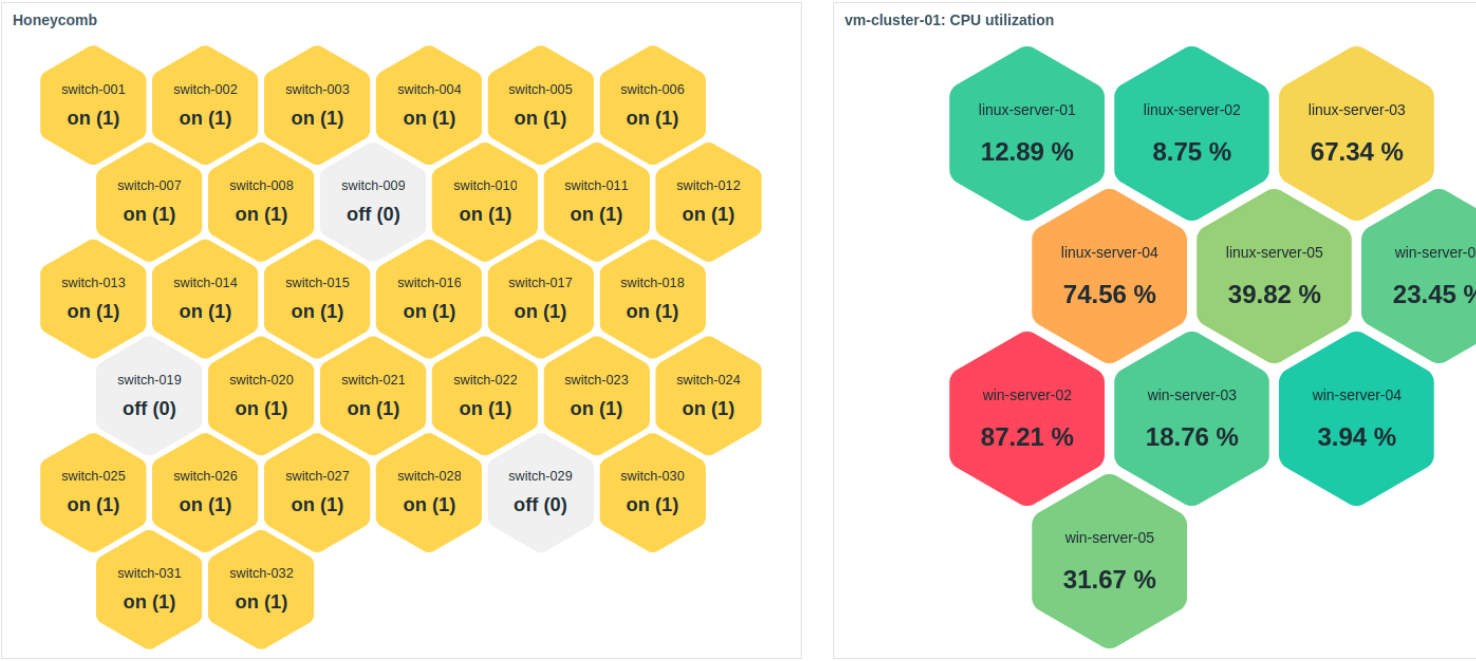
While the *Columns* and *Rows* parameters allow fitting more than one graph in the widget, there still may be more discovered graphs than there are columns/rows in the widget. In this case, paging becomes available in the widget, and a slide-up header allows you to switch between pages using the left and right arrows:



Honeycomb

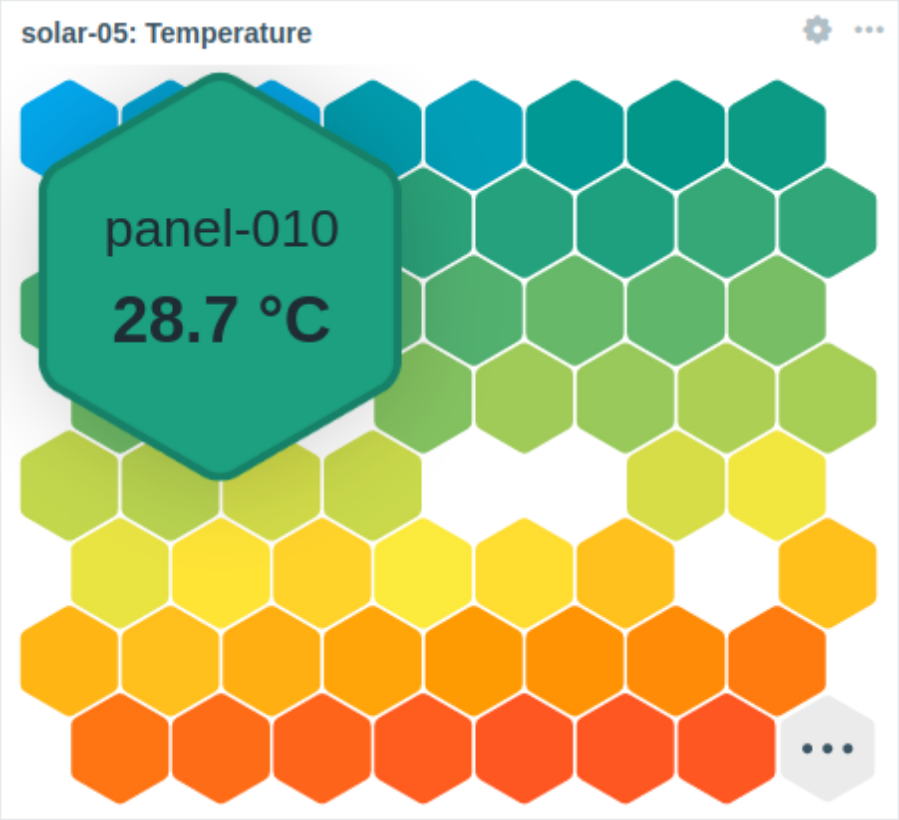
Overview

The *Honeycomb* widget displays hosts or items as hexagonal cells in a grid. This layout makes it easy to detect problem hosts/items, spot clusters of issues, compare groups at a glance, and more.



The number of displayed cells is limited by the widget size and minimum cell size (32px). If all cells do not fit, the last cell shows an ellipsis. The widget can be resized, with cell count, size, and positioning adjusting dynamically.

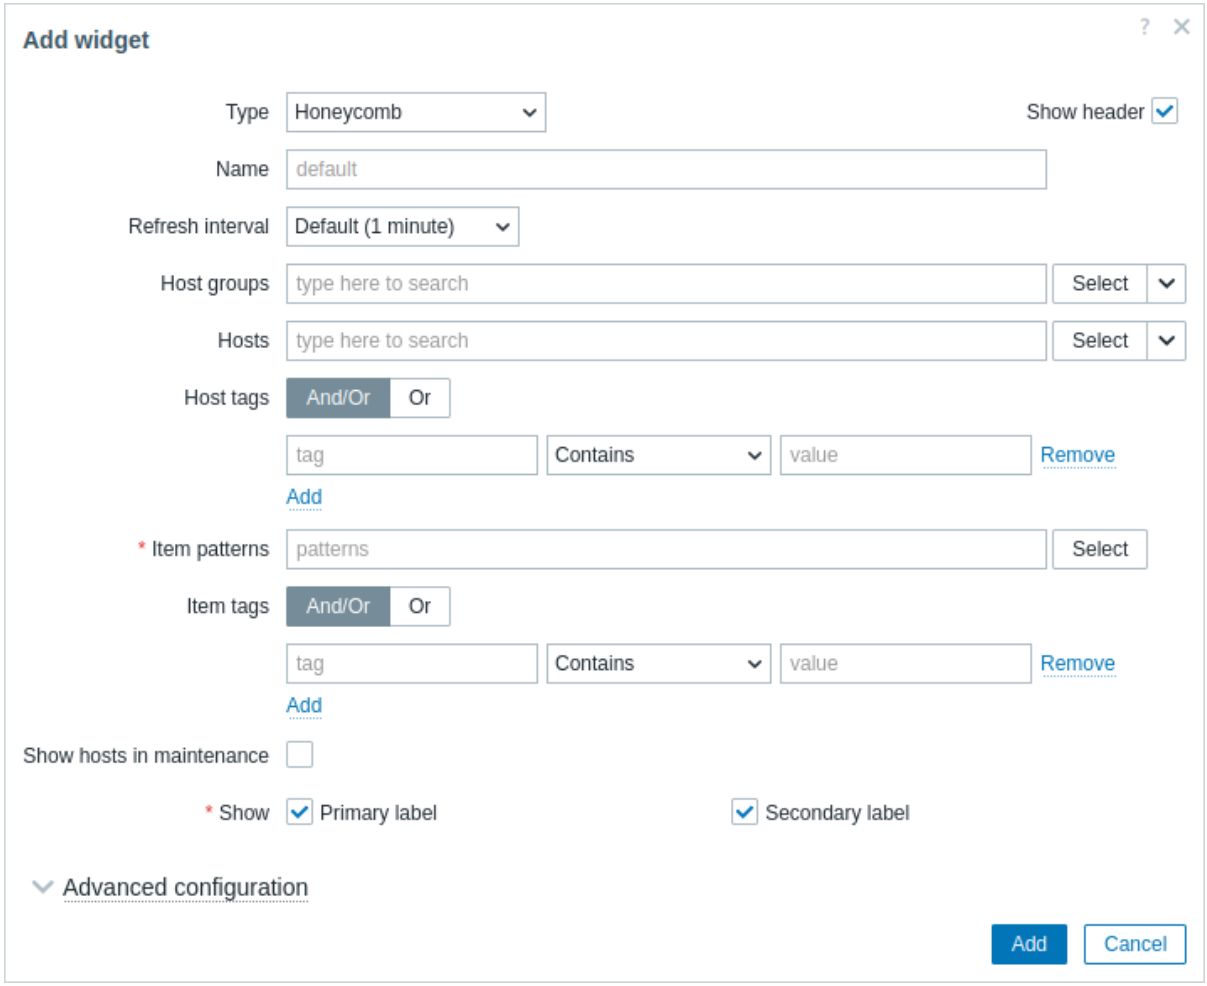
On mouseover, the focused cell enlarges for improved visibility. Clicking a cell highlights its border until another cell is selected.



The information displayed in the *Honeycomb* widget can be downloaded as a PNG image by selecting the *Download image* option in the widget *Actions* menu ⋯ that appears when you hover over the widget. Depending on your browser settings, the image is either automatically saved to the `Downloads` folder, or a *Save* dialog opens, allowing you to choose the location.

Configuration

To configure, select *Honeycomb* as type:



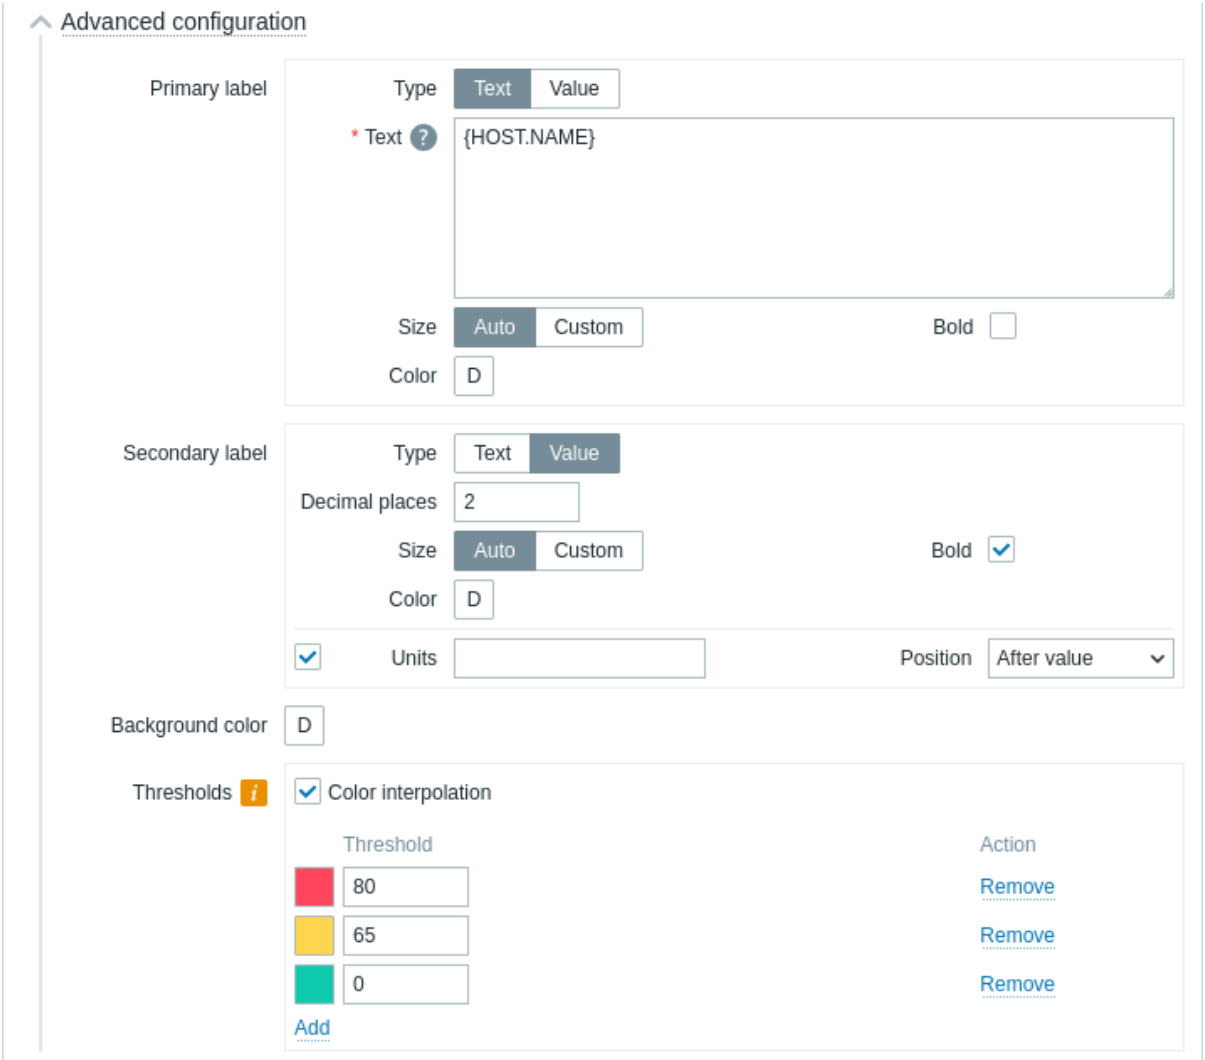
In addition to the parameters that are common for all widgets, you may set the following specific options:

Host groups	Select host groups. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Selecting a parent host group implicitly selects all nested host groups; if no host groups are selected, the widget will display all host groups containing hosts with items matching the selected *Item pattern* (see below). This parameter is not available when configuring the widget on a template dashboard.
Hosts	Select hosts. Alternatively, select a compatible widget or the dashboard as the data source for hosts. This field is auto-complete, so starting to type the name of a host will offer a dropdown of matching hosts. If no hosts are selected, the widget will display all hosts with items matching the selected *Item pattern* (see below). This parameter is not available when configuring the widget on a template dashboard.

Host tags	Specify tags to limit the number of hosts displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met. This parameter is not available when configuring the widget on a template dashboard.
Item patterns	Enter item patterns or select existing items as item patterns. Data of items that match the entered or selected patterns will be displayed on the honeycomb. The *Item patterns* parameter is mandatory. Wildcard patterns may be used for selection (for example, `*` will return items that match zero or more characters; `Zabbix*` will return items that start with "Zabbix"). To specify a wildcard pattern, enter the string manually and press *Enter*. When you start typing, a dropdown list will show matching items limited to those belonging to selected *Hosts* or hosts within selected *Host groups*, if any. The wildcard symbol is always interpreted, therefore, it is not possible to add, for example, an item named *item** individually, if there are other matching items (e.g., item2, item3). When configuring the widget on a template dashboard, this parameter allows selecting only items configured on the template.
Item tags	Specify tags to limit the number of items displayed in the widget. For more information, see *Host tags* above.
Show hosts in maintenance	Mark this checkbox to display hosts in maintenance (in this case, maintenance icon will be shown next to the host name). This parameter is labeled *Show data in maintenance* when configuring the widget on a template dashboard.
Show	Mark this checkbox to display the respective honeycomb cell element - primary label, secondary label. At least one element must be selected.
Advanced configuration	Click the *Advanced configuration* header to display advanced configuration options.

Advanced configuration

Advanced configuration options are available in the collapsible *Advanced configuration* section and only for elements selected in the *Show* field (see above) and the background color or thresholds for honeycomb cells.



Primary/Secondary label

Type — Select the label type:
Text - the label will display the text specified for the *Text* parameter;
Value - the label will display the item value with decimal places as specified for the *Decimal places* parameter.

Text — Enter the label text. This text may override the default item name.
Multiline text is supported. A combination of text and supported macros is possible.
`{HOST.\*}`, `{ITEM.\*}`, `{INVENTORY.\*}`, and user macros are supported.
Honeycomb cells are ordered alphabetically by host name, and, within each host, by item name.
This parameter is available if *Type* is set to *Text*.

Decimal places — Enter the number of decimal places to display with the value.
This parameter is available if *Type* is set to *Value*, and affects only items that return numeric (float) data.

Size — Select the label size:
Auto - use automatically adjusted label size;
Custom - enter a custom label size (in percent, relative to the honeycomb cell size).
Note that labels that do not fit the honeycomb cell size are truncated.

Bold — Mark the checkbox to display labels in bold.

Color — Select the label color from the color picker.
"D" stands for the default color, which depends on the frontend theme. To return to the default color, click the *Use default* button in the color picker.

Units

Units — Mark the checkbox to display units with the item value.
If you enter a unit name, it will override the units set in the item configuration.
This parameter is available if *Type* is set to *Value*.

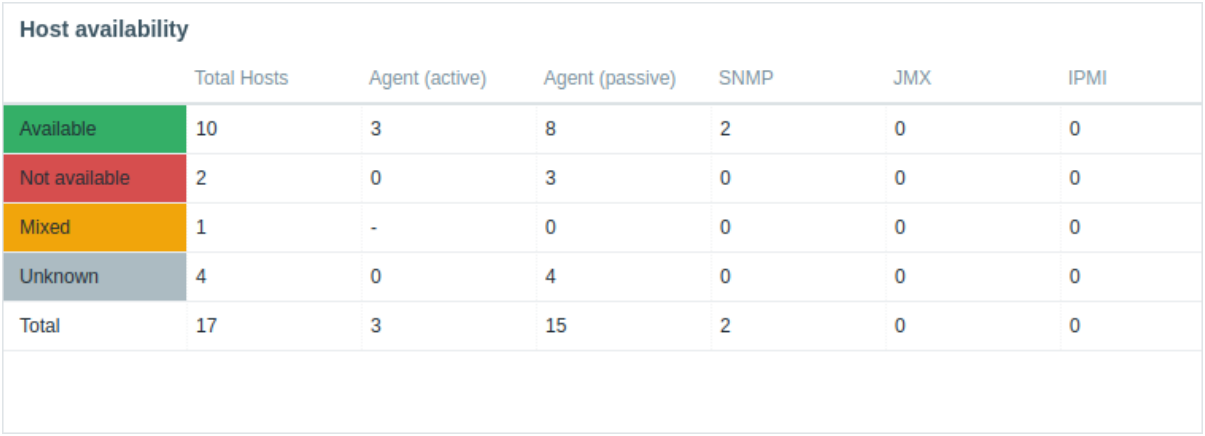
Position — Select the position of the item units (before or after the item value).
This parameter is ignored for the following time-related units: unixtime, uptime, s.
This parameter is available if *Type* is set to *Value* and the *Units* option is selected.

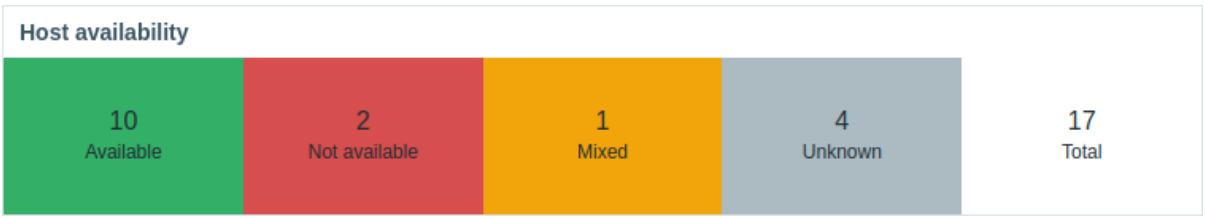
Background color	
Background color	Select the honeycomb cells background color from the color picker. "D" stands for the default color, which depends on the frontend theme. To return to the default color, click the *Use default* button in the color picker.
Thresholds	
Color interpolation	Mark the checkbox to enable smooth transitioning between threshold colors for honeycomb cells. This parameter is available if two or more thresholds are set.
Threshold	Click *Add* to add a threshold, select a threshold color from the color picker, and specify a numeric value. The thresholds list will be sorted in ascending order when saved. Note that the colors configured as thresholds will be displayed correctly only for numeric items. Suffixes (for example, "1d", "2w", "4K", "8G") are supported. Value mappings are supported.

Host availability

Overview

The *Host availability* widget displays host availability across selected host groups, allowing you to monitor the number of hosts that are up or down.



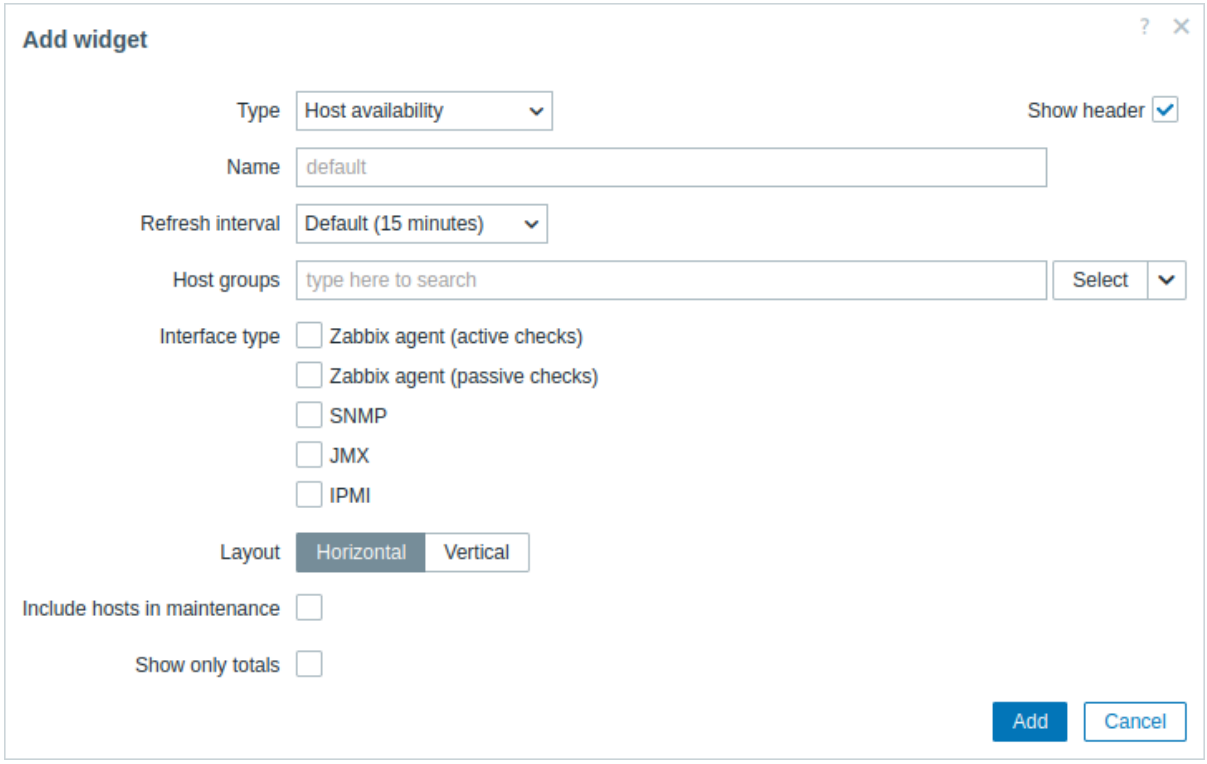


Host availability is counted as follows:

- *Available* - hosts with all interfaces available
- *Not available* - hosts with all interfaces not available
- *Mixed* - hosts with at least one interface not available and at least one available or unknown
- *Unknown* - hosts with at least one interface unknown, but none not available
- *Total* - total of all hosts

Configuration

To configure, select *Host availability* as type:



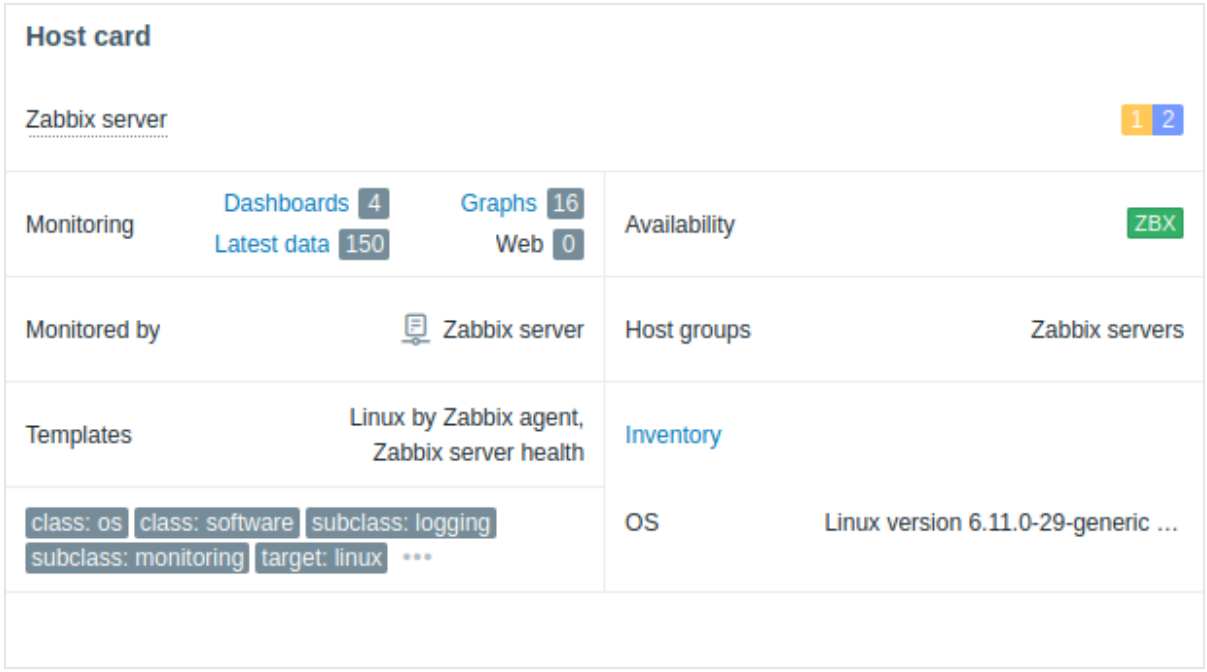
In addition to the parameters that are common for all widgets, you may set the following specific options:

Host groups	Select host groups. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. This parameter is not available when configuring the widget on a template dashboard.
Interface type	Select which host interfaces you want to see availability data for—*Zabbix agent (active checks)*, *Zabbix agent (passive checks)*, *SNMP*, *JMX*, *IPMI*. For *Zabbix agent (active checks)*, the *Mixed* cell will always be empty since this type of items cannot have multiple interfaces. Availability of all interfaces is displayed by default if nothing is selected.
Layout	Select horizontal display (columns) or vertical display (lines).
Include hosts in maintenance	Include hosts that are in maintenance in the statistics. This parameter is labeled *Show data in maintenance* when configuring the widget on a template dashboard.
Show only totals	If checked, the total of hosts, without breakdown by interfaces, is displayed. This option is disabled if only one interface is selected.

Host card

Overview

The *Host card* widget displays details of a single host, allowing you to quickly assess the status and configuration of a host (availability, problems, inventory, etc.).

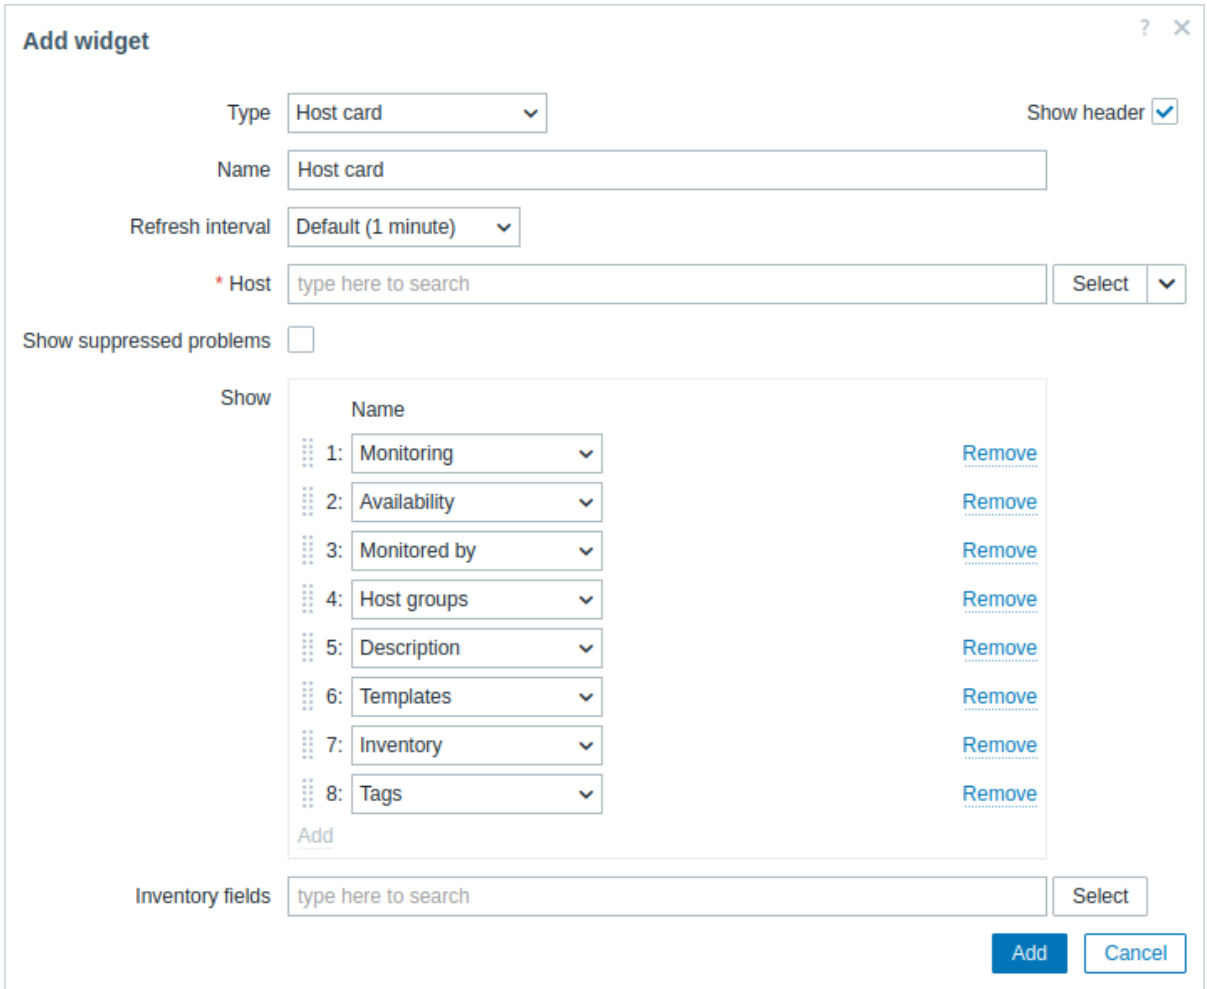


The host can be set directly in the widget configuration or selected from either a compatible widget or the host selector on the dashboard.

The widget can display multiple sections, each showing different information. The section layout is adjusted automatically based on the widget width; when expanded horizontally, sections are automatically rearranged into multiple columns.

Configuration

To configure, select *Host card* as type:



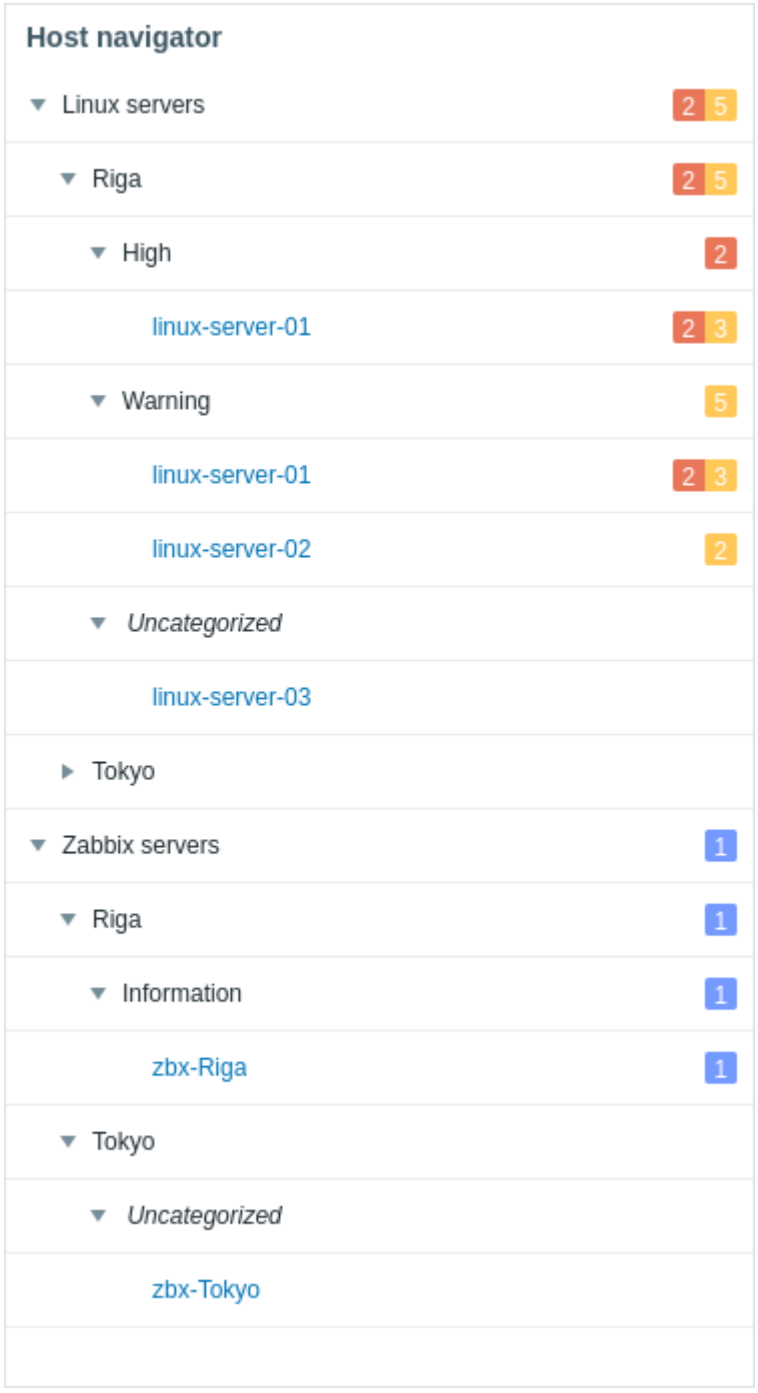
In addition to the parameters that are common for all widgets, you may set the following specific parameters:

Host	Select the host. Alternatively, select a compatible widget or the dashboard as the data source for hosts. Matching suggestions appear as you type. The top section of the widget always displays the **host name** (clicking it opens the host menu): - If the host is disabled, a red *Disabled* label is displayed next to the host name. - If the host is in maintenance, the maintenance icon is displayed next to the host name. The top section of the widget also displays the number of **host problems** grouped by severity (clicking them opens host problems). This parameter is not available when configuring the widget on a template dashboard.
Show suppressed problems	Mark this checkbox to include suppressed problems into problem calculation.
Show	Add sections to display in the widget. Sections can be rearranged by dragging and dropping.
Host groups	Displays host groups the host belongs to.
Description	Displays host description.
Monitoring	Displays quick links for navigating to host Dashboards, Latest data, Graphs, and Web scenarios. The number next to each link indicates the number of respective entities.
Availability	Displays host availability by interface.
Monitored by	Displays whether the host is monitored by Zabbix server or a specific proxy or proxy group.
Templates	Displays templates linked to the host.
Inventory	Displays host inventory fields.
Tags	Displays host tags.
Inventory fields	Select inventory fields to be displayed. Matching suggestions appear as you type. If no inventory fields are specified, all populated inventory fields will be displayed. This parameter is available if *Show* contains *Inventory*.

Host navigator

Overview

The *Host navigator* widget displays a list of hosts based on various filtering and grouping options.



Groups by which hosts are organized can be expanded or collapsed.

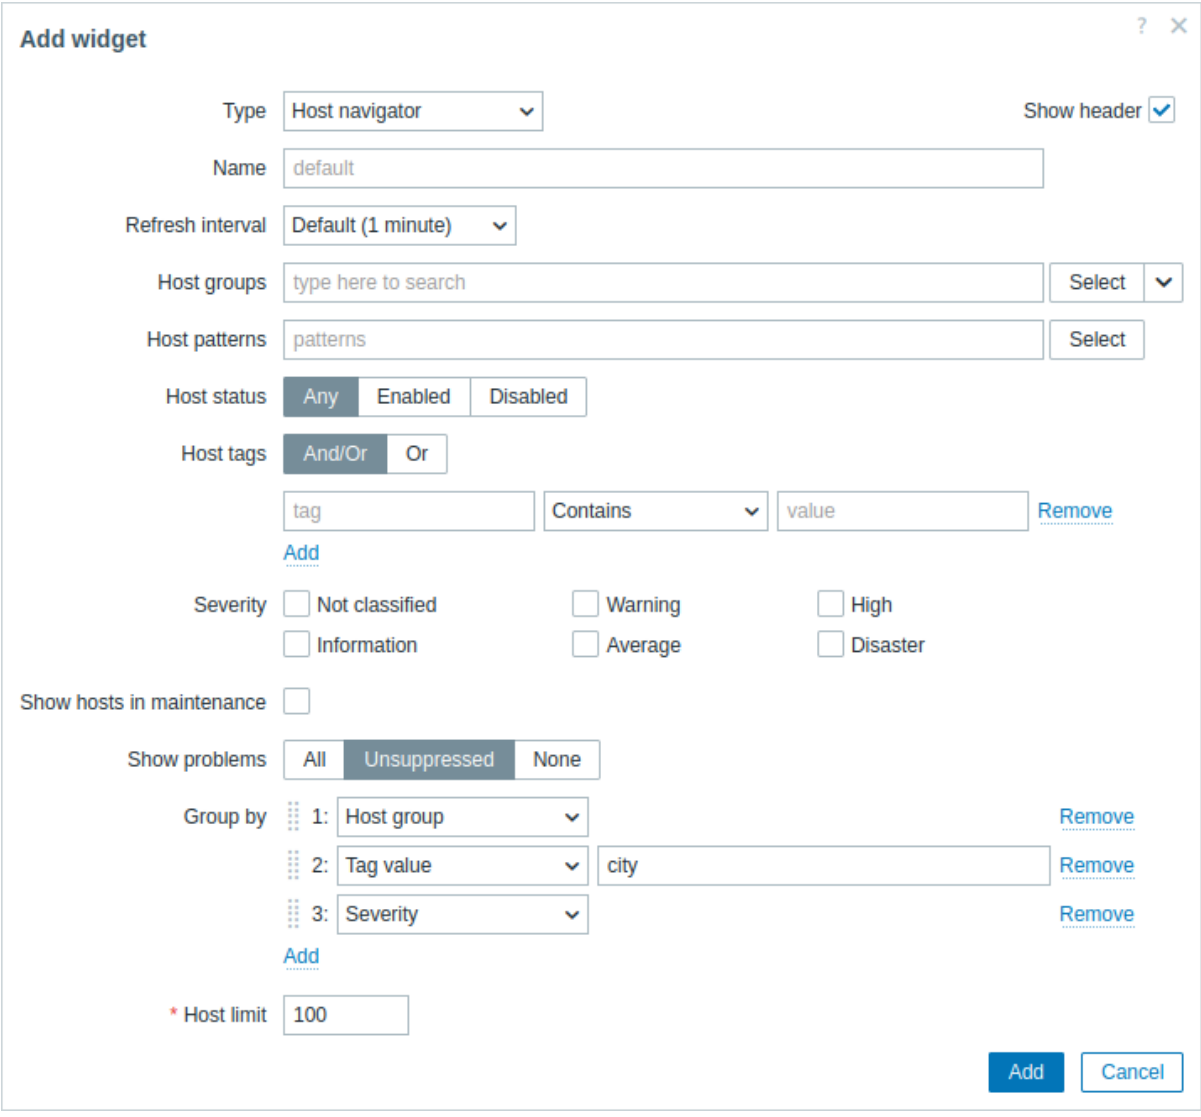
For groups, problems, and hosts in maintenance, additional details are accessible by mouseover hints.

The widget is most useful for controlling what other widgets display based on the selected host.



Configuration

To configure, select *Host navigator* as type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

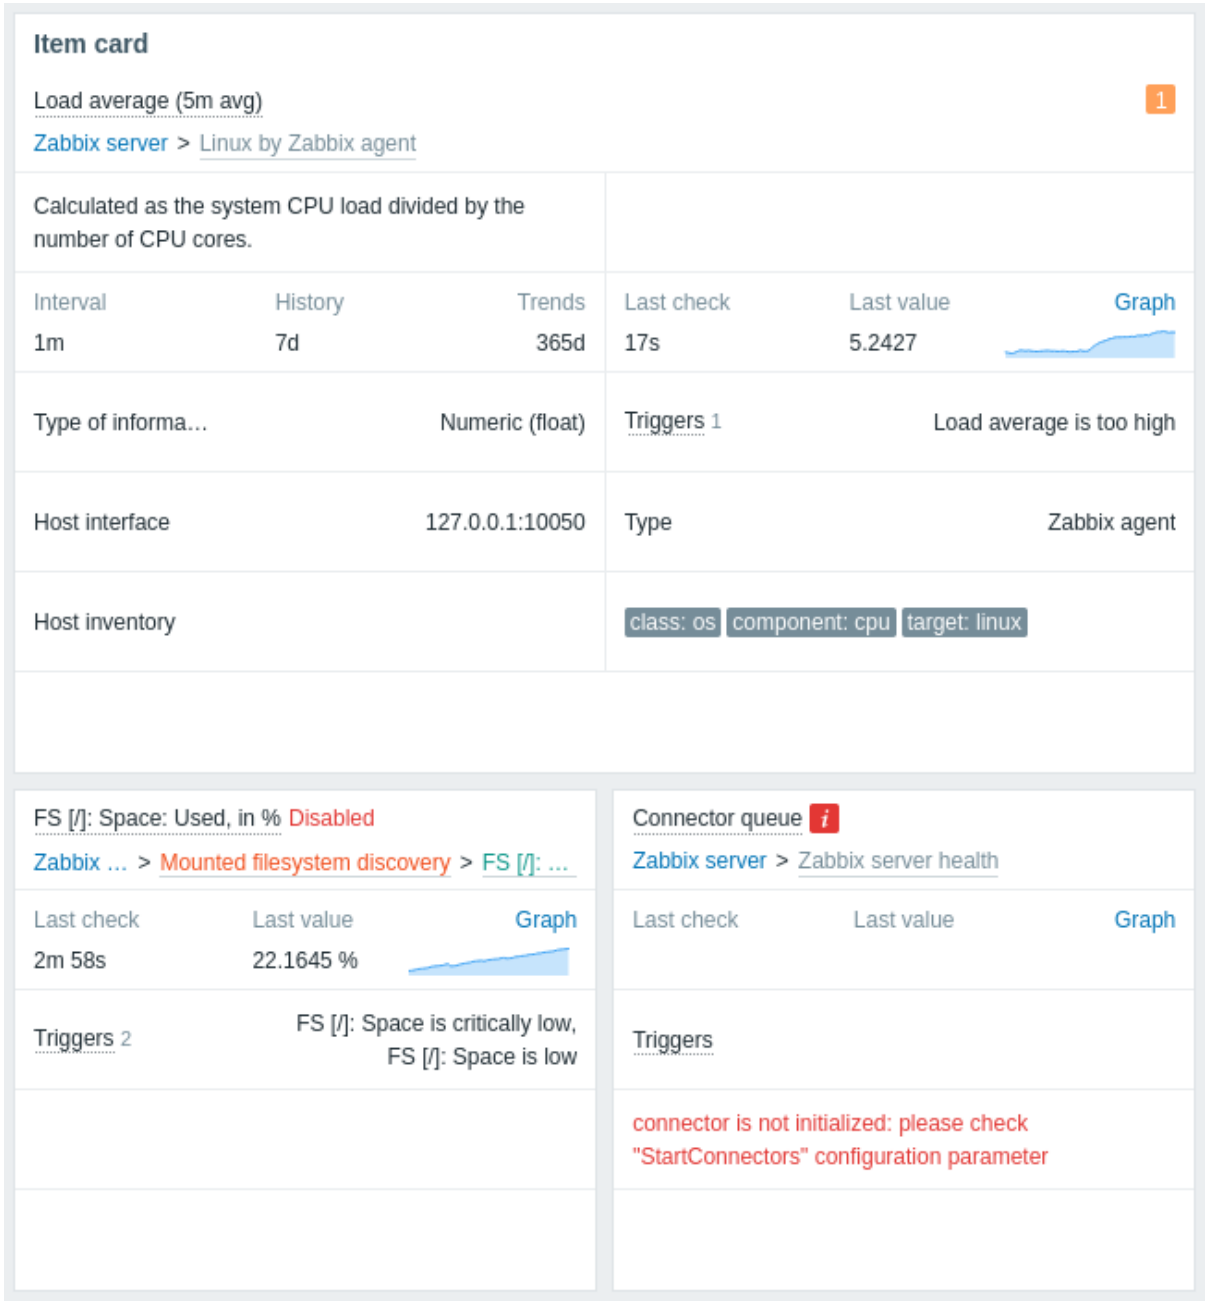
Host groups	Select host groups. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Selecting a parent host group implicitly selects all nested host groups; if no host groups are selected, the widget will display all hosts from all host groups. This parameter is not available when configuring the widget on a template dashboard.
Host patterns	Enter host patterns or select existing hosts as host patterns. Hosts that match the specified patterns will be displayed on the host navigator. This field is auto-complete, so starting to type the name of a host will offer a dropdown of matching hosts. If no hosts are selected, the widget will display all hosts. Wildcard patterns may be used for selection (for example, `*` will return hosts that match zero or more characters; `Zabbix*` will return hosts that start with "Zabbix"). To specify a wildcard pattern, enter the string manually and press *Enter*. When you start typing, a dropdown list will show matching hosts limited to those belonging to hosts within selected *Host groups*, if any. This parameter is not available when configuring the widget on a template dashboard.
Host status	Filter which hosts to display based on their status (any, enabled, disabled). This parameter is not available when configuring the widget on a template dashboard.
Host tags	Specify tags to filter the hosts displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met. This parameter is not available when configuring the widget on a template dashboard.
Severity	Mark problem severities to filter hosts with problems to be displayed in the widget. If no severities are marked, all hosts with all problems will be displayed.
Show hosts in maintenance	Mark this checkbox to display hosts in maintenance (in this case, maintenance icon will be shown next to the host name). This parameter is labeled *Show data in maintenance* when configuring the widget on a template dashboard.
Show problems	Filter which problems to display based on their status (all, unsuppressed, none). Problems are indicated with colored blocks at the end of the host row. Colors are based on problem severity, which can be customized in *Administration* > *General* > *Trigger displaying options*. Problem severities can be changed when updating problems.

Group by	Add a grouping attribute by which to group the selected hosts: **Host group** - group hosts by their host group; **Tag value** - enter a tag name to group hosts by the values of this tag (for example, enter ”city” to group hosts by values ”Riga”, ”Tokyo”, etc.); **Severity** - group hosts by their problem severities. If *Show problems* is configured to display problems, they are displayed as follows: - for each severity group, only the corresponding problem count is displayed; - for each host, all its problem counts are displayed. Note that hosts will be grouped only by those severities marked in the *Severity* parameter; if no severities are marked, all hosts will be grouped by all severities. Grouping attributes can be reordered by dragging up or down by the handle before the group name. Note that grouping attribute order determines group nesting order. For example, specifying multiple tag names (1: color, 2: city) will result in hosts being grouped by color (red, blue, etc.) and then by city (Riga, Tokyo, etc.). A host may be displayed in multiple groups depending on the configured grouping attributes (for example, when grouping by host group and the host belongs to multiple host groups). Clicking such hosts selects and highlights them in all groups. Hosts that do not match the configured grouping attributes are displayed in the *Uncategorized* group. Up to 10 grouping attributes can be specified, and all must be unique. If no grouping attributes are specified, hosts will not be grouped.
Host limit	Enter the maximum number of hosts to be displayed. Possible values range from 1 to 9999. When more hosts are available for displaying than the set limit, a corresponding message is shown below the displayed hosts (for example, ”100 of 100+ hosts are shown”). Note that the configured host limit also affects the display of configured groups; for example, if host limit is set to 100 and hosts are grouped by tag values (more than 200), only the first 100 tag values with the corresponding hosts will be displayed in the widget. This parameter is not affected by the *Limit for search and filter results* parameter in *Administration > General > GUI*. This parameter is not available when configuring the widget on a template dashboard.

Item card

Overview

The *Item card* widget displays details of a single item, allowing you to quickly assess the status and configuration of an item (latest data, errors, description, etc.).

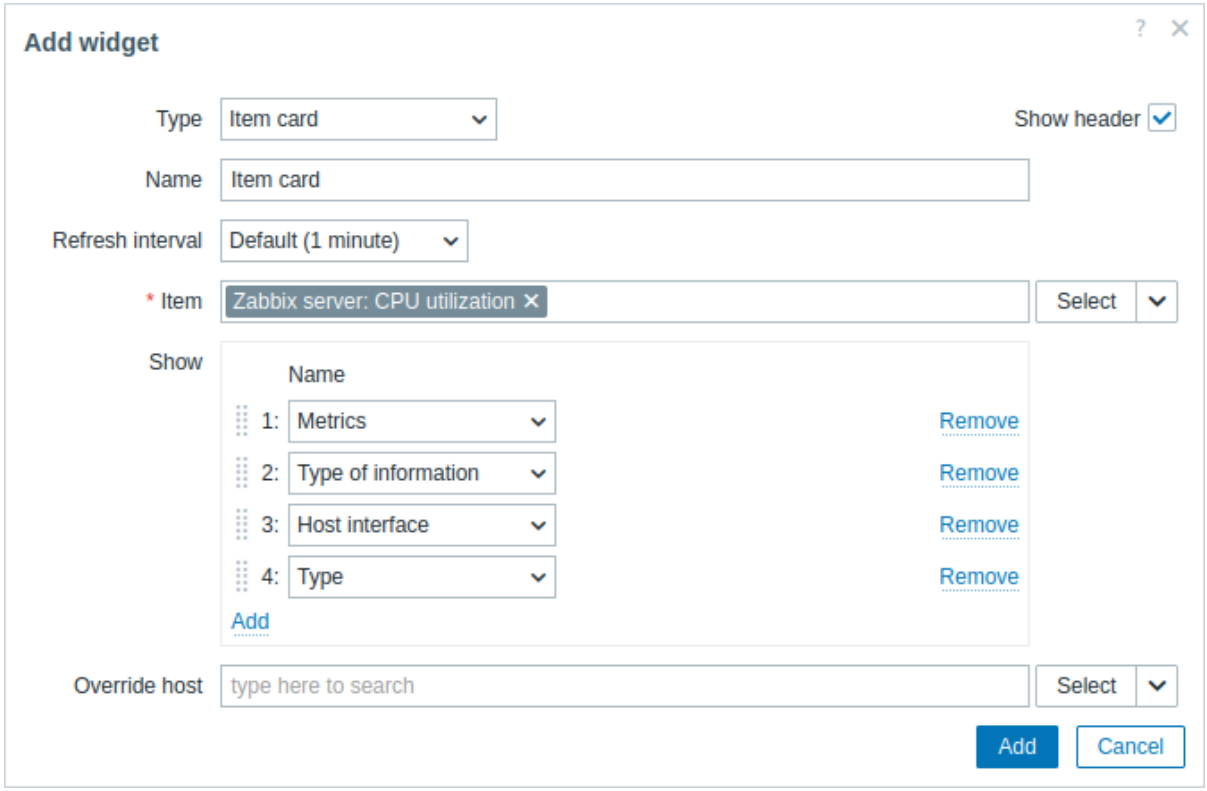


The item can be set directly in the widget configuration or selected from a compatible widget on the dashboard.

The widget can display multiple sections, each showing different information. The section layout is adjusted automatically based on the widget width; when expanded horizontally, sections are automatically rearranged into multiple columns.

Configuration

To configure, select *Item card* as type:



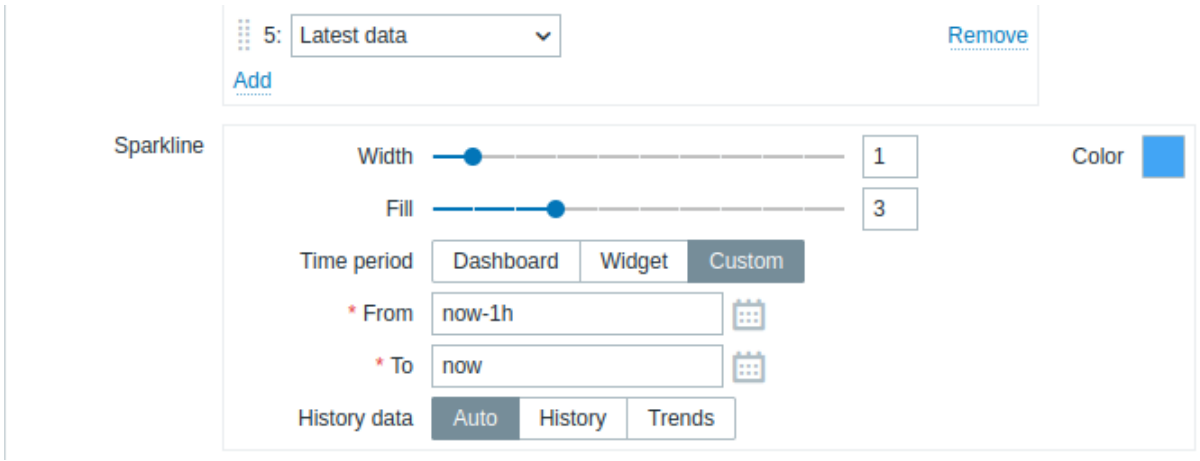
In addition to the parameters that are common for all widgets, you may set the following specific options:

Item	Select the item. Alternatively, select a compatible widget as the data source for items. Matching suggestions appear as you type. The top section of the widget always displays the **item name** (clicking it opens the item menu): - If the item is disabled, a red *Disabled* label is displayed next to the item name. - In case of errors, a square icon with the letter "i" is displayed. In case of problems, a square icon with the problem count is displayed for each problem severity. Hovering over an icon opens a pop-up with details. - On a template dashboard, the technical name (instead of visible name) of the item is displayed. The top section of the widget also displays the **item path**—its host and, if applicable, template, low-level discovery (LLD) rule, or master item (clicking an entity opens its configuration). Item path is displayed only to users with permissions to the respective entities.
Show	Add sections to display in the widget. Sections can be rearranged by dragging and dropping.
Description	Displays item description. Macros are resolved.
Error text	Displays item error (if any).
Metrics	Displays the following information: - Item update interval (except for trapper items); - Item history configuration (if any); - Item trends configuration (if any). Macros are resolved. If the item uses custom intervals, a help icon is displayed next to item update interval. Hovering over the icon opens a pop-up with custom interval details.

Latest data	Displays the following information: - Time since last item check; - Item last value; - Link to simple graph/history of item values (if history or trends are stored); - Configurable sparkline chart (if item is numeric). For binary type items, a thumbnail or the *Show* option is displayed instead of last value. Hovering over the thumbnail opens a pop-up with the image. Hovering over *Show* opens a pop-up with the item value (Base64 string). On a template dashboard, *No data* is displayed instead of item last value.
Type of information	Displays item type of information.
Triggers	Displays the following information: - *Triggers* label with trigger count (hovering over *Triggers* opens a pop-up with trigger details); - List of trigger names (macros are resolved).
Host interface	Displays host interface used by the item, or *No data* if the item does not have a host interface (e.g., dependent items, simple checks). On a template dashboard, *No data* is displayed.
Type	Displays item type.
Host inventory	Displays host inventory field that item value populates.
Tags	Displays item tags.
Override host	Select a compatible widget or the dashboard host selector as the data source for hosts. This parameter is not available when configuring the widget on a template dashboard.

Sparkline

Sparkline chart configuration options are available only if the *Latest data* section is added to the widget.



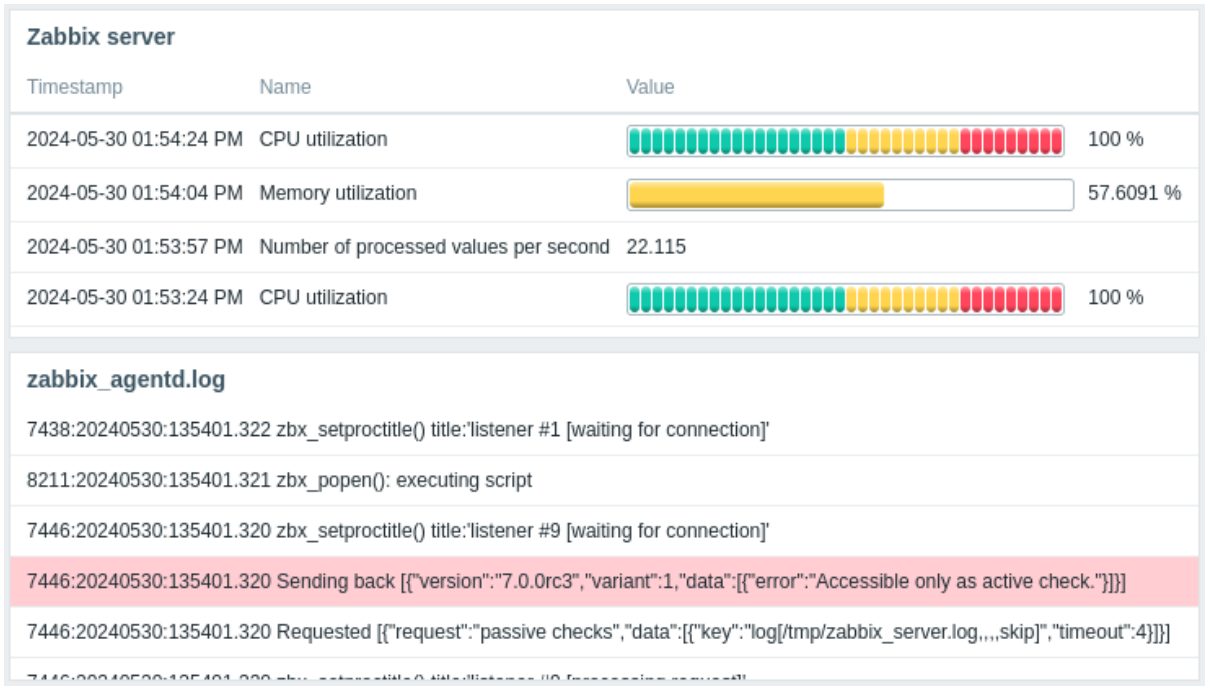
Width	Set the graph line thickness by using the slider or manually entering a value in the range from 0 to 10.
Color	Select line and fill color.
Fill	Set fill color transparency level by using the slider or manually entering a value in the range from 0 to 10.
Time period	Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period.
Widget	Enter or select a compatible widget as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".

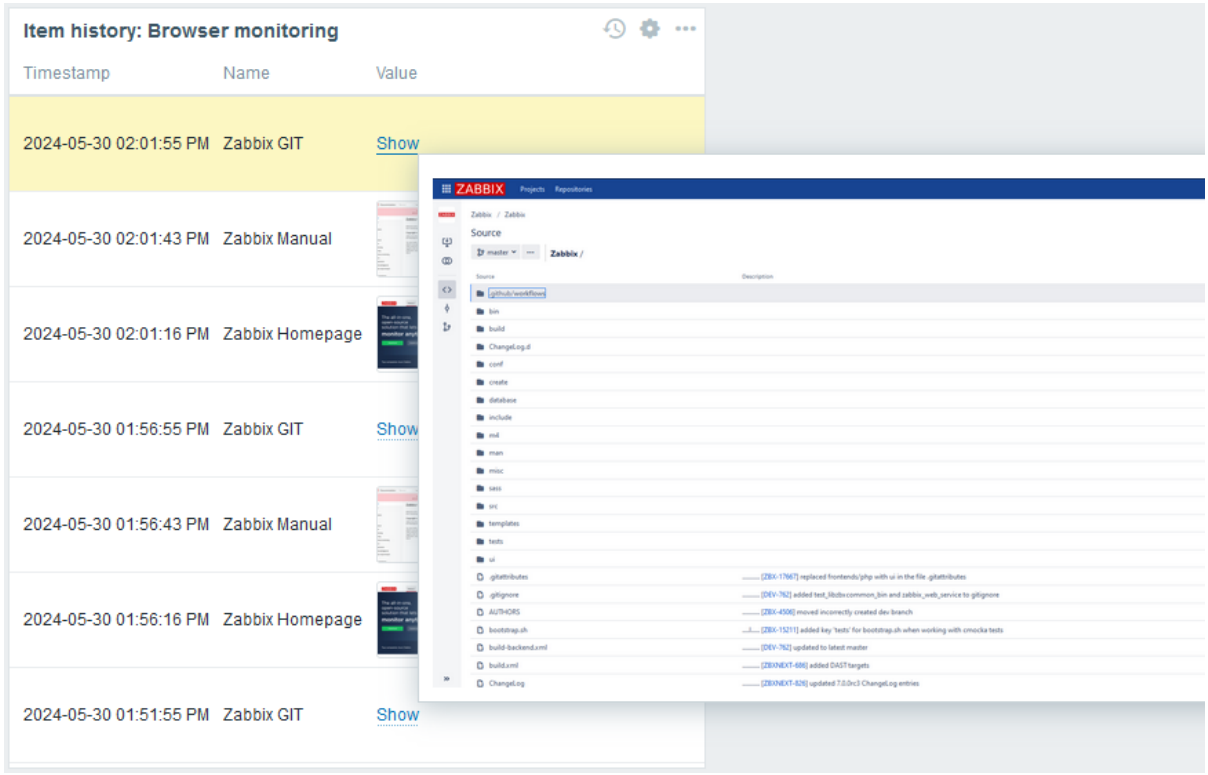
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
History data	Take data from history or trends: **Auto** - automatic selection; **History** - take history data; **Trends** - take trend data. This parameter applies only to numeric data. Non-numeric data will always be taken from history.

Item history

Overview

The *Item history* widget displays the latest data for all item types (numeric, text, etc.) in a table format. It can show progress bars and images (useful for browser items) and can highlight values (useful for log file monitoring).

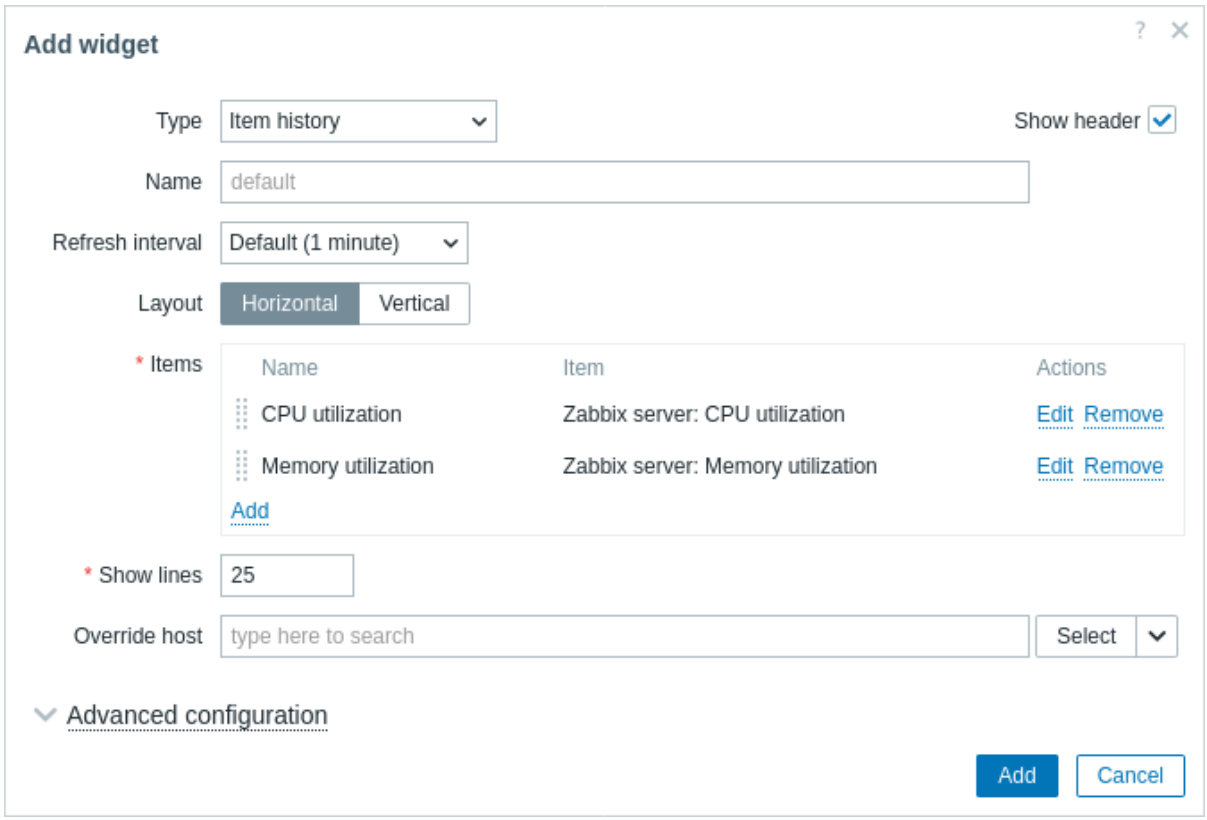




Up to 1000 records can be displayed.

Configuration

To configure, select *Item history* as type:



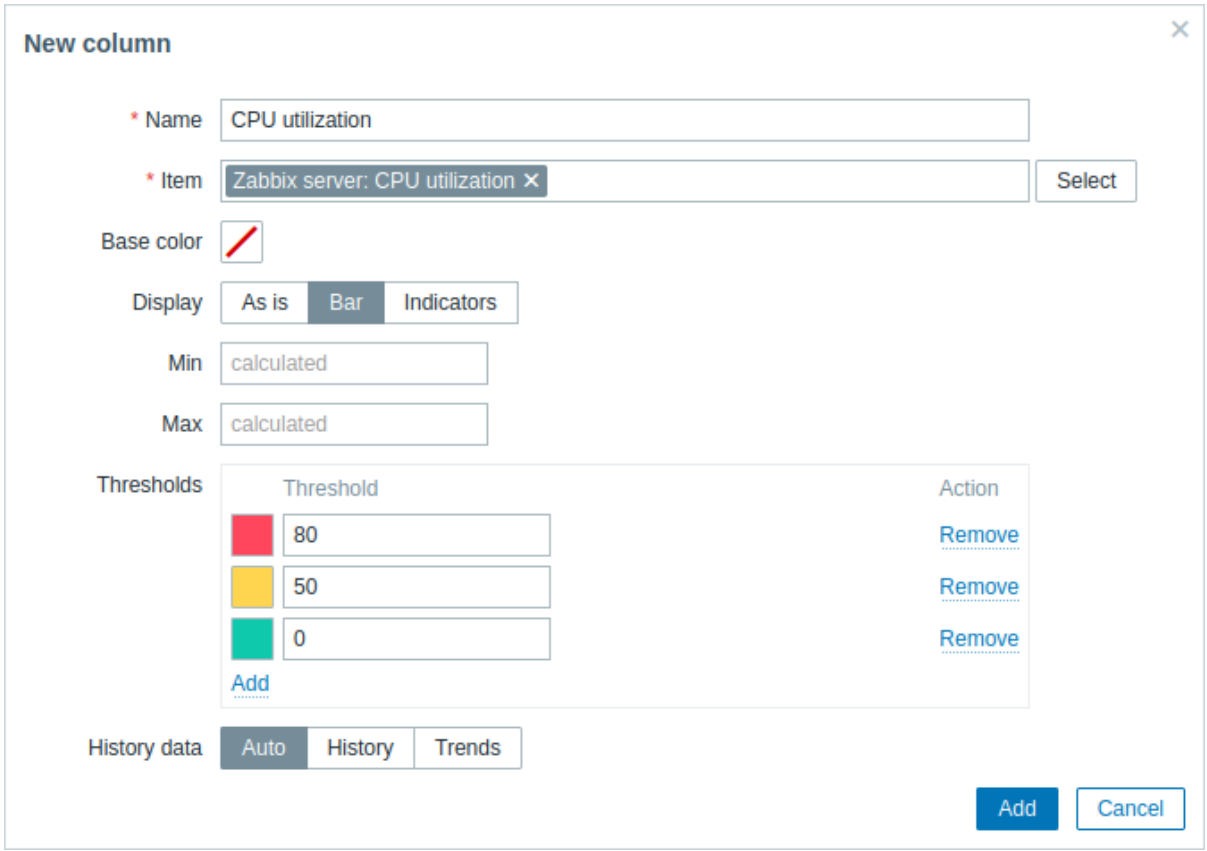
In addition to the parameters that are common for all widgets, you may set the following specific options:

Layout	Select the layout option for item columns: **Horizontal** - items will be displayed horizontally, values vertically; **Vertical** - items will be displayed vertically, values horizontally.
Items	Add item columns to display. The order of items determines their display order. Items can be reordered by dragging them up or down by the handle before the item name.

Show lines	Specify the number of item value lines to display.
Override host	Select a compatible widget or the dashboard host selector as the data source for hosts. This parameter is not available when configuring the widget on a template dashboard.
Advanced configuration	Click the *Advanced configuration* label to display advanced configuration options.

Column configuration

To configure item columns, click *Add* in the *Items* parameter:



Common column parameters:

Name	Enter the name of the column. If left empty, the item name from the *Item* parameter is used.
Item	Select the item. Note that column configuration parameters vary based on the information type of the selected item; for more information, see individual parameters below. When configuring the widget on a template dashboard, only items configured on the template can be selected.
Base color	Select the column's background color or fill color if *Display* is set to "Bar" or "Indicators". Note that the base color can be overridden by threshold or highlight colors.

Column parameters specific to numeric type items:

Display	Select how the item value should be displayed: **As is** - as regular text; **Bar** - as solid, color-filled bar; **Indicators** - as segmented, color-filled bar.
Min	Enter the minimum value for bar/indicators. If left empty, the widget will use the minimum value of the item. This parameter is available only when *Display* is set to "Bar" or "Indicators".
Max	Enter the maximum value for bar/indicators. If left empty, the widget will use the maximum value of the item. This parameter is available only when *Display* is set to "Bar" or "Indicators".

Thresholds	Click *Add* to add a threshold, select a threshold color from the color picker, and specify a numeric value. The thresholds list will be sorted in ascending order when saved. Suffixes (for example, ”1d”, ”2w”, ”4K”, ”8G”) are supported. Value mappings are supported.
History data	Select whether to take data from history or trends: **Auto** - automatic selection; **History** - take history data; **Trends** - take trend data.

Column parameters specific to character, text, log, and JSON type items:

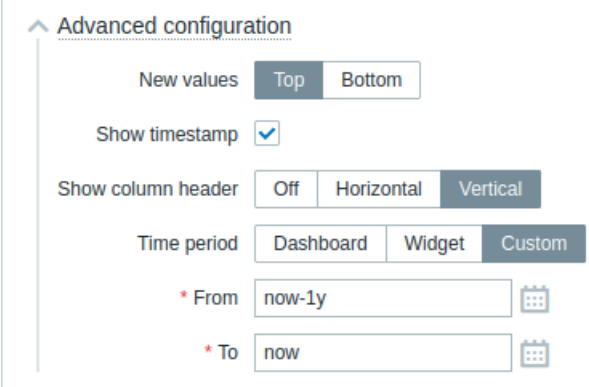
Highlights	Click *Add* to add a highlight, select a highlight color from the color picker, and specify a regular expression. The selected color will be used as the background color for item values where the specified regular expression matches the text.
Display	Select how the item value should be displayed: **As is** - exactly as received. Line breaks are preserved if the value contains line returns. Word wrapping is applied if *Layout* is set to ”Vertical” and more than one column is configured; **HTML** - as HTML-formatted text; **Single line** - as a single line, truncated to a specified length (1-500 characters). Hovering or clicking the truncated value opens a pop-up with the full value.
Use monospace font	Mark this checkbox to display the item value in monospace font (unmarked by default).
Display local time	Mark this checkbox to display local time instead of timestamp in the timestamp column. Note that the *Show timestamp* checkbox in advanced configuration must also be marked. This parameter is available only for log type items.

Column parameters specific to binary type items:

Show thumbnail	Mark this checkbox to display a thumbnail for image binaries or a ”Show” option for non-image binaries. Unmark this checkbox to display a ”Show” option for all binary item values. Hovering or clicking the ”Show” option opens a pop-up window with the item value (image or Base64 string). If the item value is an empty string, the ”Show” option is displayed; hovering or clicking it opens a pop-up containing ”Empty string”.

Advanced configuration

Advanced configuration options are available in the collapsible *Advanced configuration* section:



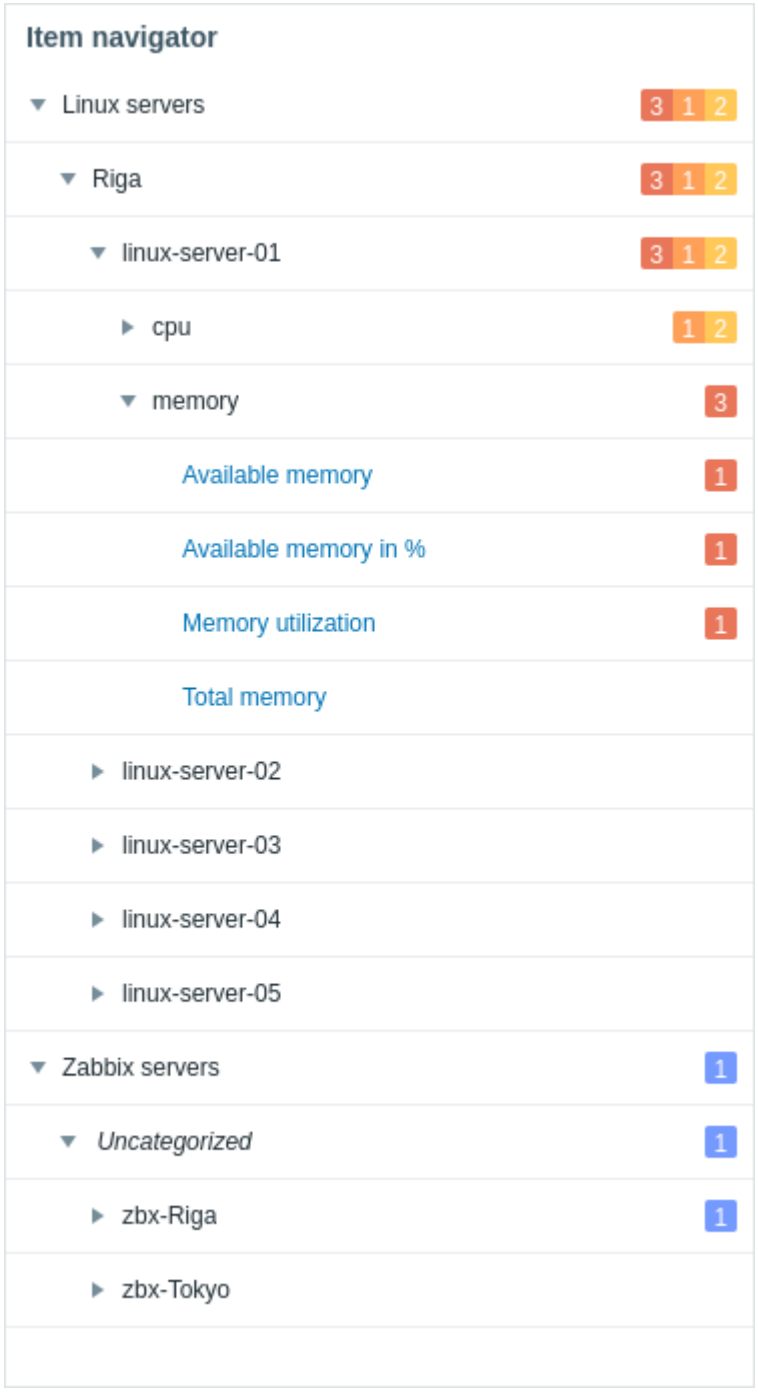
New values	Select where new item values should be added: **Top** - at the top of columns; **Bottom** - at the bottom of columns.
Show timestamp	Mark this checkbox to display the timestamp column (unmarked by default).

Show column header	Select the column header orientation: **Off** - hide the header; **Horizontal** - display the header horizontally; **Vertical** - display the header vertically.
Time period	Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period.
Widget	Enter or select a compatible widget (*Graph*, *Graph (classic)*, *Graph prototype*) as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".

Item navigator

Overview

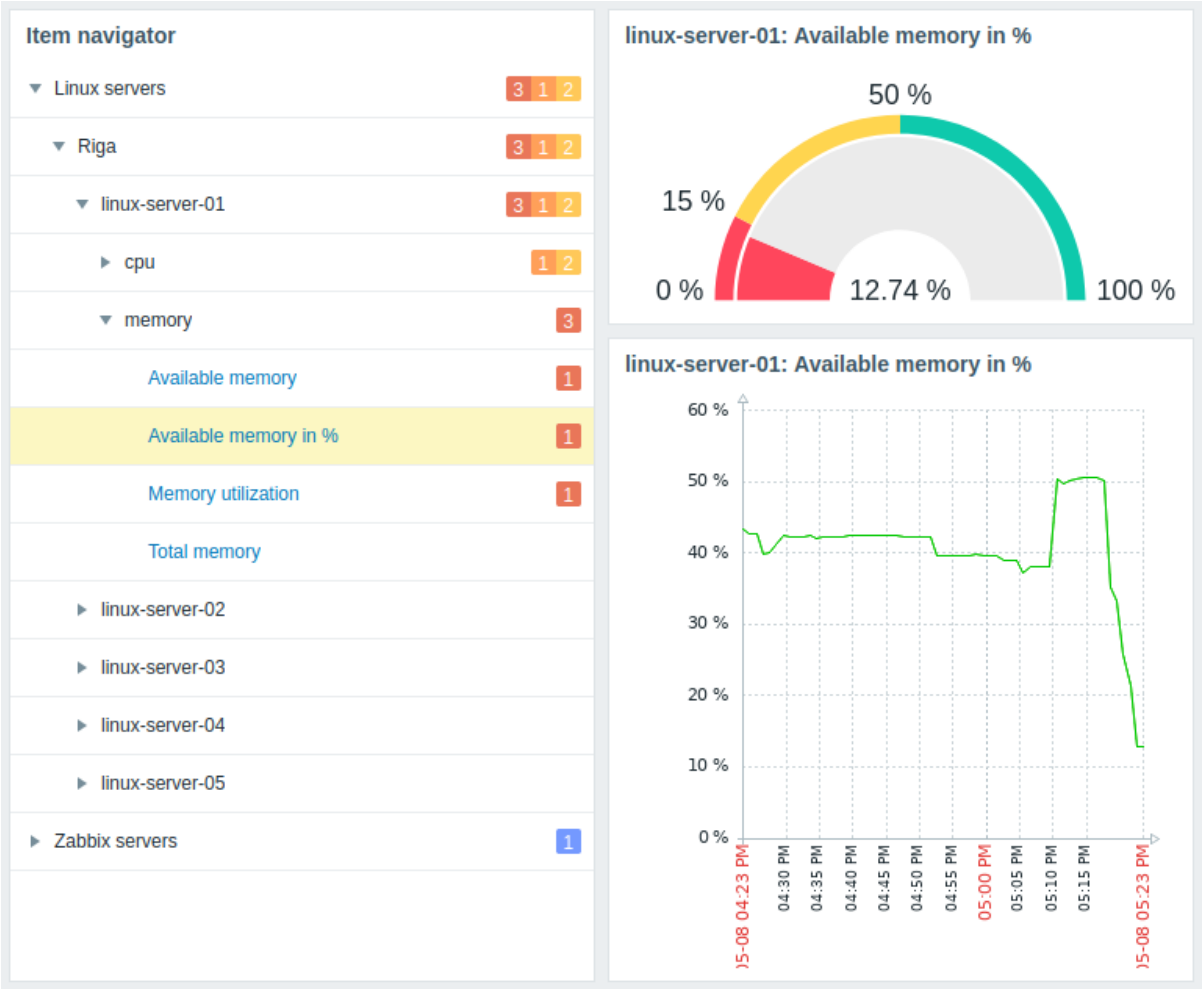
The *Item navigator* widget displays a list of items based on various filtering and grouping options.



Groups by which items are organized can be expanded or collapsed.

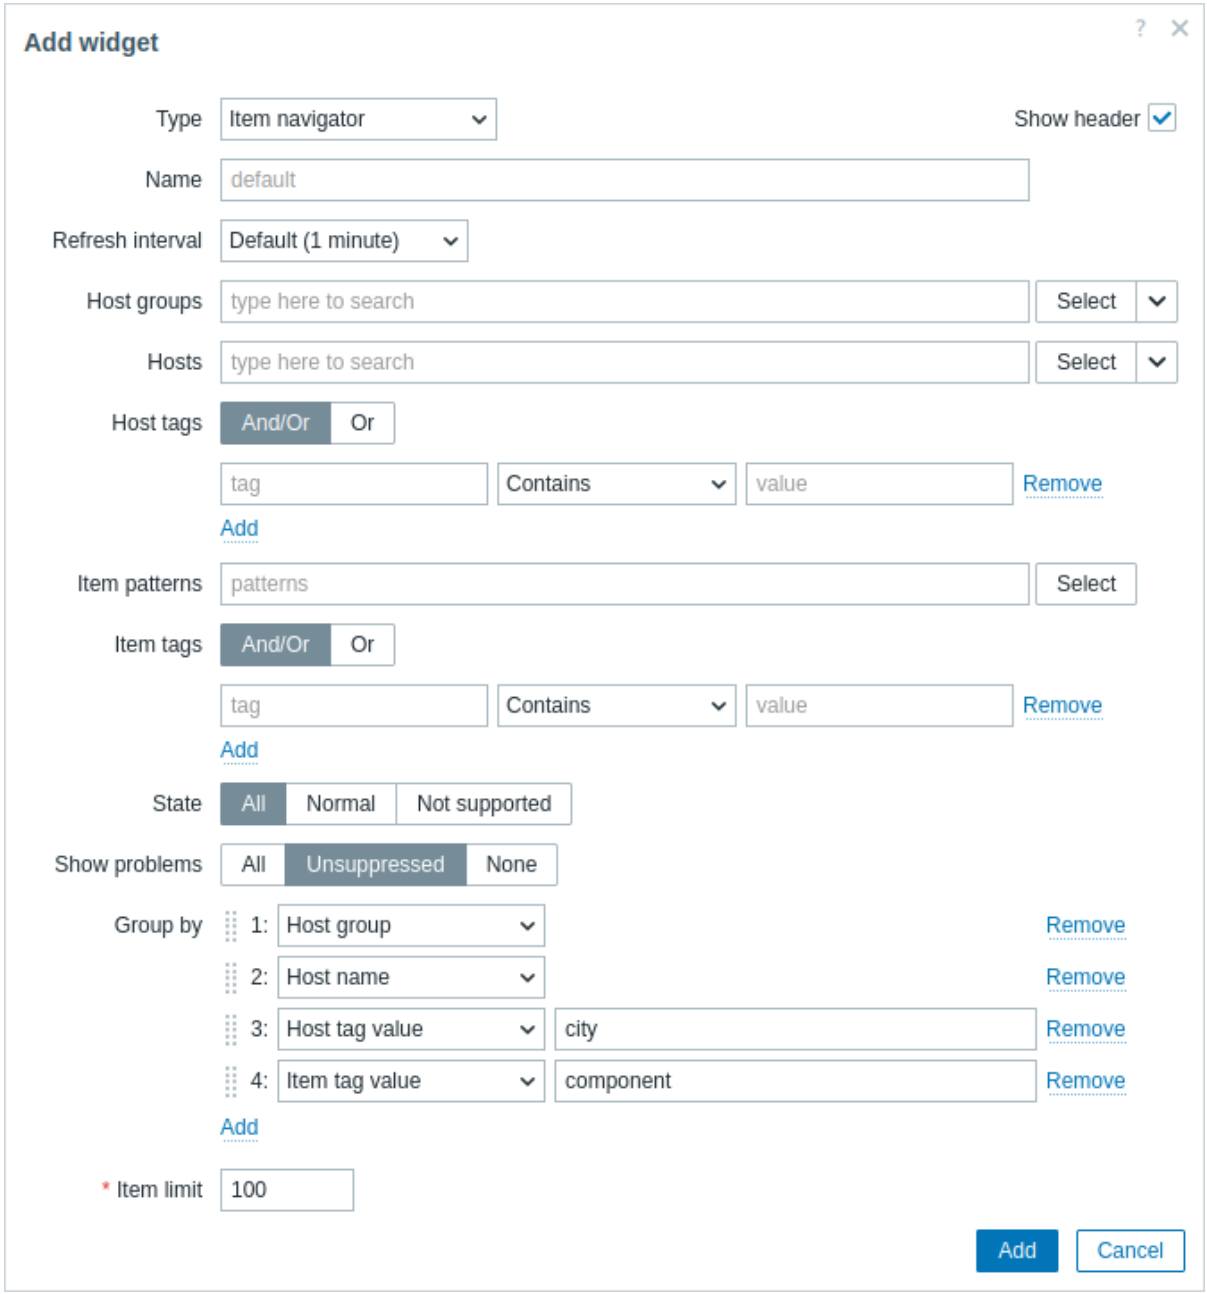
For groups and problems, additional details are accessible by mouseover hints.

The widget is most useful for controlling what other widgets display based on the selected item.



Configuration

To configure, select *Item navigator* as type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

Host groups	Select host groups. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Selecting a parent host group implicitly selects all nested host groups; if no host groups are selected, the widget will display items belonging to all hosts from all host groups. This parameter is not available when configuring the widget on a template dashboard.
Hosts	Select hosts. Alternatively, select a compatible widget or the dashboard as the data source for hosts. This field is auto-complete, so starting to type the name of a host will offer a dropdown of matching hosts. If no hosts are selected, the widget will display items belonging to all hosts. This parameter is not available when configuring the widget on a template dashboard.

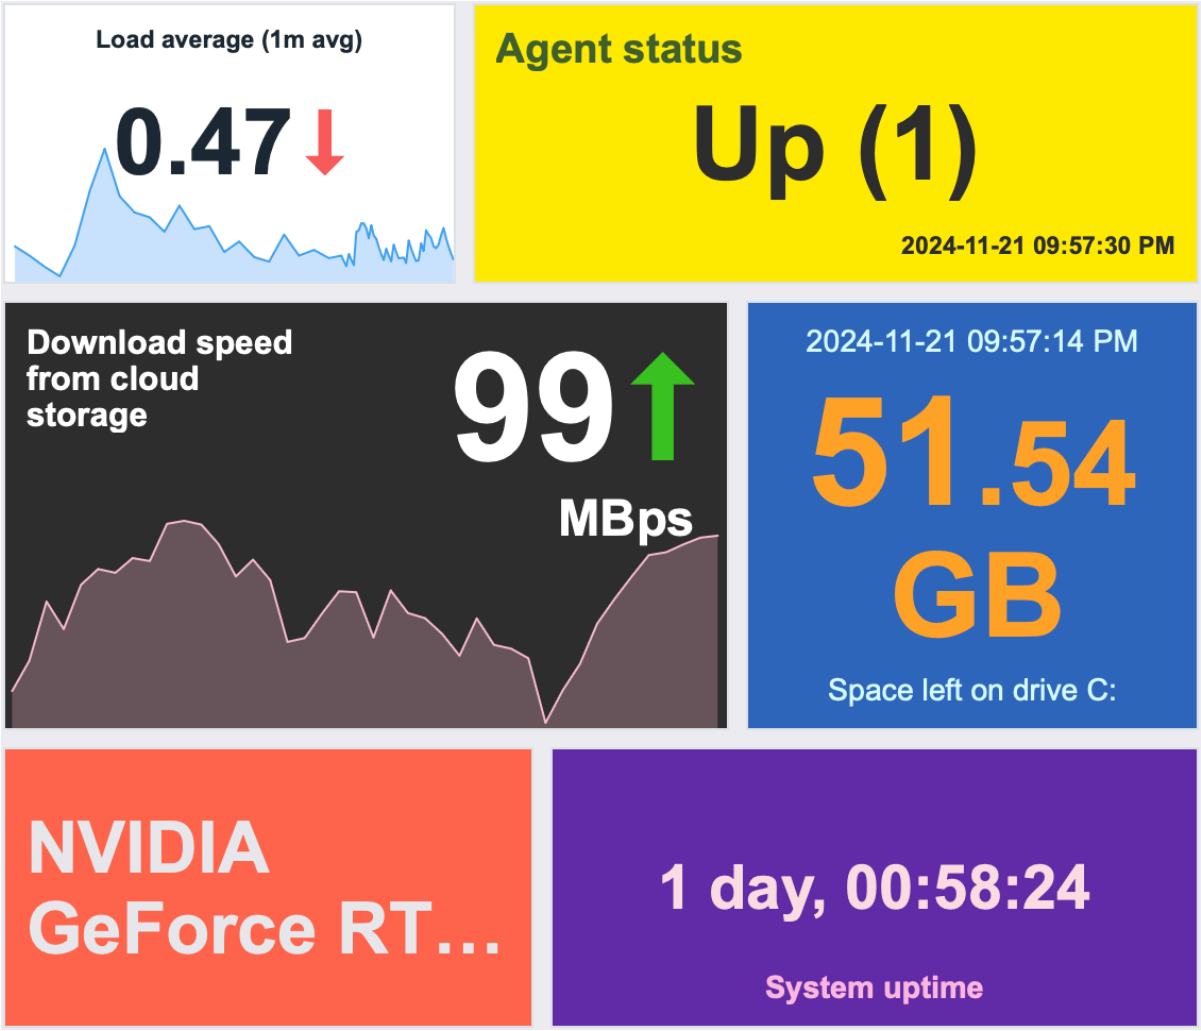
Host tags	Specify host tags to filter the items displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met. This parameter is not available when configuring the widget on a template dashboard.
Item patterns	Enter item patterns or select existing items as item patterns. Items that match the specified patterns will be displayed on the item navigator. Wildcard patterns may be used for selection (for example, ∗ will return items that match zero or more characters; `Zabbix*` will return items that start with "Zabbix"). To specify a wildcard pattern, enter the string manually and press *Enter*. When you start typing, a dropdown list will show matching items limited to those belonging to selected *Hosts* or hosts within selected *Host groups*, if any. The wildcard symbol is always interpreted, therefore, it is not possible to add, for example, an item named *item** individually, if there are other matching items (e.g., item2, item3). When configuring the widget on a template dashboard, this parameter allows selecting only items configured on the template.
Item tags	Specify item tags to filter the items displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met.
State	Filter which items to display based on their state (all, normal, not supported).
Show problems	Filter which problems to display based on their status (all, unsuppressed, none). Problems are indicated with colored blocks at the end of the item row. Colors are based on problem severity, which can be customized in *Administration > General > Trigger displaying options*. Problem severities can be changed when updating problems.

Group by	Add a grouping attribute by which to group items: **Host group** - group items by host groups of their hosts; **Host name** - group items by their hosts; **Host tag value** - enter a tag name to group items by the values of this host tag (for example, enter ”city” to group items by values ”Riga”, ”Tokyo”, etc.); **Item tag value** - enter a tag name to group items by the values of this item tag (for example, enter ”component” to group items by values ”cpu”, ”memory”, etc.). Grouping attributes can be reordered by dragging up or down by the handle before the group name. Note that grouping attribute order determines group nesting order. For example, specifying multiple host tag names (1: color, 2: city) will result in items being grouped by color (red, blue, etc.) and then by city (Riga, Tokyo, etc.). An item may be displayed in multiple groups depending on the configured grouping attributes (for example, when grouping by host group and the item’s host belongs to multiple host groups). Clicking such items selects and highlights them in all groups. Items that do not match the configured grouping attributes are displayed in the *Uncategorized* group. If *Show problems* is configured to display problems, they are displayed as follows: - for each group, all subordinate item problem count is displayed; - for each item, only its problem count is displayed. Up to 10 grouping attributes can be specified, and all must be unique. If no grouping attributes are specified, items will not be grouped.
Item limit	Enter the maximum number of items to be displayed. Possible values range from 1 to 9999. When more items are available for displaying than the set limit, a corresponding message is shown below the displayed items (for example, ”100 of 100+ items are shown”). Note that the configured item limit also affects the display of configured groups; for example, if item limit is set to 100 and items are grouped by their hosts (each containing 200 items), only the first host with its 100 items will be displayed in the widget. This parameter is not affected by the *Limit for search and filter results* parameter in *Administration > General > GUI*.

Item value

Overview

The *Item value* widget displays the value of a single numeric or string item. It is useful for keeping an eye on important metrics, visualizing thresholds, and detecting sudden changes in the data.



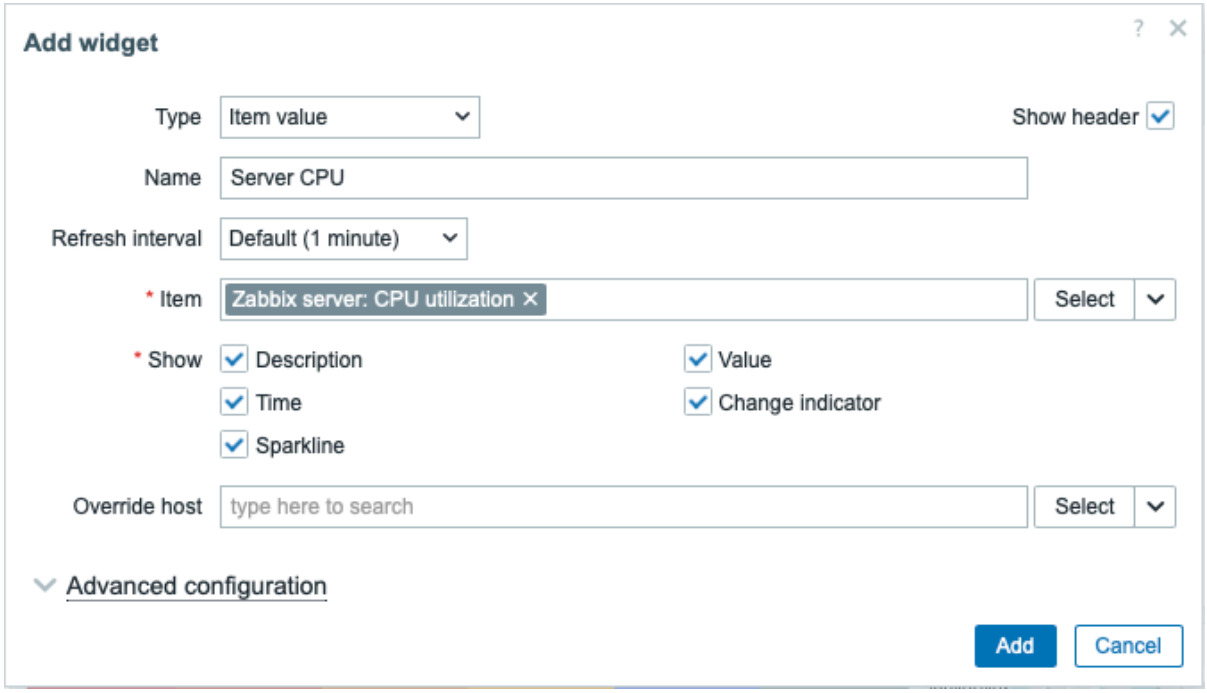
You can configure the widget to display the following:

- Metric timestamp (2024-11-21 09:57:30)
- Item description (Load average (1m), Agent status)
- Item value, change indicator, and units (0.47 ↓, 99 ↑ Mbps)
- Background color
- Sparkline chart for values from the specified time period

Clicking the widget opens an ad-hoc graph for numeric items or latest data for string items.

Configuration

To configure, select *Item value* as the widget type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

Item	Select the item. Alternatively, select a compatible widget as the data source for the item. Items returning binary data are not supported.
Show	Mark the checkbox to display the respective element: description, value, time, change indicator, sparkline. Unmark to hide. At least one element must be selected.
Override host	Select a compatible widget or the dashboard host selector as the data source for hosts. This parameter is not available when configuring the widget on a template dashboard.
Advanced configuration	Click the *Advanced configuration* header to display advanced configuration options.

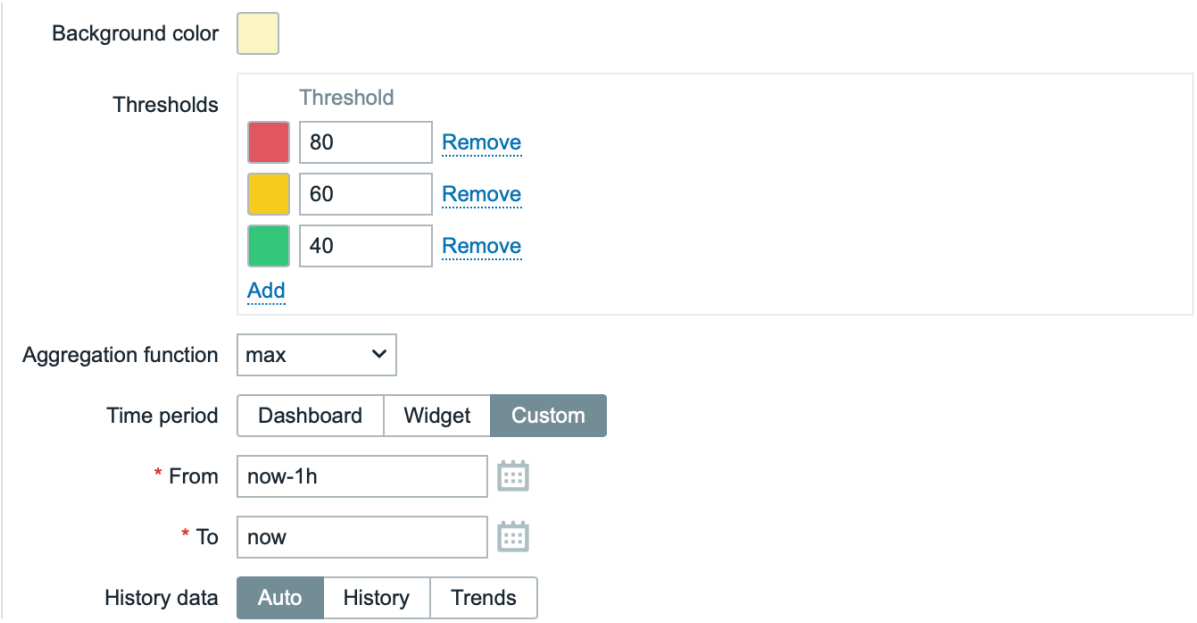
Advanced configuration

Advanced configuration options are available in the collapsible *Advanced configuration* section, and only for those elements that are selected in the *Show* field (see above).

Note that multiple elements cannot occupy the same space; if they are placed in the same space, an error message will be displayed.

Generic widget parameters

These parameters determine the background color (static or dynamic) for the whole widget and an aggregation function for displaying values.

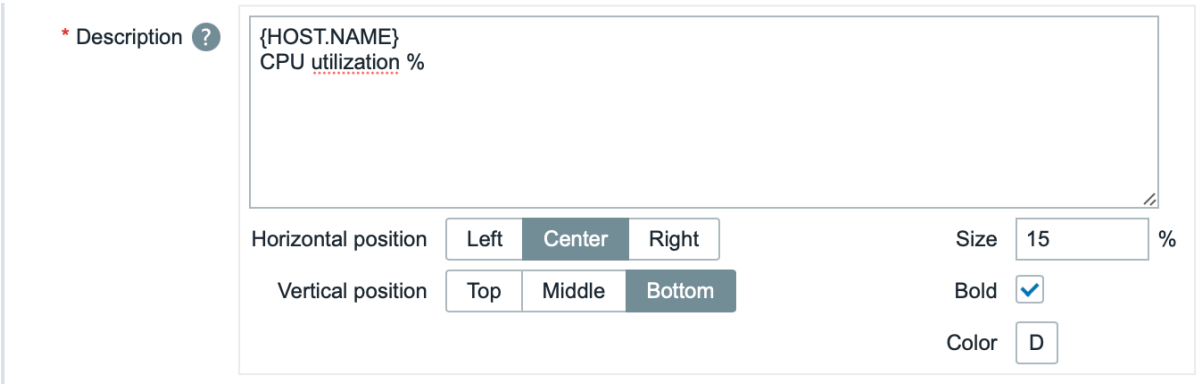


Background color	Select the background color for the whole widget from the color picker. D stands for default color (depends on the frontend theme). To return to the default value, click the *Use default* button in the color picker.
Thresholds	Configure the dynamic background color for the whole widget. Click *Add* to add a threshold, select the background color from the color picker, and specify a numeric value. Once the item value equals or is greater than the threshold value, the background color will change. The list will be sorted in ascending order when saved. Note that the dynamic background color will be displayed correctly only for numeric items.
Aggregation function	Specify which aggregation function to use within the selected *Time period*: **min** - display the smallest value; **max** - display the largest value; **avg** - display the average value; **count** - display the count of values; **sum** - display the sum of values; **first** - display the first value; **last** - display the last value. If **not used** is selected, the most recent value is displayed (no aggregation). Aggregation allows an aggregated value to be displayed for the specified interval (5 minutes, an hour, a day), instead of the most recent value. Only numeric data can be displayed for *min*, *max*, *avg* and *sum*. For *count*, non-numeric data will be changed to numeric.
Time period	Time period to use for aggregating values. Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period. This parameter is hidden when *Aggregation function* is set to "not used".
Widget	Enter or select a compatible widget as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".

History data	Take data from history or trends: **Auto** - automatic selection; **History** - take history data; **Trends** - take trend data. This setting applies only to numeric data. Non-numeric data will always be taken from history.

Description

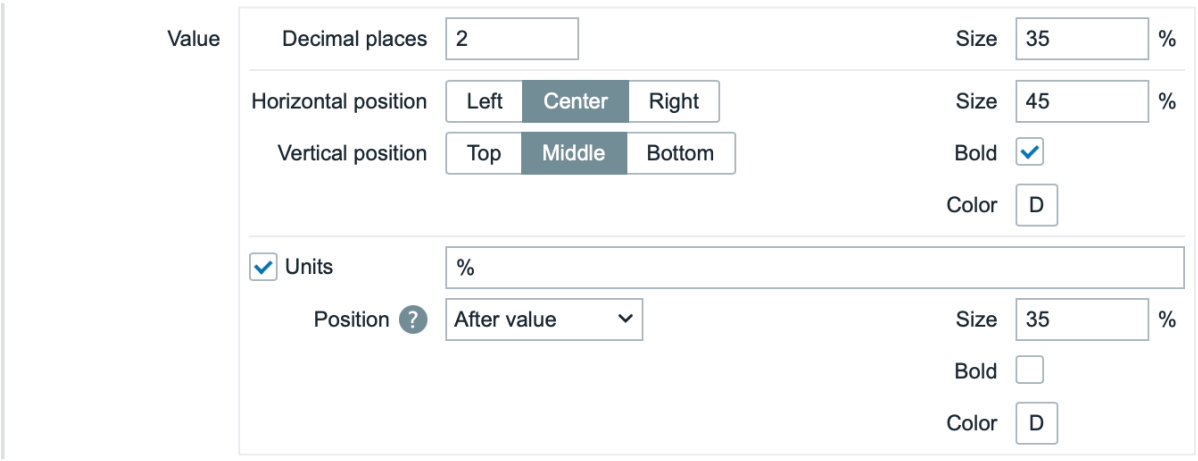
These parameters determine how the item description should be displayed.



Description	Enter the item description. This description may override the default item name. Multiline descriptions are supported. A combination of text and supported macros is possible. {HOST.*}, {ITEM.*}, {INVENTORY.*} and user macros are supported.
Horizontal position	Select horizontal position of the item description - left, right or center.
Vertical position	Select vertical position of the item description - top, bottom or middle.
Size	Enter font size height for the item description (in percent relative to total widget height).
Bold	Mark the checkbox to display item description in bold type.
Color	Select the item description color from the color picker. D stands for default color (depends on the frontend theme). To return to the default value, click the *Use default* button in the color picker.

Value

These parameters determine how the item value should be displayed.

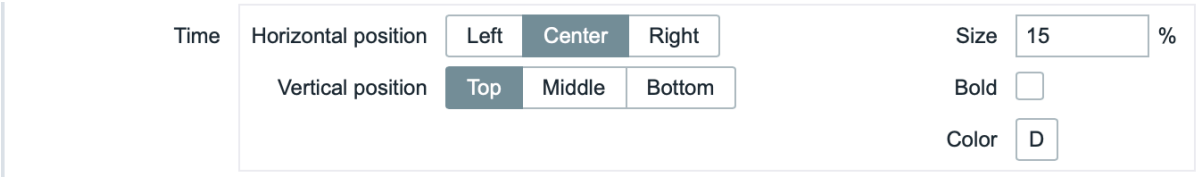


Decimal places	Select how many decimal places will be displayed with the value. This value will affect only float items. For items with units set to "s", if this value is greater than 0 (the default is 2), the widget rounds the most significant time unit and displays a numeric value (e.g., "10.43m") instead of converting the value into a full time string. When set to 0, the seconds are converted to a human-readable time string (for example, "4h 56m 30s").
Size	Enter font size height for the decimal places (in percent relative to total widget height).
Horizontal position	Select horizontal position of the item value - left, right or center.
Vertical position	Select vertical position of the item value - top, bottom or middle.

Size	Enter font size height for the item value (in percent relative to total widget height). Note that the size of item value is prioritized; other elements have to concede space for the value. With the change indicator though, if the value is too large, it will be truncated to show the change indicator.
Bold	Mark the checkbox to display item value in bold type.
Color	Select the item value color from the color picker. D stands for default color (depends on the frontend theme). To return to the default value, click the *Use default* button in the color picker.
Units	
Units	Mark the checkbox to display units with the item value. If you enter a unit name, it will override the unit from item configuration.
Position	Select the item unit position - above, below, before or after the value.
Size	Enter font size height for the item unit (in percent relative to total widget height).
Bold	Mark the checkbox to display item unit in bold type.
Color	Select the item unit color from the color picker. D stands for default color (depends on the frontend theme). To return to the default value, click the *Use default* button in the color picker.

Time

These parameters determine how the time (clock value from the item history) should be displayed.



Horizontal position	Select horizontal position of the time - left, right or center.
Vertical position	Select vertical position of the time - top, bottom or middle.
Size	Enter font size height for the time (in percent relative to total widget height).
Bold	Mark the checkbox to display time in bold type.
Color	Select the time color from the color picker. D stands for default color (depends on the frontend theme). To return to the default value, click the *Use default* button in the color picker.

Change indicator

The section allow you to select the color of change indicators from the color picker.



The change indicators are as follows:

- ↑ - item value is up (for numeric items)
- ↓ - item value is down (for numeric items)
- ↕ - item value has changed (for string items and items with value mapping)

Note:

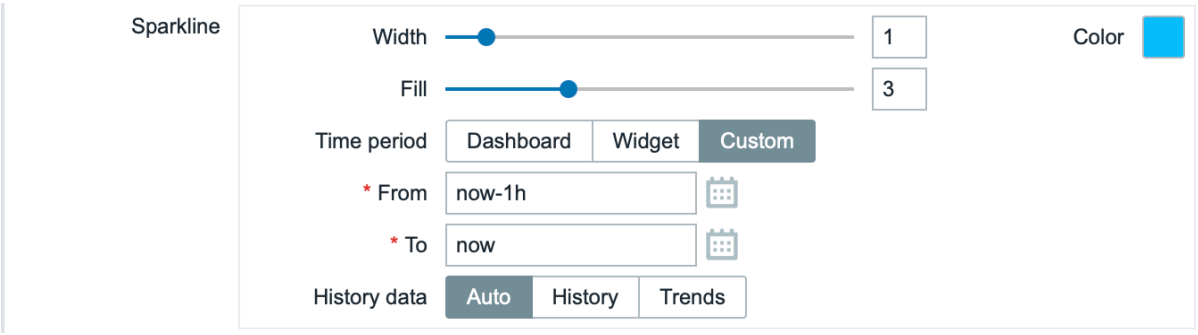
- Up (↑) and down (↓) indicators are not shown if there's only one value.
- Values are always compared within the same scope; for example:
  - The latest value is compared to the previous value.
  - A monthly value is compared to the previous month.
  - For aggregations, the previous period has the same duration as the selected period, ending immediately before the selected period begins.

D stands for default color (depends on the frontend theme). To return to the default value, click the *Use default* button in the color picker.

Vertical size of the change indicator is equal to the size of the value (integer part of the value for numeric items).

Sparkline

These parameters determine how the sparkline chart should be displayed.

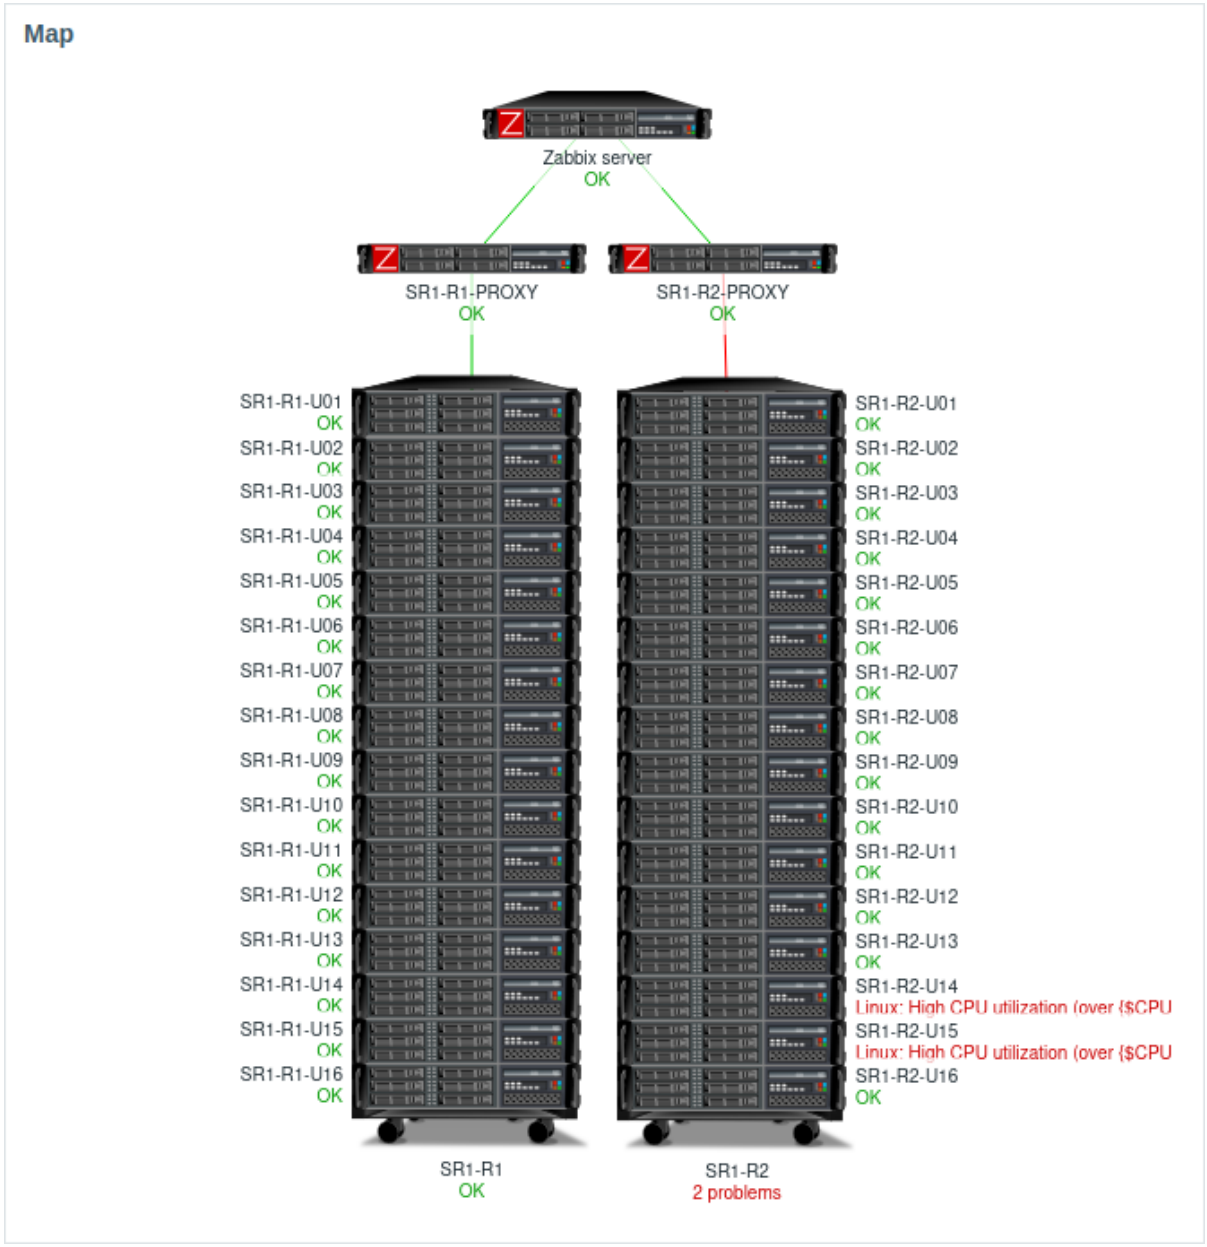


Width	Set the graph line thickness by using the slider or manually entering a value in the range from 0 to 10.
Color	Select line and fill color.
Fill	Set fill color transparency level by using the slider or manually entering a value in the range from 0 to 10.
Time period	Specify the time period for values to be included into the sparkline chart. Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period. This parameter is not available if *Aggregation function* is set to "not used".
Widget	Enter or select a compatible widget as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
History data	Take data from history or trends: **Auto** - automatic selection; **History** - take history data; **Trends** - take trend data.

Map

Overview

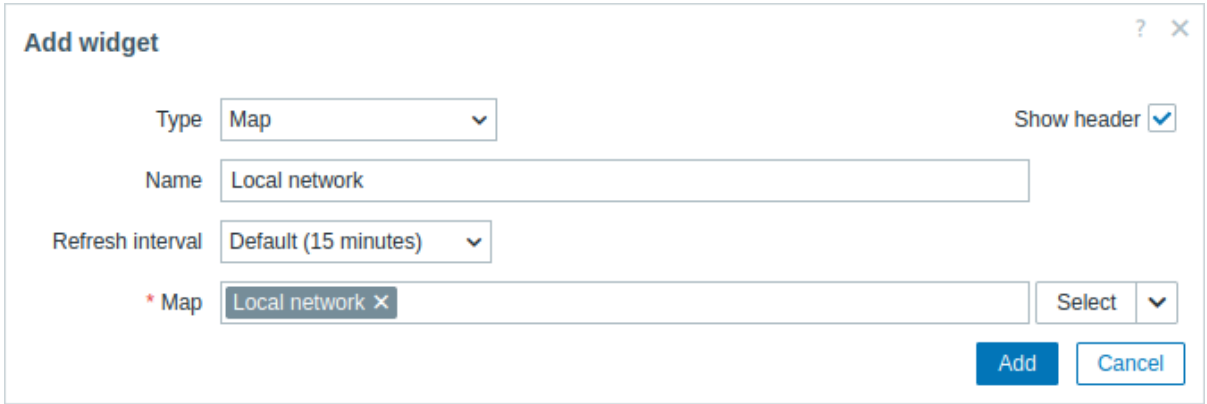
The *Map* widget can display a network map or a similar visualization, providing you a dynamic overview of your network.



You can pair it with the Map navigation tree widget to dynamically display the map selected in the tree.

Configuration

To configure, select *Map* as type:



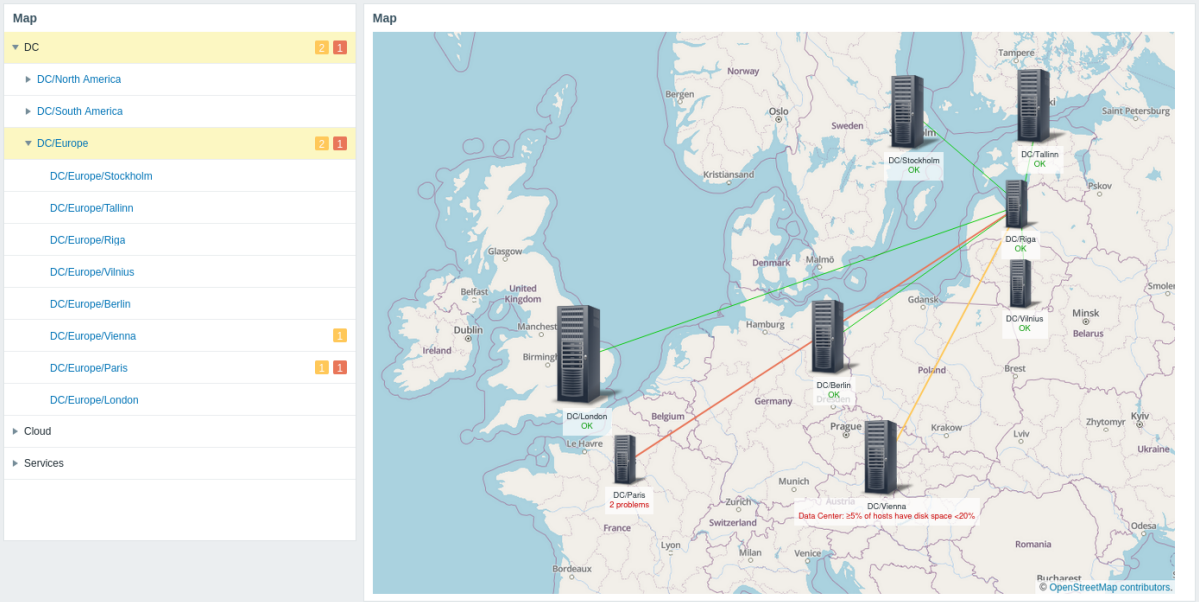
In addition to the parameters that are common for all widgets, you may set the following specific options:

Map	Set a map to display. Alternatively, select a compatible widget as the data source for the map to display. This field is auto-complete, so starting to type the name of the map or widget will offer a dropdown of matching maps or widgets.

Map navigation tree

Overview

The *Map navigation tree* widget displays a hierarchy of existing maps and the problem count for each map and map group.

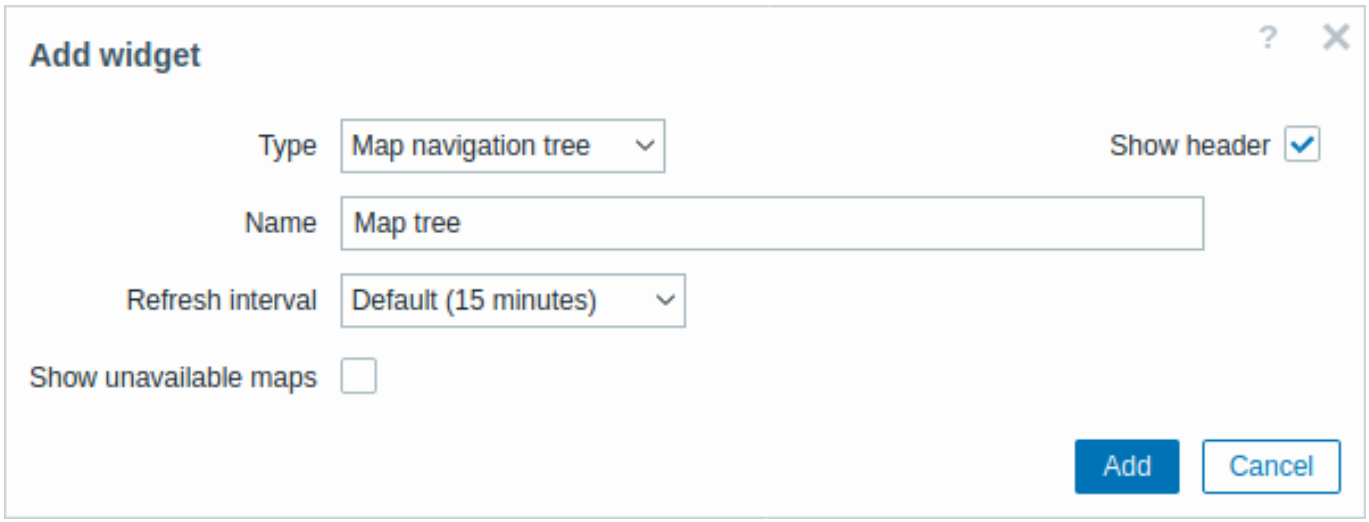


You can pair it with the Map widget to dynamically display the map selected in the tree.

The problem count for the top-level map in the hierarchy includes all submap problems and its own.

Configuration

To configure, select *Map navigation tree* as type:

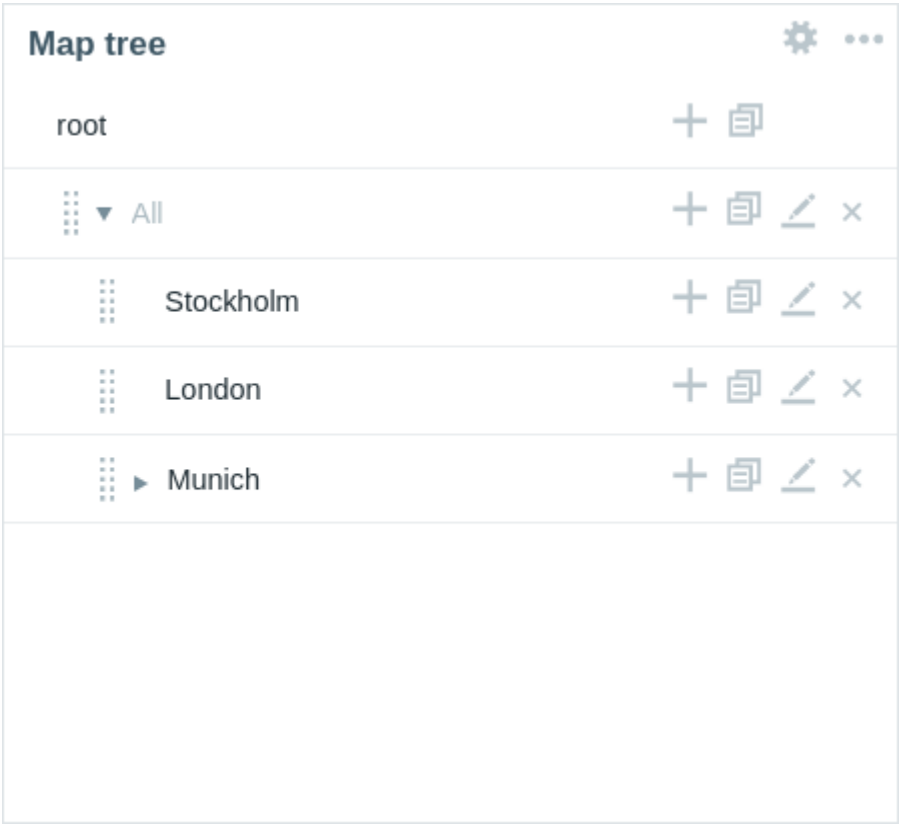


In addition to the parameters that are common for all widgets, you may set the following specific options:

Show unavailable maps	Mark this checkbox to display maps that the user does not have read permission to. Unavailable maps in the navigation tree will be displayed with a grayed-out icon. Note that if this checkbox is marked, available submaps are displayed even if the parent level map is unavailable. If unmarked, available submaps to an unavailable parent map will not be displayed at all. Problem count is calculated based on available maps and available map elements.

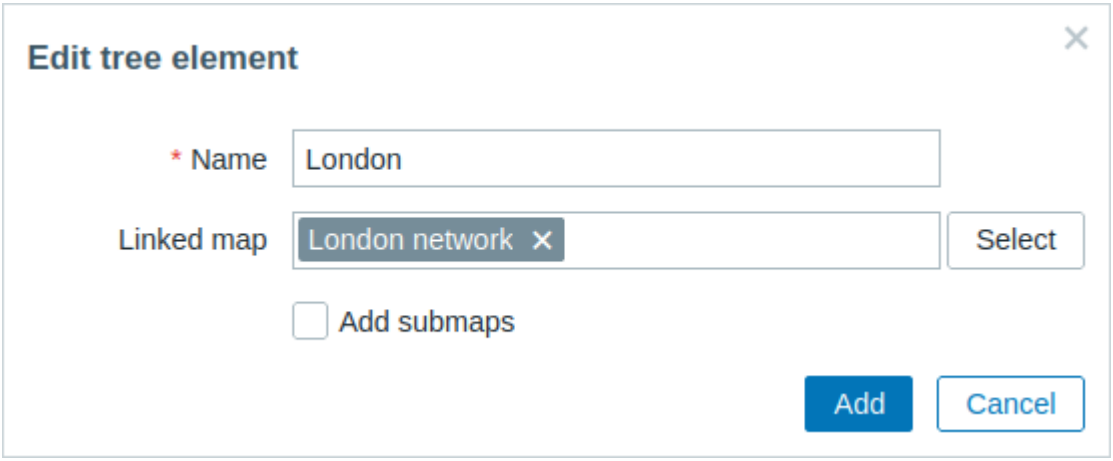
Navigation tree elements are displayed in a list. You can:

- drag an element (including its child elements) to a new place in the list;
- expand or collapse an element to display or hide its child elements;
- add a child element (with or without a linked map) to an element;
- add multiple child elements (with linked maps) to an element;
- edit an element;
- remove an element (including its child elements).



Element configuration

To configure a navigation tree element, either add a new element or edit an existing element.



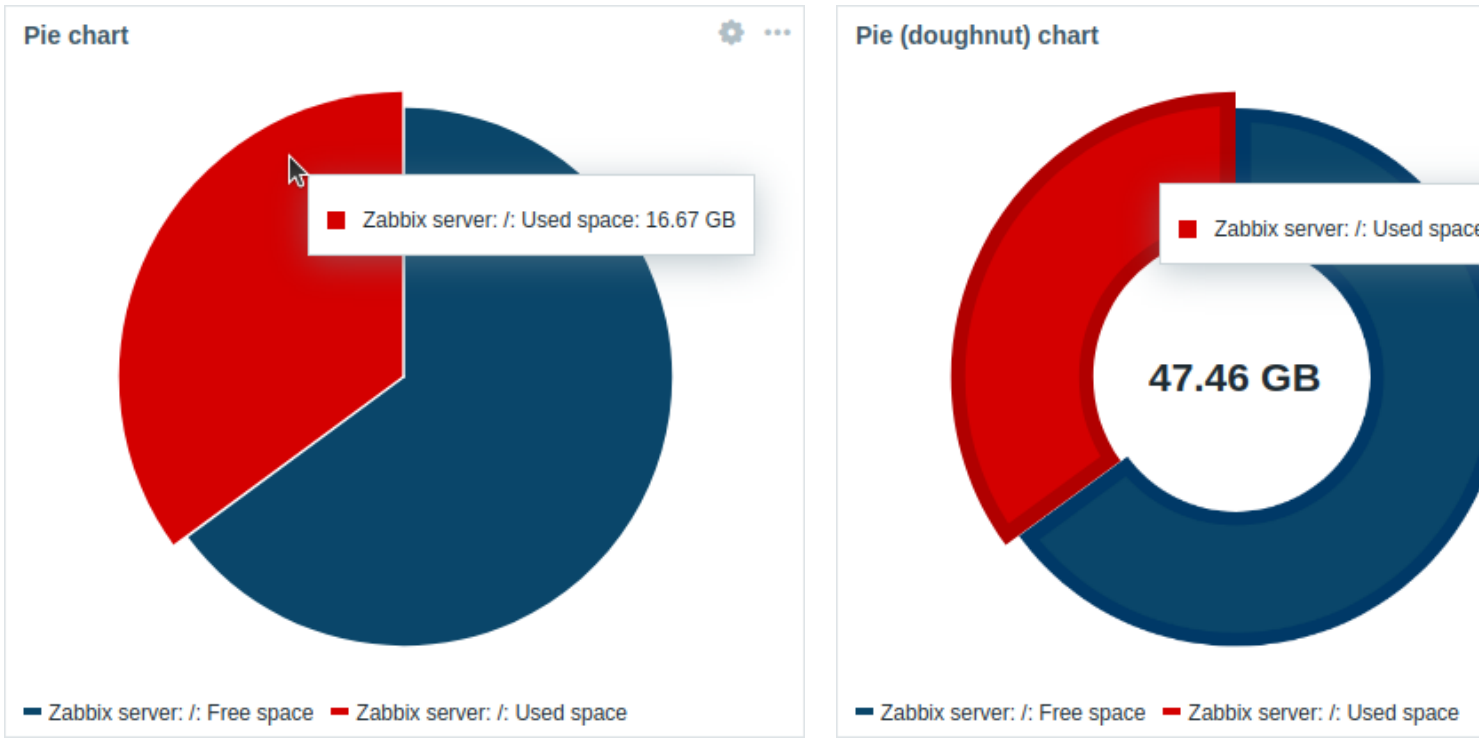
The following navigation tree element configuration parameters are available:

Name	Enter the navigation tree element name.
Linked map	Select the map to link to the navigation tree element. This field is auto-complete, so starting to type the name of a map will offer a dropdown of matching maps.
Add submaps	Mark this checkbox to add the submaps of the linked map as child elements to the navigation tree element.

Pie chart

Overview

The *Pie chart* widget displays numeric item data as a vector-based pie or doughnut chart. This widget can help you visualize how items or hosts contribute to the overall dataset.

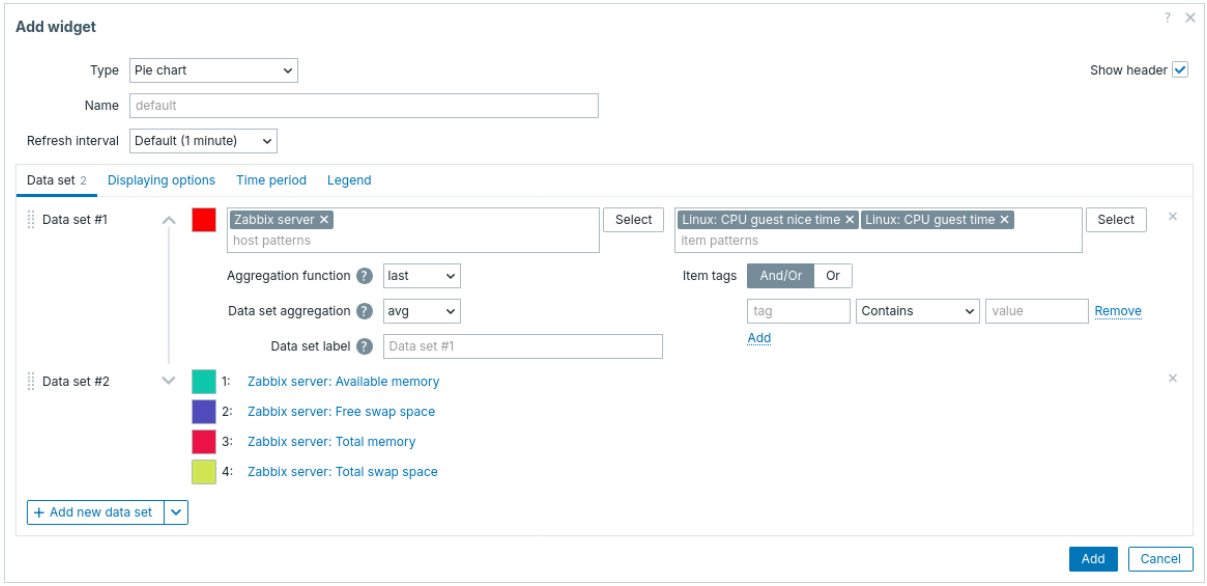


On mouseover, the hovered sector expands and shows its legend; clicking the sector keeps it expanded and its legend visible until closed.

The information displayed in the *Pie chart* widget can be downloaded as a PNG image by selecting the *Download image* option in the widget *Actions* menu ⋯ that appears when you hover over the widget. Depending on your browser settings, the image is either automatically saved to the `Downloads` folder, or a *Save* dialog opens, allowing you to choose the location.

Configuration

To configure, select *Pie chart* as type:



Data set

The **Data set** tab allows selecting data for the pie chart by adding data sets. Two types of data sets can be added:

- *Item patterns* - data from matching items is displayed. You can pick a single base color or select a palette row to assign distinct colors to each matched item.
- *Item list* - data from selected items is displayed. You can choose each item's color individually from the picker.

By default, an *Item patterns* data set is added.

Data set	For **Item patterns** data set: Select or enter host and item patterns; data of items that match the entered patterns will be displayed on the pie chart; up to 50 items may be displayed. Wildcard patterns may be used for selection (for example, ∗ will return results that match zero or more characters). To specify a wildcard pattern, enter the string manually and press *Enter*. The wildcard symbol is always interpreted, so it is not possible to add, for example, an item named *item** individually if there are other matching items (for example, *item2*, *item3*). Specifying host and item patterns is mandatory for "Item patterns" data sets. See also: Data set configuration details. For **Item list** data set: Select items for the pie chart by clicking the *Add item* button. The *Type* dropdown after the item name allows you to select the display type for each item: **Normal** - item value is represented proportionally on the pie chart (default); **Total** - item value takes up the whole pie chart. Note that only one "Total" item can exist per pie chart, and it will be placed first on the pie chart legend. If an item is set to "Total", the *Data set aggregation* parameter (see below) will be disabled and set to "not used". You may also select compatible widgets as the data source for items by clicking the *Add widget* button. The same options apply to widgets as to individual items. Specifying items or widgets is mandatory for "Item list" data sets. See also: Data set configuration details. Note that only numeric item types are allowed. When configuring the widget on a template dashboard, the parameter for specifying host patterns is not available, and the parameter for specifying an item list allows you to select only the items configured on the template.

Aggregation function	Specify which aggregation function to use for each item in the data set within the selected *Time period*: **min** - display the smallest value; **max** - display the largest value; **avg** - display the average value; **sum** - display the sum of values; **count** - display the count of values; **first** - display the first value; **last** - display the last value (default). Aggregation allows an aggregated value to be displayed for the specified interval (5 minutes, an hour, a day). Only numeric data can be displayed for *min*, *max*, *avg* and *sum*. For *count*, non-numeric data will be changed to numeric.
Data set aggregation	Specify which aggregation function to use for the whole data set: **not used** - no aggregation, items are displayed separately (default); **min** - display the smallest value; **max** - display the largest value; **avg** - display the average value; **sum** - display the sum of values; **count** - display the count of values. Aggregation allows an aggregated value to be displayed for the specified interval (5 minutes, an hour, a day) selected in the *Time period* tab or used for the whole dashboard.
Data set label	Specify a custom label for the data set. The label is displayed in the data set configuration and the pie chart legend (for aggregated data sets). All data sets are numbered including those with a specified *Data set label*. If no label is specified, the data set will be labeled automatically according to its number (e.g., ”Data set #2”, ”Data set #3”, etc.). Data set numbering is recalculated after reordering/dragging data sets. Data set labels that are too long will be shortened to fit where displayed (e.g., ”Number of proc...”).
Item tags	Specify tags to filter the items displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met.

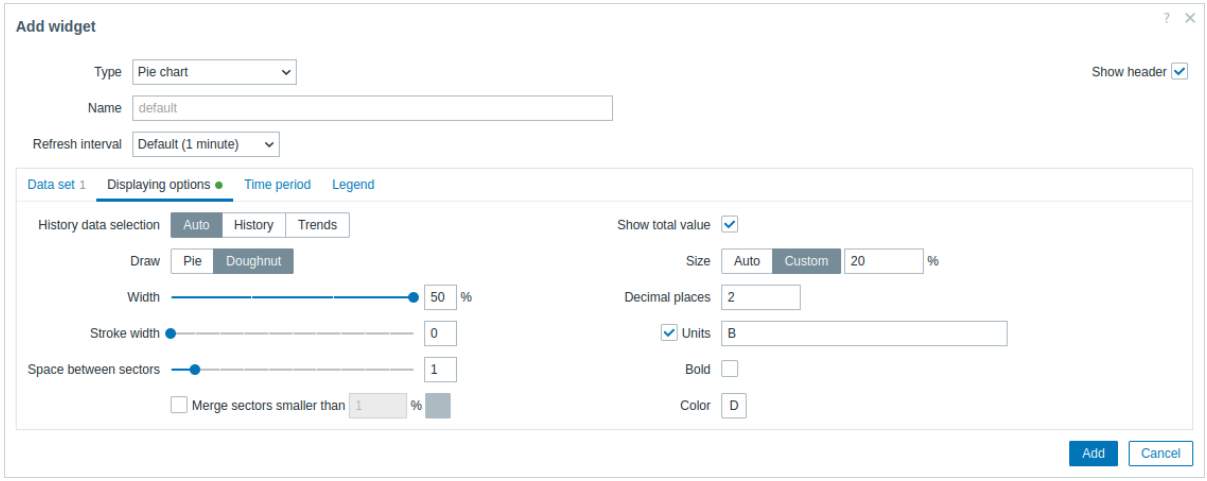
Data set configuration details

Existing data sets are displayed in a list. You can rearrange, expand/collapse, change colors, and clone these data sets.

For more information, see data set configuration details in the *Graph* widget. These details also apply to the *Pie chart* widget.

Displaying options

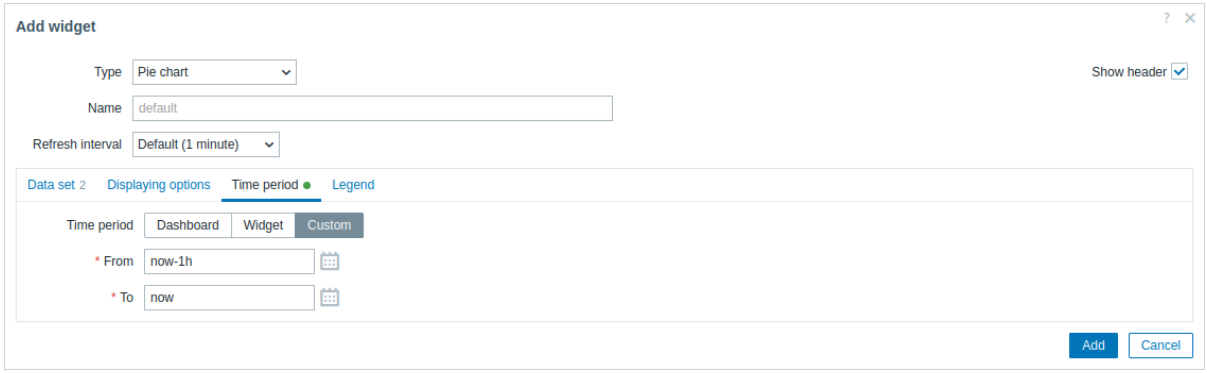
The **Displaying options** tab allows you to define history data selection and visualization options for the pie chart:



History data selection	Select the data source: **Auto** - data is sourced according to the classic algorithm (default); **History** - data from history; **Trends** - data from trends.
Draw	Select the visualization style of the pie chart: **Pie** - a full pie (sectors take up 100% of the radius); **Doughnut** - a pie with empty circle in the middle (sectors use up to 50% of radius).
Space between sectors	Select the space size (in units of 0-10) between sectors (default is ”1”).
Merge sectors smaller than N%	Mark the checkbox to merge sectors smaller than the N%. If enabled, select the color for the merged sectors and the percentage threshold (N) for merging small sectors.
Draw style: Doughnut	
Width	Select the doughnut width: 20, 30, 40 or 50% (default) of the radius.
Stroke width	Select the width of the doughnut sector border (0-10).
Show total value	Mark the checkbox to show the total value in the middle of the doughnut chart.
Size	Select the size option for the total value: **Auto** - the text is sized automatically to fit the doughnut middle readably; **Custom** - specify the text size as height percentage from the total widget height.
Decimal places	Specify the number of decimal places for the total value (0-6).
Units	Specify the units for the total value.
Bold	Mark the checkbox to display the total value in bold.
Color	Select the color for the total value.

Time period

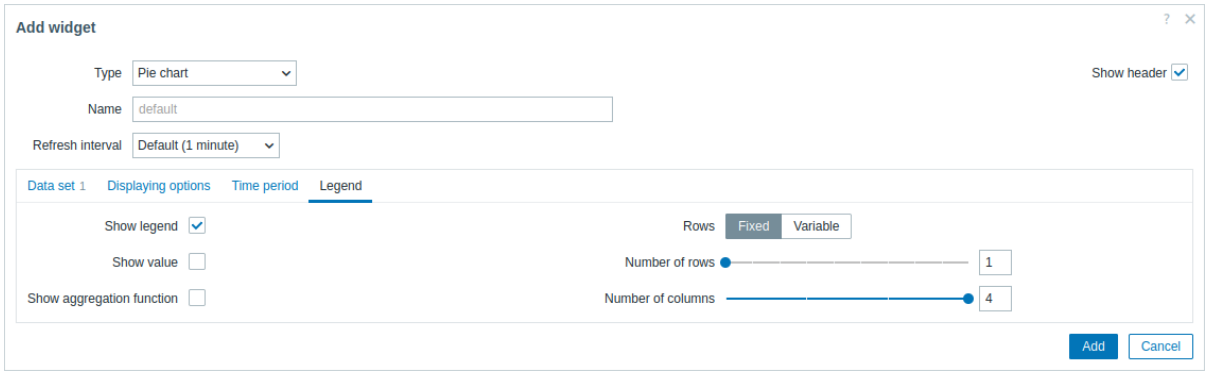
The **Time period** tab allows you to set a custom time period for the aggregation settings of the pie chart:



Time period	Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period specified in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period.
Widget	Enter or select a compatible widget (*Graph*, *Graph (classic)*, *Graph prototype*) as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".

Legend

The **Legend** tab allows you to customize the pie chart legend:

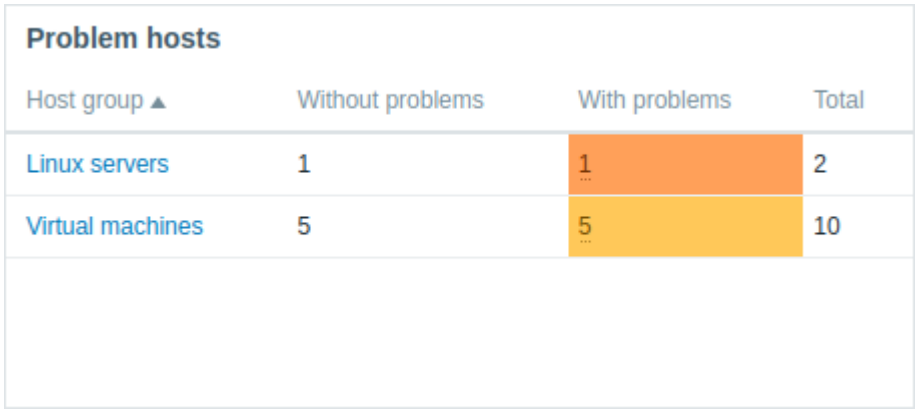


Show legend	Unmark this checkbox to hide the legend on the pie chart (marked by default).
Show value	Mark this checkbox to show the value of the item in the legend.
Show aggregation function	Mark this checkbox to show the aggregation function in the legend.
Rows	Select the display mode for legend rows: **Fixed** - the number of rows displayed is determined by the *Number of rows* parameter value; **Variable** - the number of rows displayed is determined by the amount of configured items while not exceeding the *Maximum number of rows* parameter value.
Number of rows/ Maximum number of rows	If *Rows* is set to "Fixed", set the number of legend rows to be displayed (1-10). If *Rows* is set to "Variable", set the maximum number of legend rows to be displayed (1-10).
Number of columns	Set the number of legend columns to be displayed (1-4). This parameter is available if *Show value* is unmarked.

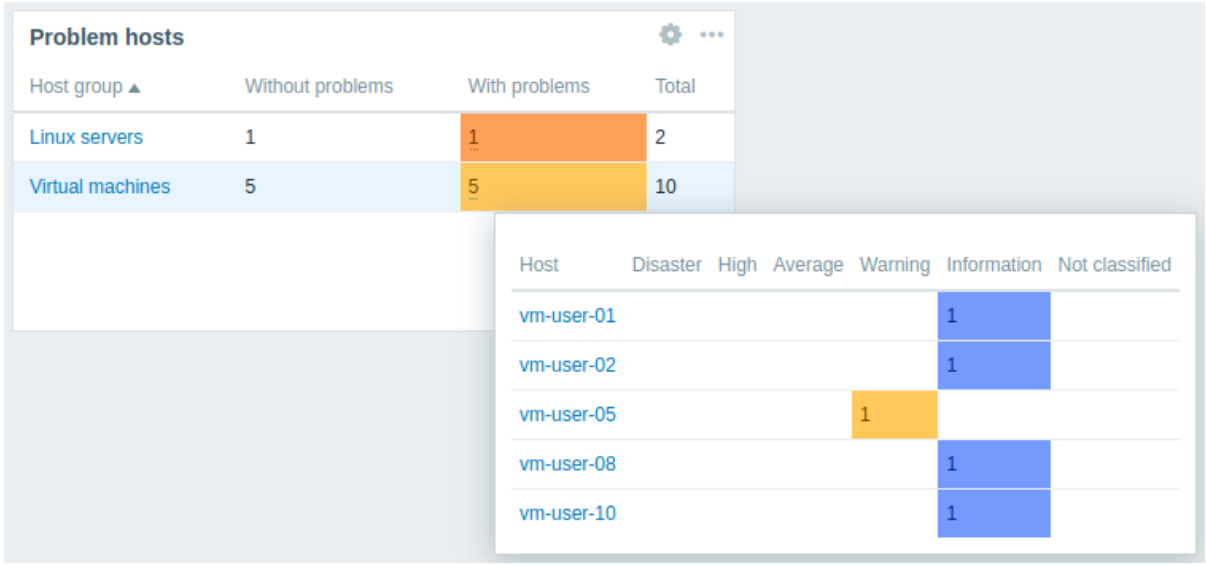
Problem hosts

Overview

The *Problem hosts* widget displays the number of cause problems per host group and shows the highest problem severity in each group. It allows filtering problems by different parameters (host groups, hosts, problem names, etc.).

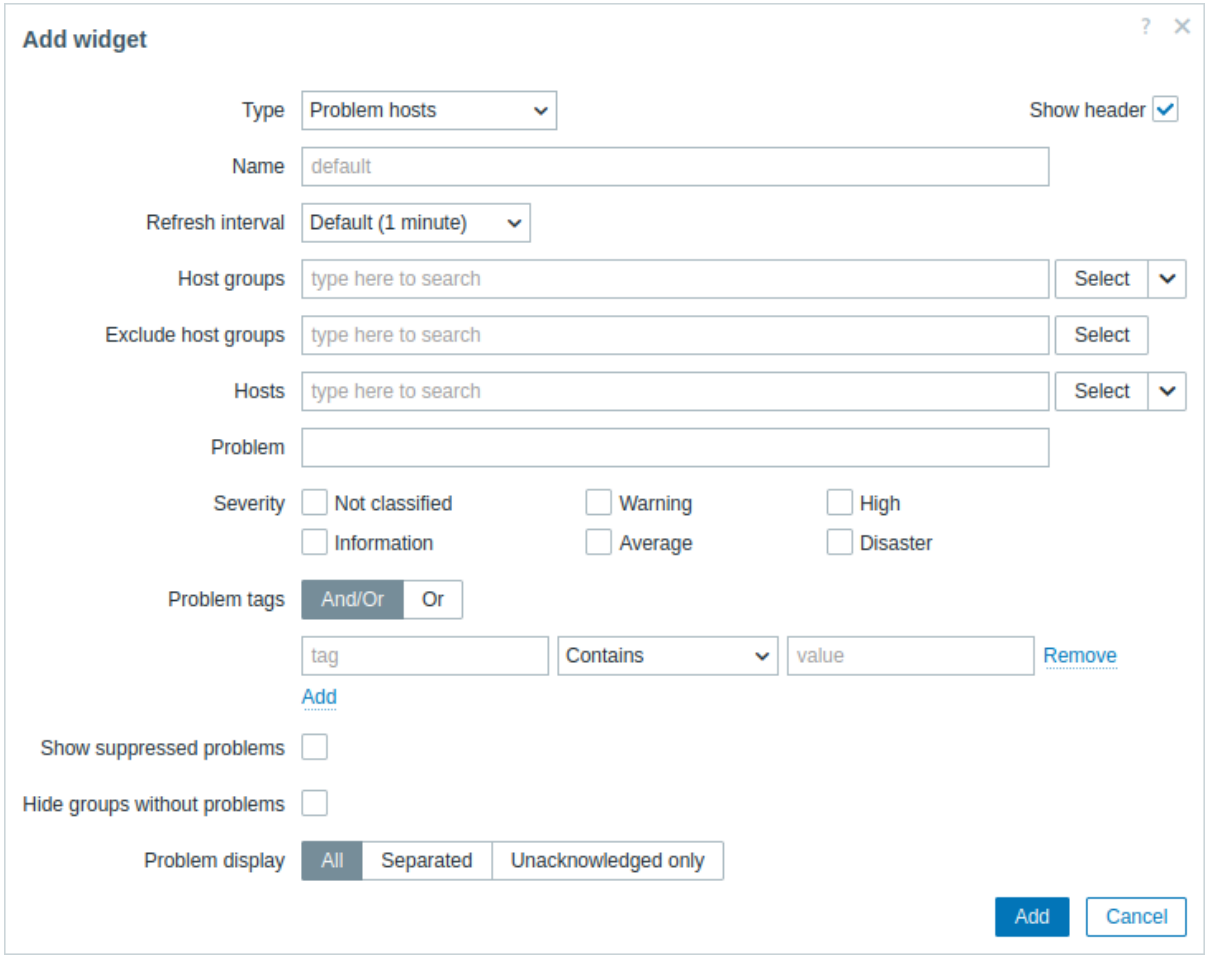


Hovering over a problem count shows the list of hosts with problems:



Configuration

To configure, select *Problem hosts* as type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

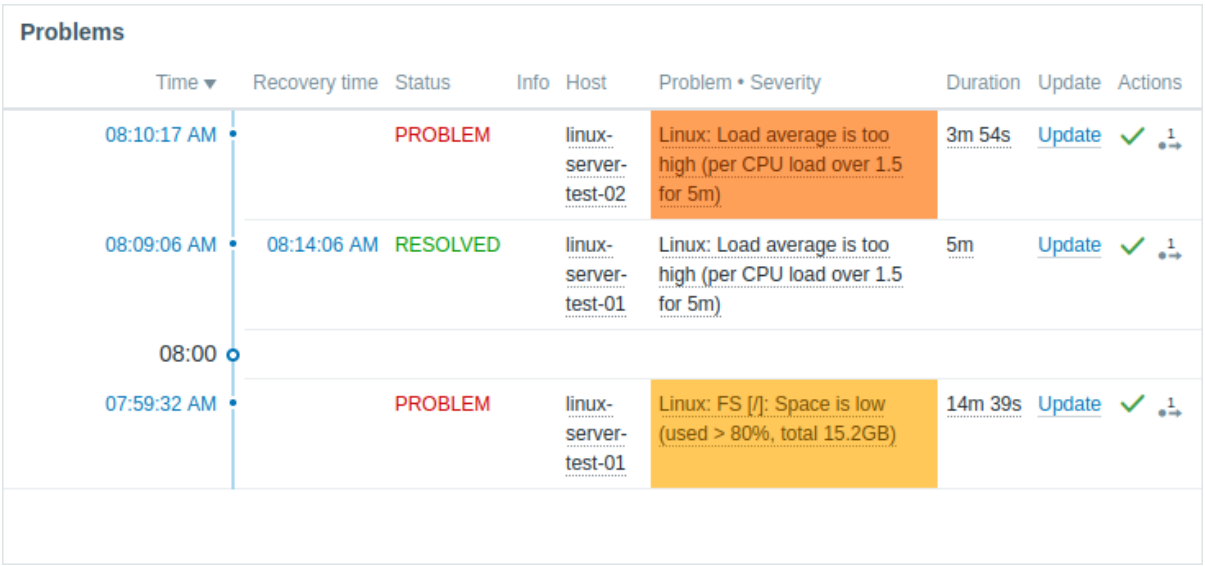
Host groups	Select host groups to display in the widget. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Specifying a parent host group implicitly selects all nested host groups. Host data from these host groups will be displayed in the widget; if no host groups are entered, all host groups will be displayed. This parameter is not available when configuring the widget on a template dashboard.
Exclude host groups	Select host groups to hide from the widget. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Specifying a parent host group implicitly selects all nested host groups. Host data from these host groups will not be displayed in the widget. For example, hosts 001, 002, 003 may be in Group A and hosts 002, 003 in Group B as well. If we select to *show* Group A and *exclude* Group B at the same time, only data from host 001 will be displayed in the dashboard. This parameter is not available when configuring the widget on a template dashboard.
Hosts	Select hosts to display in the widget. Alternatively, select a compatible widget or the dashboard as the data source for hosts. This field is auto-complete, so starting to type the name of a host will offer a dropdown of matching hosts. If no hosts are entered, all hosts will be displayed. This parameter is not available when configuring the widget on a template dashboard.
Problem	You can limit the number of problem hosts displayed by the problem name. If you enter a string here, only those hosts with problems whose name contains the entered string will be displayed. Macros are not expanded.
Severity	Mark problem severities to filter problems to be displayed in the widget. If no severities are marked, all problems will be displayed.

Problem tags	Specify problem tags to limit the number of problems displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the *Or* condition; **Or** - enough if one condition is met.
Show suppressed problems	Mark the checkbox to display problems that would otherwise be suppressed (not shown) because of host maintenance.
Hide groups without problems	Mark the *Hide groups without problems* option to hide data from host groups without problems in the widget. This parameter is not available when configuring the widget on a template dashboard.
Problem display	Display problem count as: **All** - full problem count will be displayed; **Separated** - unacknowledged problem count will be displayed separated as a number of the total problem count; **Unacknowledged only** - only the unacknowledged problem count will be displayed.

Problems

Overview

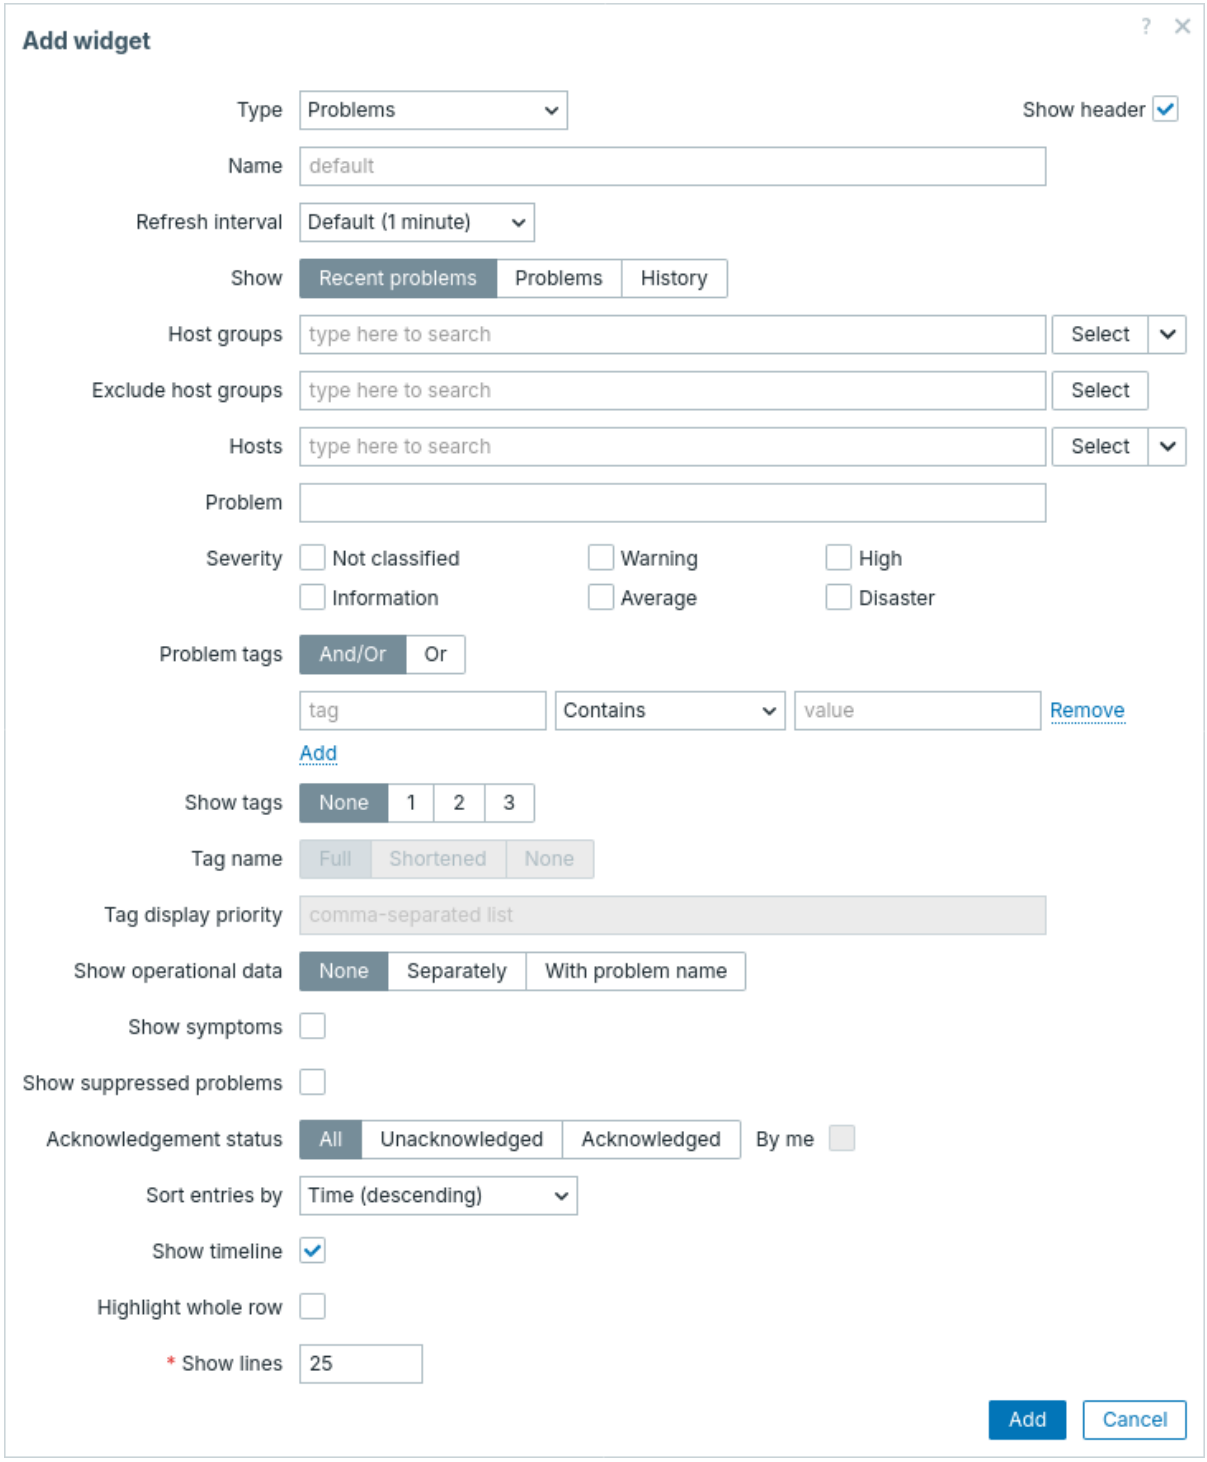
The *Problems* widget displays current problems filtered by different parameters (host groups, hosts, problem names, etc.), giving you a clear view of what needs attention.



It displays the same data as *Monitoring* > *Problems* and can display up to 1000 records.

Configuration

To configure, select *Problems* as type:



You can limit how many problems are displayed in the widget in various ways - by problem status, problem name, severity, host group, host, event tag, acknowledgment status, etc.

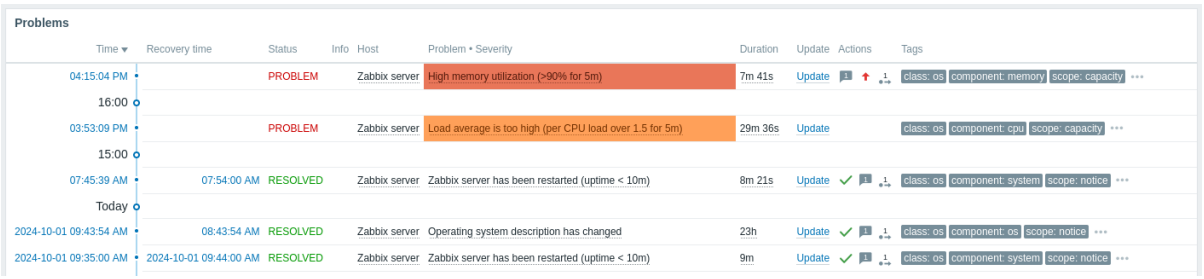
In addition to the parameters that are common for all widgets, you may set the following specific options:

Show	Filter by problem status: **Recent problems** - unresolved and recently resolved problems are displayed (default); **Problems** - unresolved problems are displayed; **History** - history of all events is displayed.
Host groups	Select host groups to display problems of in the widget. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Specifying a parent host group implicitly selects all nested host groups. Problems from these host groups will be displayed in the widget; if no host groups are entered, problems from all host groups will be displayed. This parameter is not available when configuring the widget on a template dashboard.

Exclude host groups	Select host groups to hide problems of from the widget. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Specifying a parent host group implicitly selects all nested host groups. Problems from these host groups will not be displayed in the widget. For example, hosts 001, 002, 003 may be in Group A and hosts 002, 003 in Group B as well. If we select to *show* Group A and *exclude* Group B at the same time, only problems from host 001 will be displayed in the widget. This parameter is not available when configuring the widget on a template dashboard.
Hosts	Select hosts to display problems of in the widget. Alternatively, select a compatible widget or the dashboard as the data source for hosts. This field is auto-complete, so starting to type the name of a host will offer a dropdown of matching hosts. If no hosts are entered, problems of all hosts will be displayed. This parameter is not available when configuring the widget on a template dashboard.
Problem	You can limit the number of problems displayed by their name. If you enter a string here, only those problems whose name contains the entered string will be displayed. Macros are not expanded.
Severity	Mark problem severities to filter problems to be displayed in the widget. If no severities are marked, all problems will be displayed.
Problem tags	Specify problem tags to limit the number of problems displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the *Or* condition; **Or** - enough if one condition is met. When filtered, the tags specified here will be displayed first with the problem, unless overridden by the *Tag display priority* (see below) list.
Show tags	Select the number of displayed tags: **None** - no *Tags* column; **1** - *Tags* column contains one tag; **2** - *Tags* column contains two tags; **3** - *Tags* column contains three tags. To see all tags for the problem roll your mouse over the three dots icon.
Tag name	Select tag name display mode: **Full** - tag names and values are displayed in full; **Shortened** - tag names are shortened to 3 symbols, but tag values are displayed in full; **None** - only tag values are displayed; no names.
Tag display priority	Enter tag display priority for a problem, as a comma-separated list of tags. Only tag names should be used, no values. Example: `Services,Applications,Application` The tags of this list will always be displayed first, overriding the natural ordering by alphabet.
Show operational data	Select the mode for displaying operational data: **None** - no operational data is displayed; **Separately** - operational data is displayed in a separate column; **With problem name** - append operational data to the problem name, using parentheses for the operational data. For details, see the description of the *Operational data* field in trigger configuration.

Show symptoms	Mark the checkbox to display in its own line problems classified as symptoms.
Show suppressed problems	Mark the checkbox to display problems that would otherwise be suppressed (not shown) because of host maintenance or single problem suppression.
Acknowledgement status	Filter to display all problems, unacknowledged problems only, or acknowledged problems only. Mark the additional checkbox to filter out those problems ever acknowledged by you.
Sort entries by	Sort entries by: **Time** (descending or ascending); **Severity** (descending or ascending); **Problem name** (descending or ascending); **Host** (descending or ascending). Sorting entries by **Host** (descending or ascending) is not available when configuring the widget on a template dashboard.
Show timeline	Mark the checkbox to display a visual timeline.
Highlight whole row	Mark the checkbox to highlight the full line for unresolved problems. The problem severity color is used for highlighting. *Highlight whole row* is not available in the high-contrast themes.
Show lines	Specify the number of problem lines to display.

Using the widget

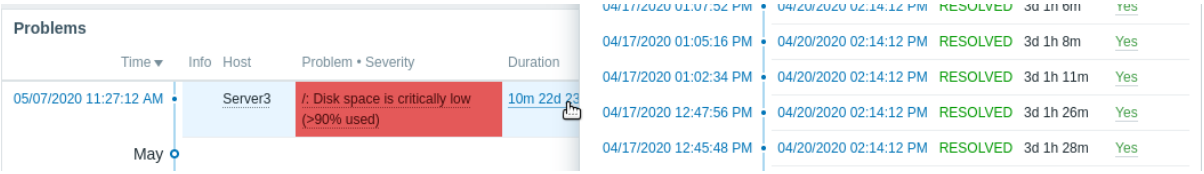


Problems widget offers quick access to additional information:

- Click the problem date (in the Time column) or the recovery date (in the Recovery Time column) to view event details.
- If the Info field is not empty, hover over the displayed icon to view additional details.
- Click the host name to open the host menu.
- Click the problem name to open the event menu.
- Hover over or click the problem duration to view problem event popup.
- Clicking *Update* opens an Update problem window.
- Hover over or click the gray arrow icon in Actions column to view list of executed actions.

Problem event popup

The problem event popup includes the list of problem events for this trigger and, if defined, the trigger description and a clickable URL.



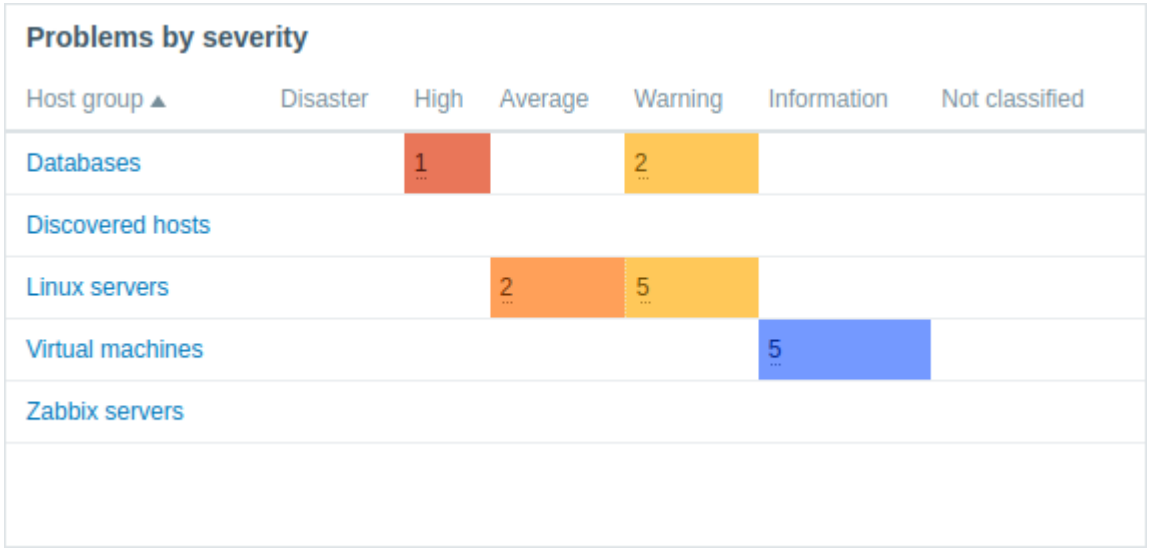
To bring up the problem event popup:

- Roll a mouse over the problem duration in the *Duration* column of the *Problems* widget. The popup disappears once you remove the mouse from the duration.
- Click the duration in the *Duration* column of the *Problems* widget. The popup disappears only if you click the duration again.

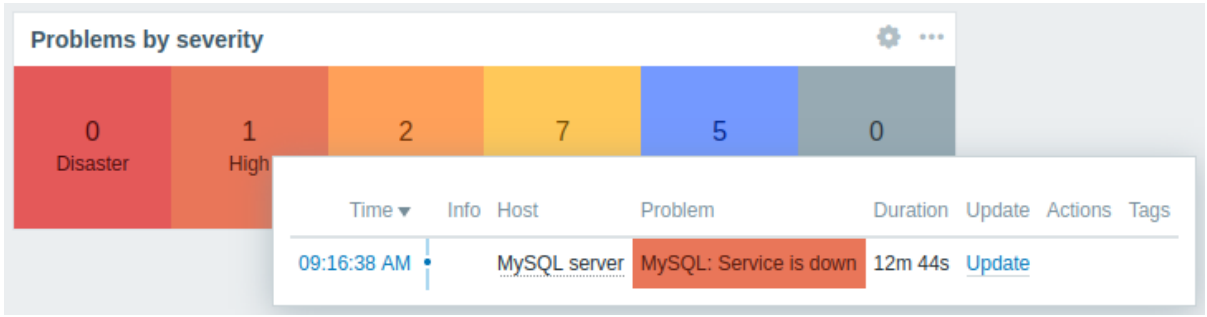
Problems by severity

Overview

The *Problems by severity* widget displays the number of cause problems grouped by severity. It allows filtering problems by different parameters (host groups, hosts, problem names, etc.).

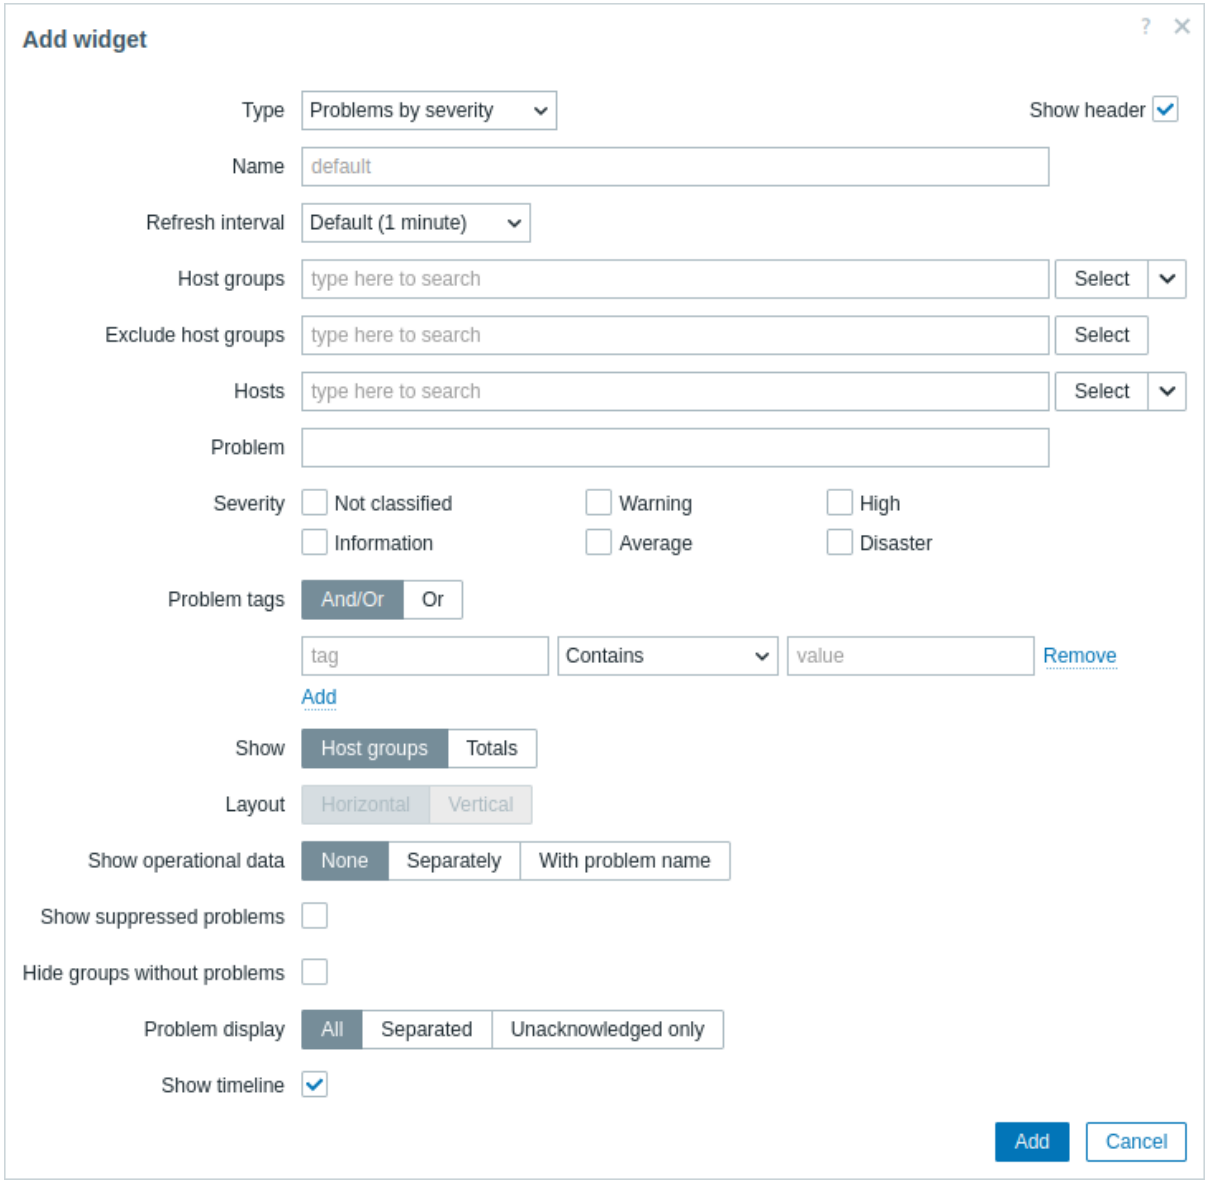


Hovering over a problem count shows the list of related problems:



Configuration

To configure, select *Problems by severity* as type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

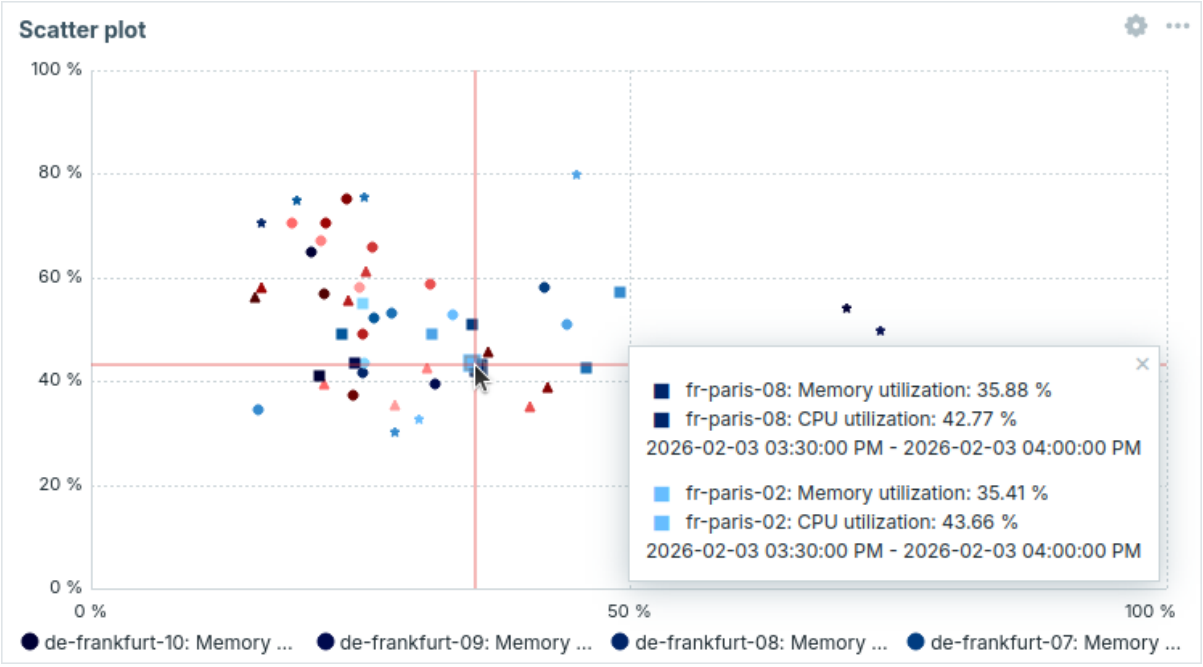
Host groups	Select host groups to display in the widget. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Specifying a parent host group implicitly selects all nested host groups. Host data from these host groups will be displayed in the widget; if no host groups are entered, all host groups will be displayed. This parameter is not available when configuring the widget on a template dashboard.
Exclude host groups	Select host groups to hide from the widget. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Specifying a parent host group implicitly selects all nested host groups. Host data from these host groups will not be displayed in the widget. For example, hosts 001, 002, 003 may be in Group A and hosts 002, 003 in Group B as well. If we select to *show* Group A and *exclude* Group B at the same time, only data from host 001 will be displayed in the dashboard. This parameter is not available when configuring the widget on a template dashboard.
Hosts	Select hosts to display in the widget. Alternatively, select a compatible widget or the dashboard as the data source for hosts. This field is auto-complete, so starting to type the name of a host will offer a dropdown of matching hosts. If no hosts are entered, all hosts will be displayed. This parameter is not available when configuring the widget on a template dashboard.

Problem	You can limit the number of problem hosts displayed by the problem name. If you enter a string here, only those hosts with problems whose name contains the entered string will be displayed. Macros are not expanded.
Severity	Mark problem severities to filter problems to be displayed in the widget. If no severities are marked, all problems will be displayed.
Problem tags	Specify problem tags to limit the number of problems displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the *Or* condition; **Or** - enough if one condition is met.
Show	Select the show option: **Host groups** - display problems per host group; **Totals** - display a problem total for all selected host groups in colored blocks corresponding to the problem severity. This parameter is not available when configuring the widget on a template dashboard, and only a problem total will be displayed.
Layout	Select the layout option: **Horizontal** - colored blocks of totals will be displayed horizontally; **Vertical** - colored blocks of totals will be displayed vertically. This parameter is available if *Show* is set to ”Totals”.
Show operational data	Display operational data as: **None** - operational data will not be displayed **Separately** - operational data will be displayed as a separate line **With problem name** - operational data will be displayed together with the problem name. For details, see the description of the *Operational data* field in trigger configuration.
Show suppressed problems	Mark the checkbox to display problems that would otherwise be suppressed (not shown) because of host maintenance.
Hide groups without problems	Mark the *Hide groups without problems* option to hide data from host groups without problems in the widget. This parameter is not available when configuring the widget on a template dashboard.
Problem display	Display problem count as: **All** - full problem count will be displayed; **Separated** - unacknowledged problem count will be displayed separated as a number of the total problem count; **Unacknowledged only** - only the unacknowledged problem count will be displayed.
Show timeline	Mark the checkbox to display a visual timeline.

Scatter plot

Overview

The *Scatter plot* widget displays the relationship between two metrics by plotting individual data points along an X and Y axis. This helps reveal patterns, clusters, correlations, and outliers in the data set.



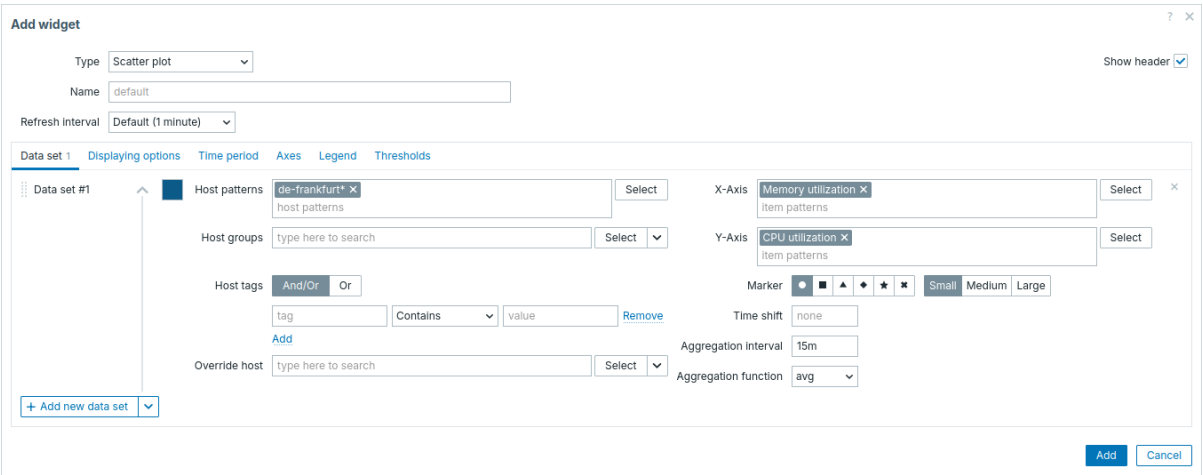
Hovering over a data point displays a tooltip with a list of items from that data point and other points within a 6px radius; sorted first by the scatter plot X-axis values in ascending order and then by Y-axis values in descending order.

Selecting an item in the tooltip broadcasts its data to other widgets.

The information displayed in the *Scatter plot* widget can be downloaded as a PNG image by selecting the *Download image* option in the widget *Actions* menu ... that appears when you hover over the widget. Depending on your browser settings, the image is either automatically saved to the `Downloads` folder, or a *Save* dialog opens, allowing you to choose the location.

Configuration

To configure, select *Scatter plot* as type:



Data set

The **Data set** tab allows selecting data for the scatter plot by adding data sets. Two types of data sets can be added:

- *Item patterns* - hosts with matching items are displayed as data points. You can pick a single base color or select a palette row to assign distinct colors to each host with the matched items.
- *Item list* - selected items are displayed as data points, where each point represents a pair of items (X and Y). You can choose each data set's color individually from the picker.

Data points are created only for those hosts/item pairs that have at least one value on both axes.

By default, an *Item patterns* data set is added.

Data set	For **Item patterns** data set: Select or enter host and item patterns; data of items that match these patterns will be displayed on the scatter plot. Note that if the number of data points exceeds 8k, `memory_limit` should be increased in the `php.ini` configuration file. The recommended memory limit for 20k data points is 512 MB. Wildcard patterns may be used for selection (for example, * will return results that match zero or more characters). To specify a wildcard pattern, enter the string manually and press *Enter*. The wildcard symbol is always interpreted, so it is not possible to add, for example, an item named *item** individually if there are other matching items (for example, *item2*, *item3*). Specifying host and item patterns is mandatory for "Item patterns" data sets. See also: Data set configuration details. For **Item list** data set: Select items for the scatter plot by clicking the *Add item* button. You may also select compatible widgets as the data source for items by clicking the *Add widget* button. Specifying items or widgets is mandatory for "Item list" data sets. See also: Data set configuration details. Note that only numeric item types are allowed. When configuring the widget on a template dashboard, the parameter for specifying host patterns is not available, and the parameter for specifying an item list allows the selection of only the items configured on the template.
Host groups	Select host groups. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Specifying a parent host group implicitly selects all nested host groups. Host data from these host groups will be displayed in the widget; if no host groups are entered, all host groups will be displayed. This parameter is available only for *Item patterns* data sets and is not available when configuring the widget on a template dashboard.
Host tags	Specify tags to filter the hosts displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met. This parameter is available only for *Item patterns* data sets and is not available when configuring the widget on a template dashboard.
Override host	Select a compatible widget or the dashboard host selector as the data source for hosts. This parameter is not available when configuring the widget on a template dashboard.
Marker	Choose the marker type and size for host data points in the data set.
Time shift	Specify time shift if required. You may use time suffixes in this field. Negative values are allowed.

Aggregation function	Specify which aggregation function to use for each item in the data set within the selected *Aggregation interval*: **min** - display the smallest value; **max** - display the largest value; **avg** - display the average value; **sum** - display the sum of values; **count** - display the count of values (if at least one value exists on both axes); **first** - display the first value; **last** - display the last value. Aggregation allows an aggregated value to be displayed for the specified interval (5 minutes, an hour, a day), instead of all values. Only numeric data can be displayed for *min*, *max*, *avg* and *sum*. For *count*, non-numeric data will be changed to numeric.
Aggregation interval	Specify the interval for aggregating values. You may use time suffixes in this field. A numeric value without a suffix will be regarded as seconds. Note that if the widget is configured to display historical data based on trends (*History data selection* is set to *Trends* or *Auto*), it is recommended to use an aggregation interval that is a multiple of 1 hour (e.g., 3600, 60m, 1h, 3h, etc.). Trends store hourly aggregated values, so using an aggregation interval that is not a multiple of 1 hour (e.g., 100s, 7min, 15min, 90min, etc.) may lead to results that are hard to interpret. Specifying an aggregation interval is mandatory. The aggregation interval determines the time interval and coordinates for a data point; for example, if X-axis items have 4 values and Y-axis items have 1 value, the coordinates reflect aggregating 4 on X and 1 on Y in the specified interval.

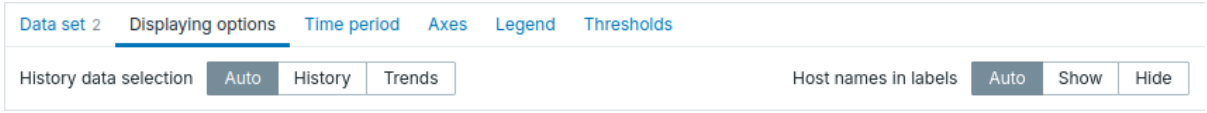
Data set configuration details

Existing data sets are displayed in a list. You can rearrange, expand/collapse, change colors, and clone these data sets.

For more information, see data set configuration details in the *Graph* widget. These details also apply to the *Scatter plot* widget.

Displaying options

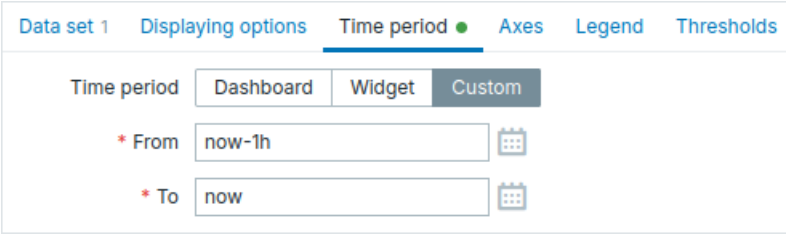
The **Displaying options** tab allows you to define history data selection:



History data selection	Set the source of scatter plot data: **Auto** - data is sourced according to the classic graph algorithm (default). **History** - data from history. **Trends** - data from trends.
Host names in labels	Choose whether to display host names in the graph legend and in the tooltip that appears when hovering over the graph data: **Auto** - host names are displayed only if more than one host is present across the data sets (default). **Show** - host names are displayed. **Hide** - host names are hidden. This parameter is not available when configuring the widget on a template dashboard.

Time period

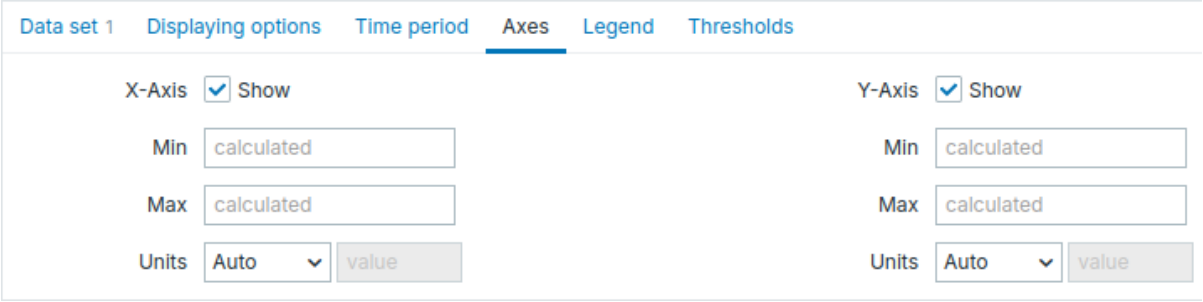
The **Time period** tab allows you to set a time period for which to display data in the scatter plot:



Time period	Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period.
Widget	Enter or select a compatible widget as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".

Axes

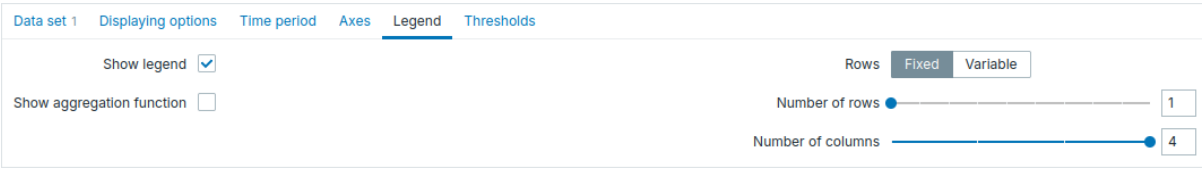
The **Axes** tab allows you to customize how axes are displayed:



X-Axis/Y-Axis	Unmark this checkbox to hide X-axis/Y-axis (marked by default).
Min	Set the minimum value of the corresponding axis.
Max	Set the maximum value of the corresponding axis.
Units	Choose the unit for the scatter plot axis values from the dropdown: **Auto** - axis values are displayed using the unit of the first item in the data set; **Static** - axis values are displayed using the unit specified in the *value* input field; if the field is left blank, only numeric values are displayed.

Legend

The **Legend** tab allows you to customize the scatter plot legend:

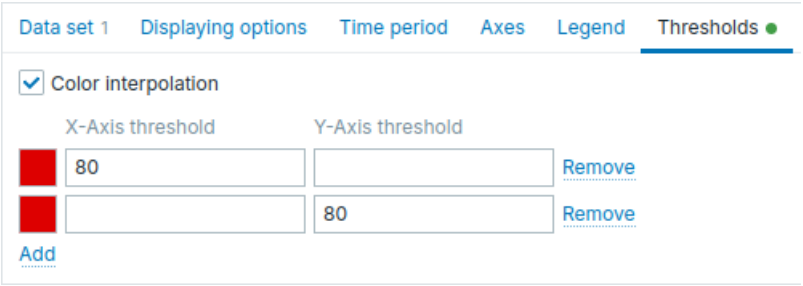


Show legend	Unmark this checkbox to hide the legend on the scatter plot (marked by default).
Show aggregation function	Mark this checkbox to show the aggregation function in the legend.

Rows	Select the display mode for legend rows: **Fixed** - the number of rows displayed is determined by the *Number of rows* parameter value; **Variable** - the number of rows displayed is determined by the amount of configured items while not exceeding the *Maximum number of rows* parameter value.
Number of rows/ Maximum number of rows	If *Rows* is set to ”Fixed”, set the number of legend rows to be displayed (1-10). If *Rows* is set to ”Variable”, set the maximum number of legend rows to be displayed (1-10).
Number of columns	Set the number of legend columns to be displayed (1-4). This parameter is available if *Display min/avg/max* is unmarked.

Thresholds

The **Thresholds** tab allows you to configure thresholds that determine host marker colors:

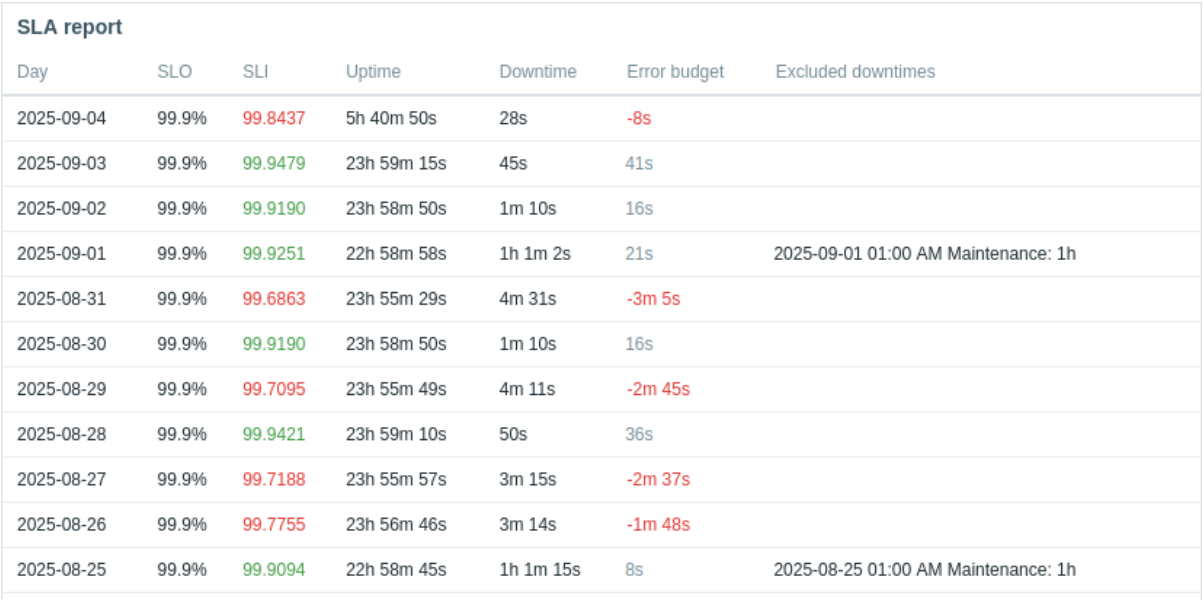


Color interpolation	Mark the checkbox to enable smooth transitioning between threshold colors for host markers. This parameter applies if two or more thresholds are set for either the X-axis, the Y-axis, or both axes simultaneously.
X-Axis/Y-Axis threshold	Click *Add* to add a threshold, select a threshold color from the color picker, and specify a numeric value. Thresholds can be applied to a single axis, both axes separately, or both axes simultaneously. For example, a single threshold with X=80 and Y=80 will color the markers when both conditions are met. The thresholds list will be sorted in ascending order when saved. Suffixes (for example, ”1d”, ”2w”, ”4K”, ”8G”) are supported. Value mappings are supported.

SLA report

Overview

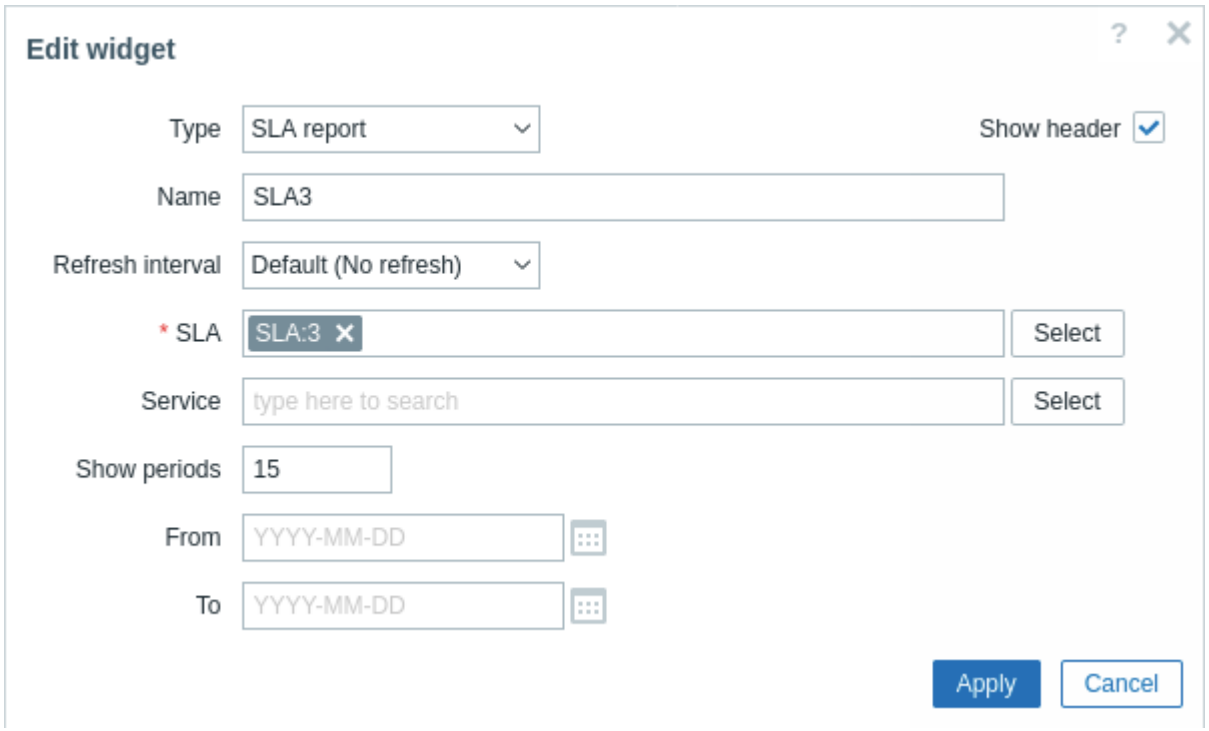
The *SLA report* widget displays SLA reports, helping you monitor service performance against defined targets.



It displays the same data as *Services* > *SLA report*.

Configuration

To configure, select *SLA report* as type:



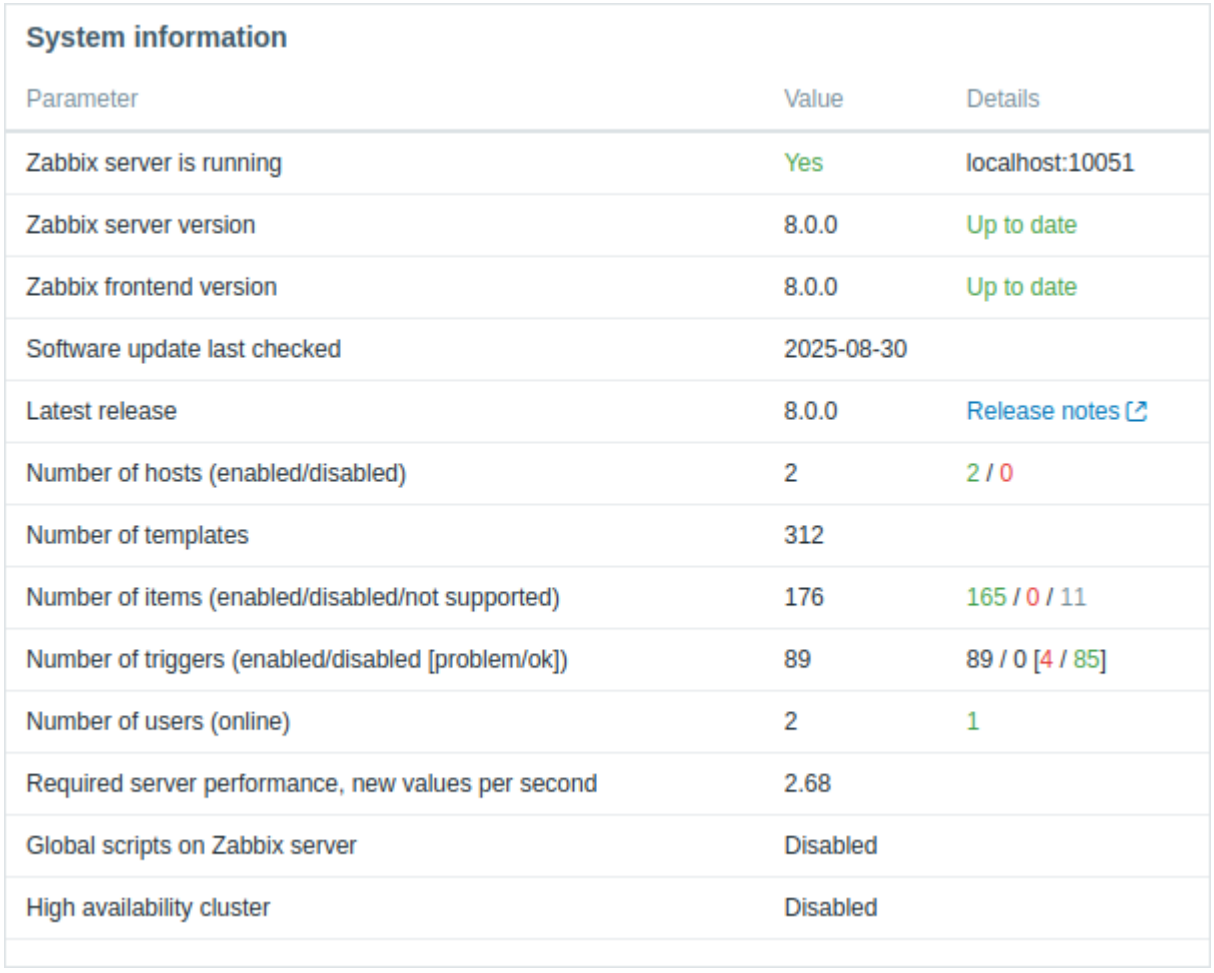
In addition to the parameters that are common for all widgets, you may set the following specific options:

SLA	Select the SLA for the report.
Service	Select the service for the report.
Show periods	Set how many periods will be displayed in the widget (20 by default, 100 maximum).
From	Select the beginning date for the report. Relative dates are supported: `now`, `now/d`, `now/w-1w` etc; supported date modifiers: d, w, M, y.
To	Select the end date for the report. Relative dates are supported: `now`, `now/d`, `now/w-1w` etc; supported date modifiers: d, w, M, y.

System information

Overview

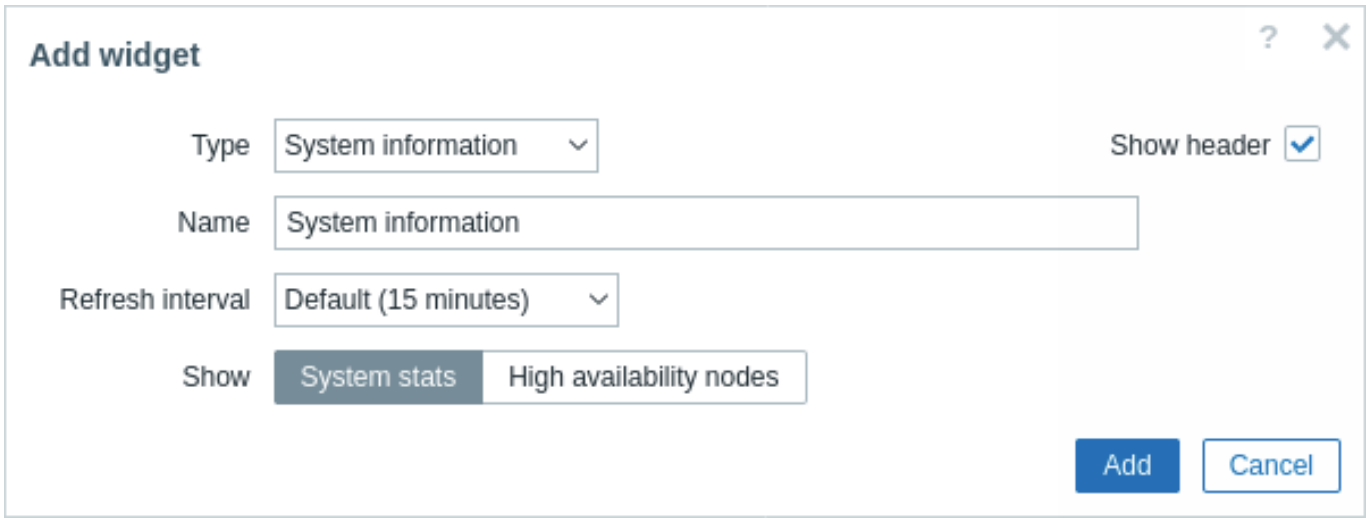
The *System information* widget displays a summary of key Zabbix server and system data, and, when high availability is enabled, details of the high availability nodes.



It displays the same data as *Reports* > *System information*, except for *Zabbix server ID*.

Configuration

To configure, select *System information* as type:



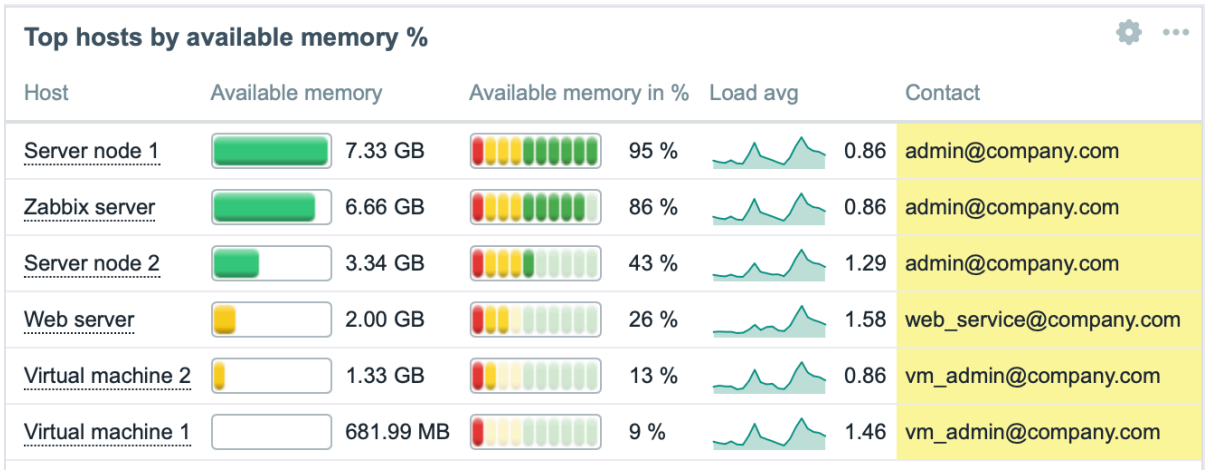
In addition to the parameters that are common for all widgets, you may set the following specific options:

Show	Select what to display: **System stats** - display a summary of key Zabbix server and system data; **High availability nodes** - display the status of high availability nodes (if high availability cluster is enabled).
Show software update check details	Mark the checkbox to display Zabbix software update check details. This option is only available if software update check is enabled in the Zabbix server configuration and "System stats" is selected in the *Show* field.

Top hosts

Overview

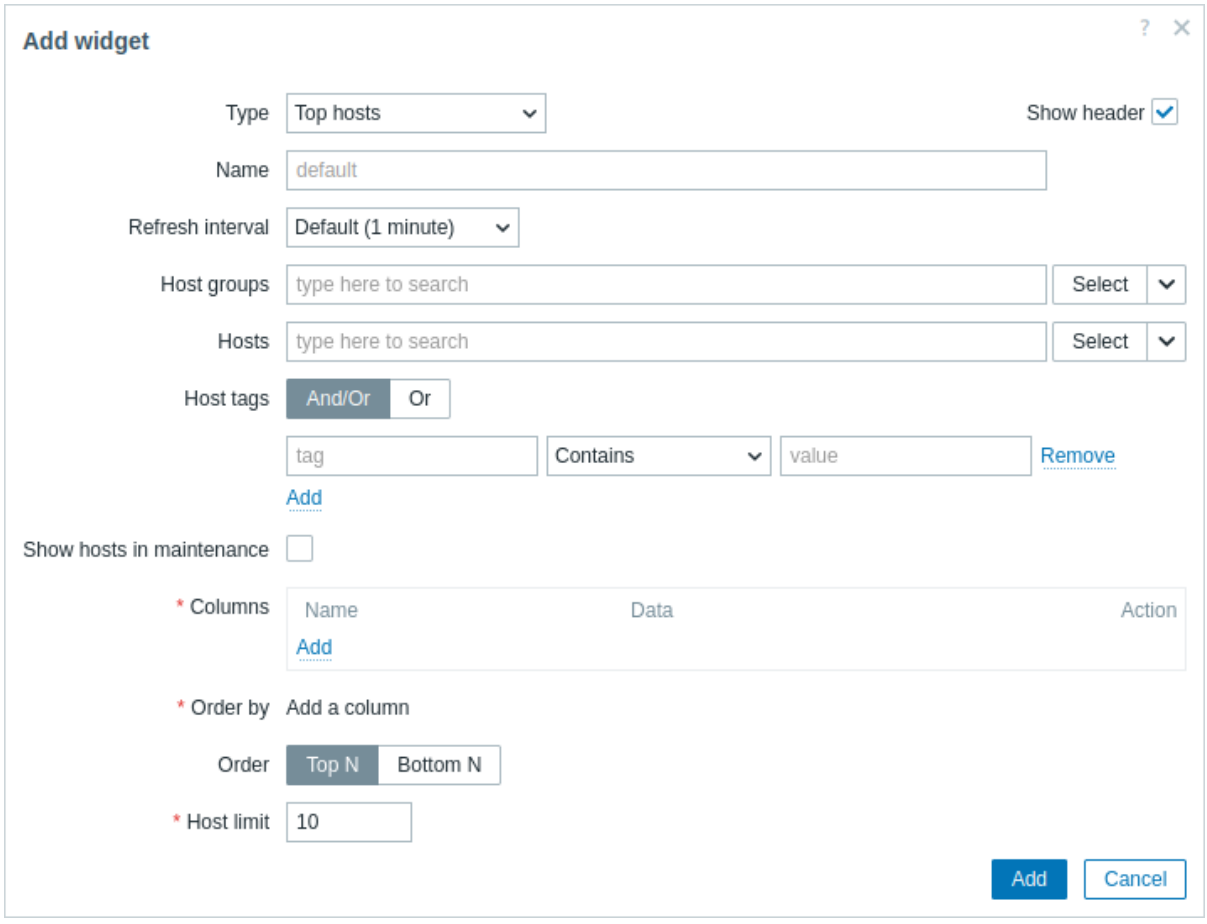
The *Top hosts* widget displays the top/bottom values for selected items from multiple hosts.



The widget supports displaying up to 1000 hosts simultaneously.

Configuration

To configure, select *Top hosts* as type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

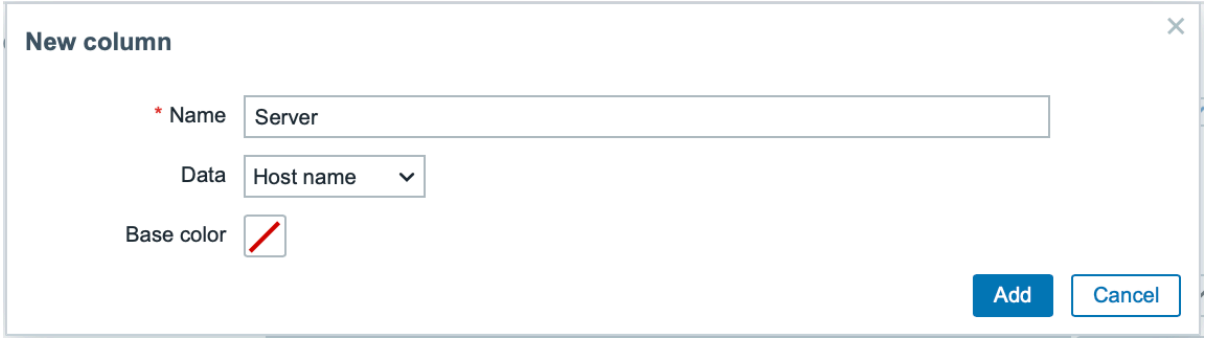
Host groups	Select host groups to display in the widget. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. This parameter is not available when configuring the widget on a template dashboard.
Hosts	Select hosts to display in the widget. Alternatively, select a compatible widget or the dashboard as the data source for hosts. This field is auto-complete, so starting to type the name of a host will offer a dropdown of matching hosts. This parameter is not available when configuring the widget on a template dashboard.
Host tags	Specify tags to limit the number of hosts displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met. This parameter is not available when configuring the widget on a template dashboard.
Show hosts in maintenance	Mark this checkbox for hosts in maintenance to be displayed as well (in this case, maintenance icon will be shown next to the host name). Unmarked by default.
Columns	Add data columns to display. The column order determines their display from left to right. Columns can be reordered by dragging up and down by the handle before the column name.
Order by	Specify the column from the defined *Columns* list to use for *Top N* or *Bottom N* ordering.
Order	Specify the ordering of rows: **Top N** - in descending order according to the *Order by* aggregated value; **Bottom N** - in ascending order according to the *Order by* aggregated value.
Host limit	Number of host rows to be shown (1-1000). This parameter is not available when configuring the widget on a template dashboard.

Column configuration

It is possible to add columns of three data types: host name, item value, or text. The list of available column parameters depends on the column data type and, for the item value type, on the value format.

Host name column

Host name column is used to display the name of the host.



Supported parameters:

Name	Name of the column.
Data	Data type to display in the column, select **Host name**.
Base color	Background color of the column.

Text column

Text column is used to display any specified text string.

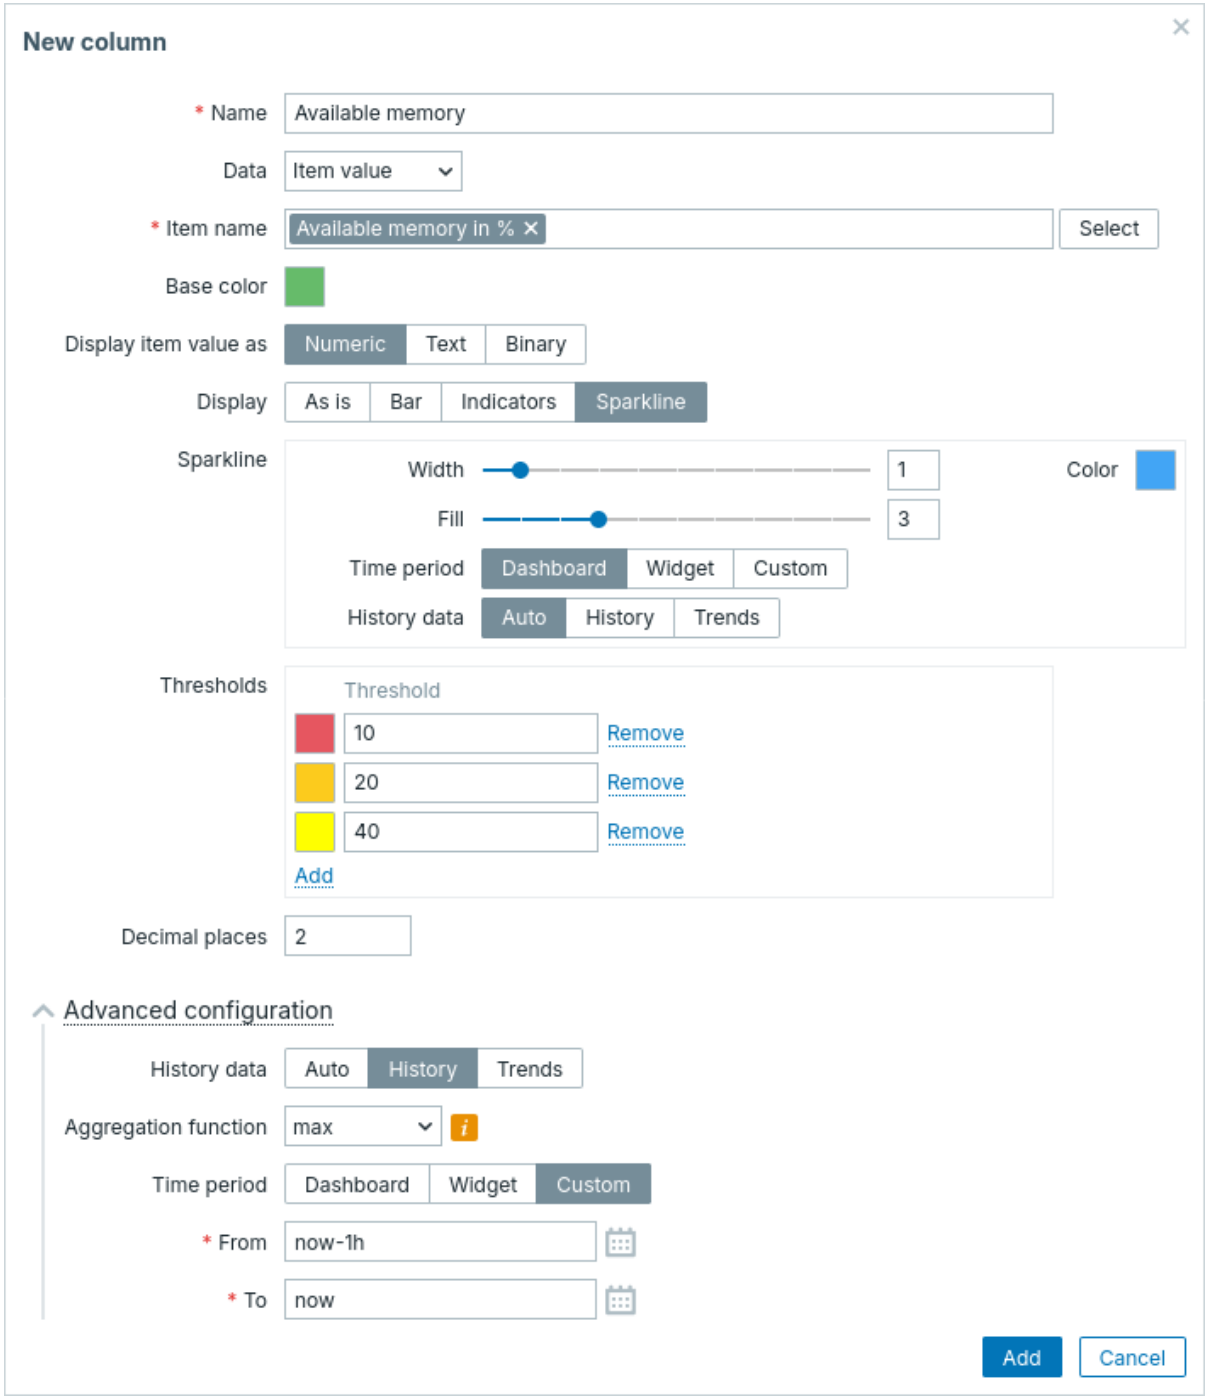


Supported parameters:

Name	Name of the column.
Data	Data type to display in the column, select **Text**.
Base color	Background color of the column.
Text	Enter the string to display. May contain host and inventory macros.

Item value column

Item value column is used to display value of the specified item; it supports multiple value display formats and options.



Supported parameters:

Name	Name of the column.
Data	Data type to display in the column, select **Item value**.
Item name	Select an item; its name is used to match and display items with the same name on all selected hosts. If a host has multiple such items, the widget displays the item whose key is first alphabetically (e.g., *CPU utilization* with key `proc.cpu.util` instead of `system.cpu.util`). When configuring the widget on a template dashboard, only items configured on the template can be selected.
Base color	Background color of the column; fill color if *Item value* data is displayed as bar/indicators. The default color can be overridden by custom color, if the item value is over one of the specified thresholds.
Display item value as	Format for displaying the item value: **Numeric**, **Text**, or **Binary**. The selected option determines which additional parameters will be available. Refer to the parameter list for each format below.
Advanced configuration	

History data	Take data from history or trends: **Auto** - automatic selection; **History** - take history data; **Trends** - take trend data. This parameter is available only for numeric item values. For text and binary values data will always be taken from history.
Aggregation function	Specify which aggregation function to use within the selected *Time period*: **min** - display the smallest value; **max** - display the largest value; **avg** - display the average value; **count** - display the count of values; **sum** - display the sum of values; **first** - display the first value; **last** - display the last value. If **not used** is selected, the most recent value is displayed (no aggregation). Aggregation allows an aggregated value to be displayed for the specified interval (5 minutes, an hour, a day), instead of the most recent value. Only numeric data can be displayed for *min*, *max*, *avg* and *sum*. For *count*, non-numeric data will be changed to numeric.
Time period	Specify the time period to use for aggregating values. Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period. This parameter is not available if *Aggregation function* is set to "not used".
Widget	Enter or select a compatible widget as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".

Numeric item value parameters:

Display	Define how the value should be displayed: **As is** - as regular text; **Bar** - as solid, color-filled bar; **Indicators** - as segmented, color-filled bar; **Sparkline** - mini line graph.
Thresholds	Specify threshold values when the background/fill color should change. The list will be sorted in ascending order when saved. For sparklines, thresholds are applied only to the item last value.
Decimal places	Specify how many decimal places will be displayed with the value.
Bar/Indicators	
Min	Minimum value.
Max	Maximum value.
Sparkline	
Width	Set the graph line thickness by using the slider or manually entering a value in the range from 0 to 10.
Color	Select line and fill color.
Fill	Set fill color transparency level by using the slider or manually entering a value in the range from 0 to 10.

Time period	Specify the time period for values to be included into the sparkline chart. Select the data source for the time period: **Dashboard** - use the dashboard time period selector **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period.
Widget	Enter or select a compatible widget as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
History data	Take data from history or trends: **Auto** - automatic selection; **History** - take history data; **Trends** - take trend data.

Text item value parameters:

Highlights	Specify the regular expressions upon matching which the background/fill color should change.

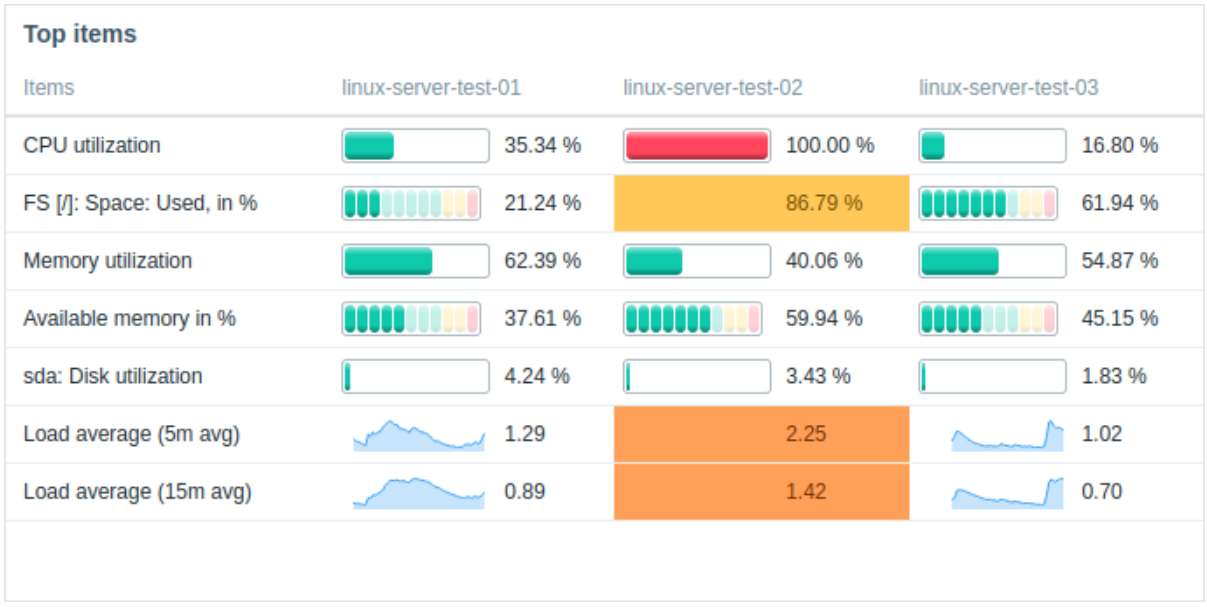
Binary item value parameters:

Show thumbnail	Specify whether to create and display a thumbnail for the image containing binary data or to display a hyperlink *Show* leading to the full-size image in the value column.

Top items

Overview

The *Top items* widget displays the top/bottom values for selected items, providing a quick overview of their performance.

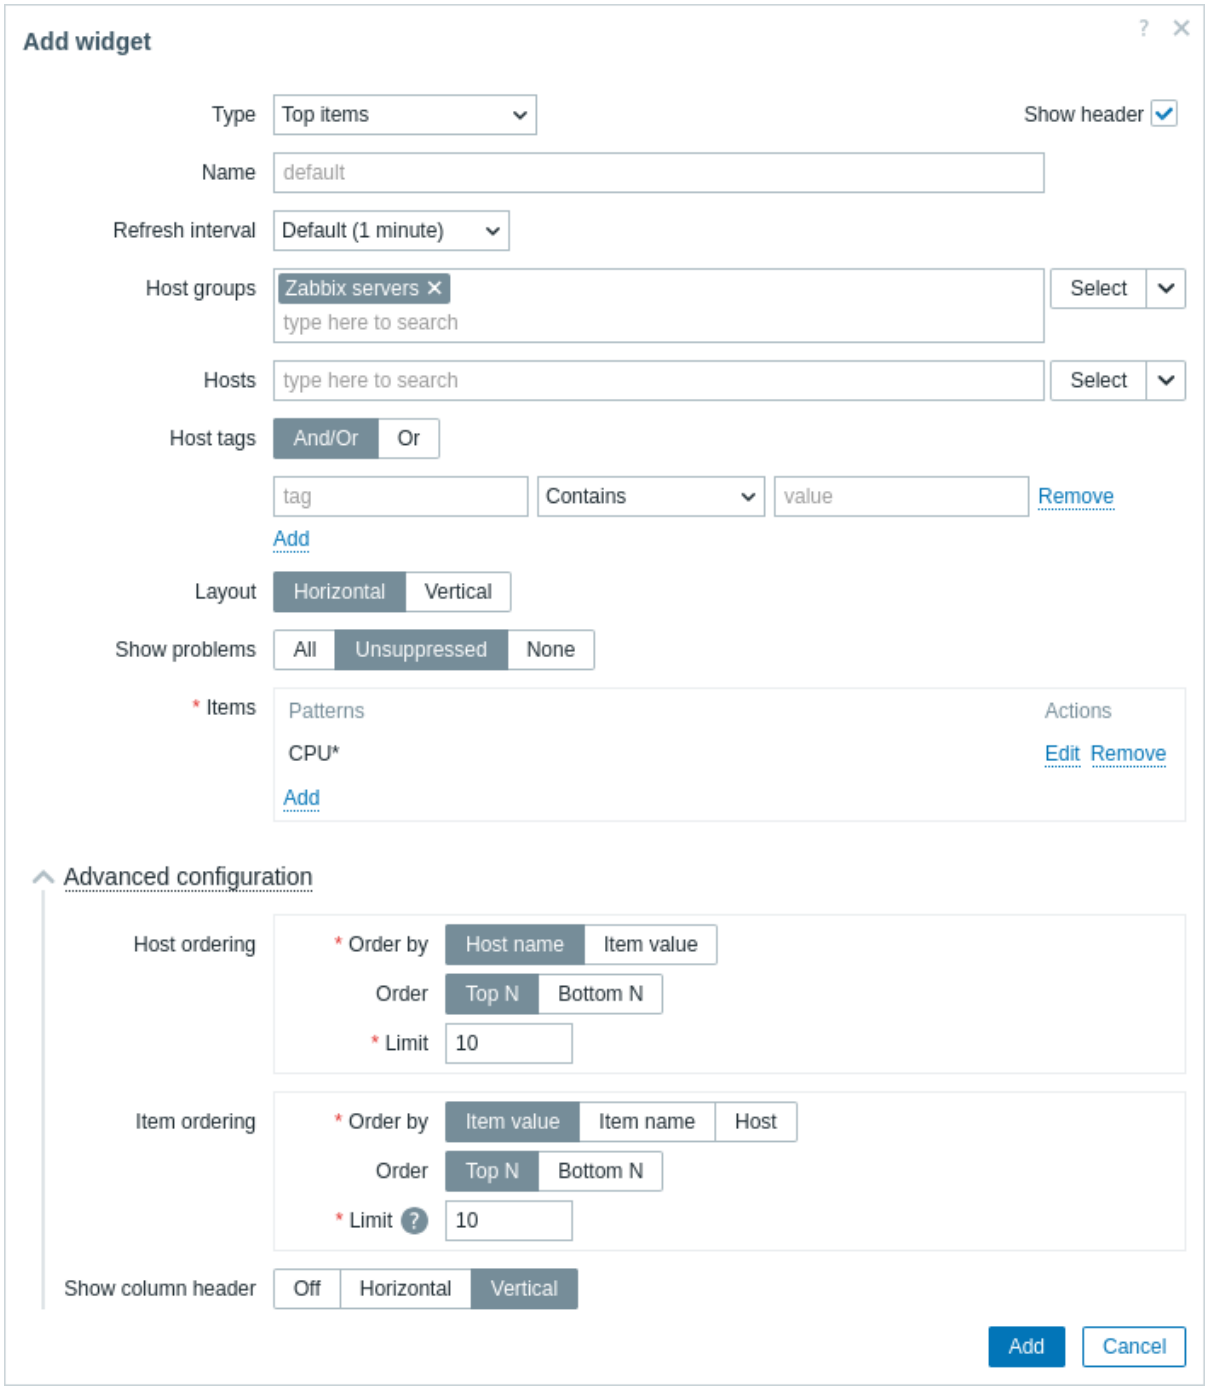


By default, the widget displays values from the last 24 hours. This period can be changed in *Administration* > *General* > *GUI* (*Max history display period*).

Clicking the item value opens the item menu.

Configuration

To configure, select *Top items* as type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

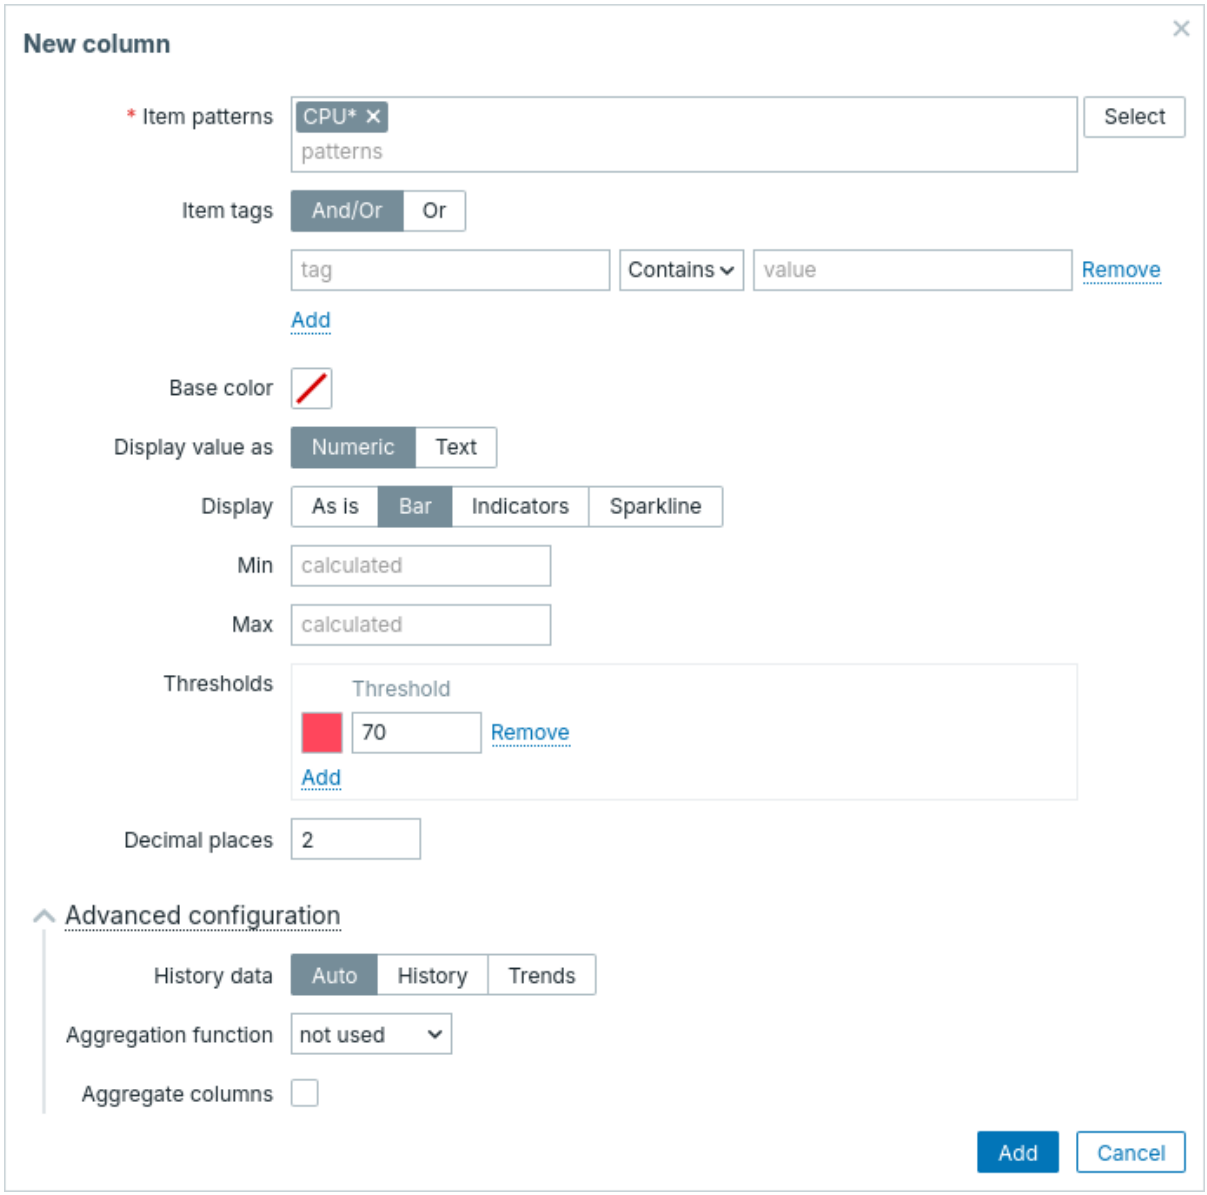
Host groups	Select host groups. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete so starting to type the name of a group will offer a dropdown of matching groups. This parameter is not available when configuring the widget on a template dashboard.
Hosts	Select hosts. Alternatively, select a compatible widget or the dashboard as the data source for hosts. This field is auto-complete so starting to type the name of a host will offer a dropdown of matching hosts. This parameter is not available when configuring the widget on a template dashboard.

Host tags		Specify tags to limit the number of host data displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the *Or* condition; **Or** - enough if one condition is met.
Layout		Select the layout option: **Horizontal** - host names will be displayed on the left; **Vertical** - host names will be displayed at the top.
Show problems		Filter which problems to display based on their status (all, unsuppressed, none). Problems are indicated with colored item values, overriding the item column display settings. Colors are based on problem severity, which can be customized in *Administration > General > Trigger displaying options*. Problem severities can be changed when updating problems.
Items		Add item patterns or specific items for display (see Column configuration).
Advanced configuration		
Host ordering		Select ordering options for the host column/row.
	Order by	Order hosts by: **Host name** - hosts will be ordered by host name; **Item value** - hosts will be ordered by the value of selected items.
	Order	Select whether to display the highest or lowest values: **TopN** - top N values; **BottomN** - bottom N values. The value of N is selected in the *Limit* field.
	Limit	Enter limit for displayable hosts (range 1-1000; default 10). This value will become the value of N in the *Order* field.
Item ordering		Select ordering options for the item column/row.
	Order by	Order items by: **Item value** - items will be ordered by item value; **Item name** - items will be ordered by item name; **Host** - items will be ordered by the selected host pattern.
	Order	Select whether to display the highest or lowest values: **TopN** - top N values; **BottomN** - bottom N values. The value of N is selected in the *Limit* field.
	Limit	Enter limit for displayable items (range 1-1000; default 10). This value will become the value of N in the *Order* field.

Show column header	Select column header display options: **Off** - do not display column header; **Horizontal** - display text horizontally in the header; **Vertical** - display text vertically in the header.

Column configuration

To configure item columns, click *Add* in the *Items* parameter:



Common column parameters:

Item patterns	Specify one or several item patterns, using the wildcard character. Alternatively, select the items. When configuring the widget on a template dashboard, only items configured on the template can be selected.

Item tags	Specify tags to limit the number of item data displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the *Or* condition; **Or** - enough if one condition is met.
Base color	Select the column's background color or fill color if *Display* is set to "Bar" or "Indicators". Note that the base color can be overridden by threshold or highlight colors.
Display value as	Format for displaying the item value: **Numeric** or **Text**. The selected option determines which additional parameters will be available. Refer to the parameter list for each format below.

Column parameters for numeric values:

Display	Define how the value should be displayed: **As is** - as regular text; **Bar** - as solid, color-filled bar; **Indicators** - as segmented, color-filled bar; **Sparkline** - mini line graph.
Min	Minimum value for bar/indicators display.
Max	Maximum value for bar/indicators display.
Width	Set the graph line thickness by using the slider or manually entering a value in the range from 0 to 10. This parameter is for sparkline display only.
Fill	Set fill color transparency level by using the slider or manually entering a value in the range from 0 to 10. This parameter is for sparkline display only.
Color	Select line and fill color. This parameter is for sparkline display only.
See *Advanced configuration* for description of fields related to time period and history data selection.	
Thresholds	Specify threshold values when the background/fill color should change. The list will be sorted in ascending order when saved. For sparklines, thresholds are applied only to the item last value.
Decimal places	Specify how many decimal places will be displayed with the value.
Advanced configuration	
History data	Take data from history or trends: **Auto** - automatic selection; **History** - take history data; **Trends** - take trend data. This parameter is available only for numeric item values. For text values data will always be taken from history.

Aggregation function	Specify which aggregation function to use within the selected *Time period*: **min** - display the smallest value; **max** - display the largest value; **avg** - display the average value; **count** - display the count of values; **sum** - display the sum of values; **first** - display the first value; **last** - display the last value. If **not used** is selected, the most recent value is displayed (no aggregation). Aggregation allows an aggregated value to be displayed for the specified interval (5 minutes, an hour, a day), instead of the most recent value. Only numeric data can be displayed for *min*, *max*, *avg* and *sum*. For *count*, non-numeric data will be changed to numeric.
Time period	Specify the time period to use for aggregating values. Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period. This parameter is not available if *Aggregation function* is set to "not used".
Widget	Enter or select a compatible widget as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
Aggregate columns	When checked, group all items that match the column's item patterns into a single combined column/row. Values from matching items will be aggregated using the selected *Column aggregation function* and displayed under the *Combined column name*. Default: off (unchecked).This parameter is shown only when *Display value as* is **Numeric**.
Column aggregation function	Choose how values from the grouped item patterns are combined. Available options: **min**, **max**, **avg**, **count**, **sum** (default). This parameter is shown only when *Aggregate columns* is checked.
Combined column name	Enter the display name for the aggregated column. This field is mandatory when *Aggregate columns* is checked and becomes the column header for the combined values.

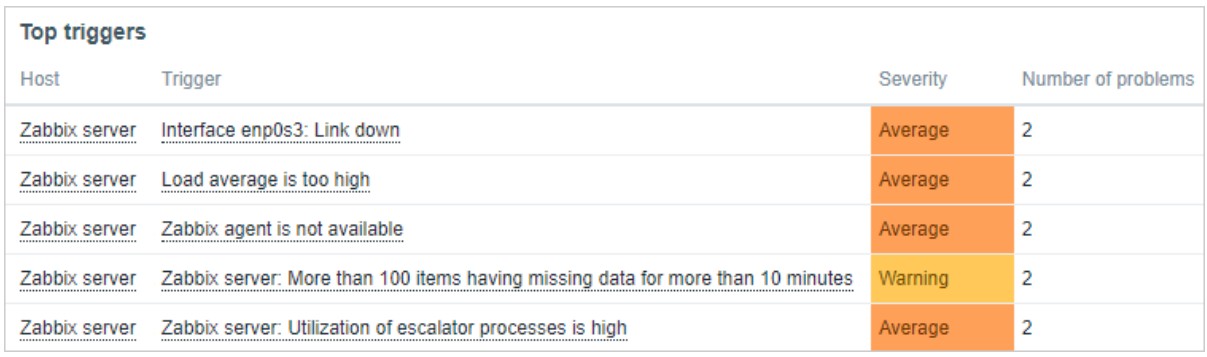
Column parameters for text values:

Highlights	Specify the regular expressions upon matching which the background/fill color should change.

Top triggers

Overview

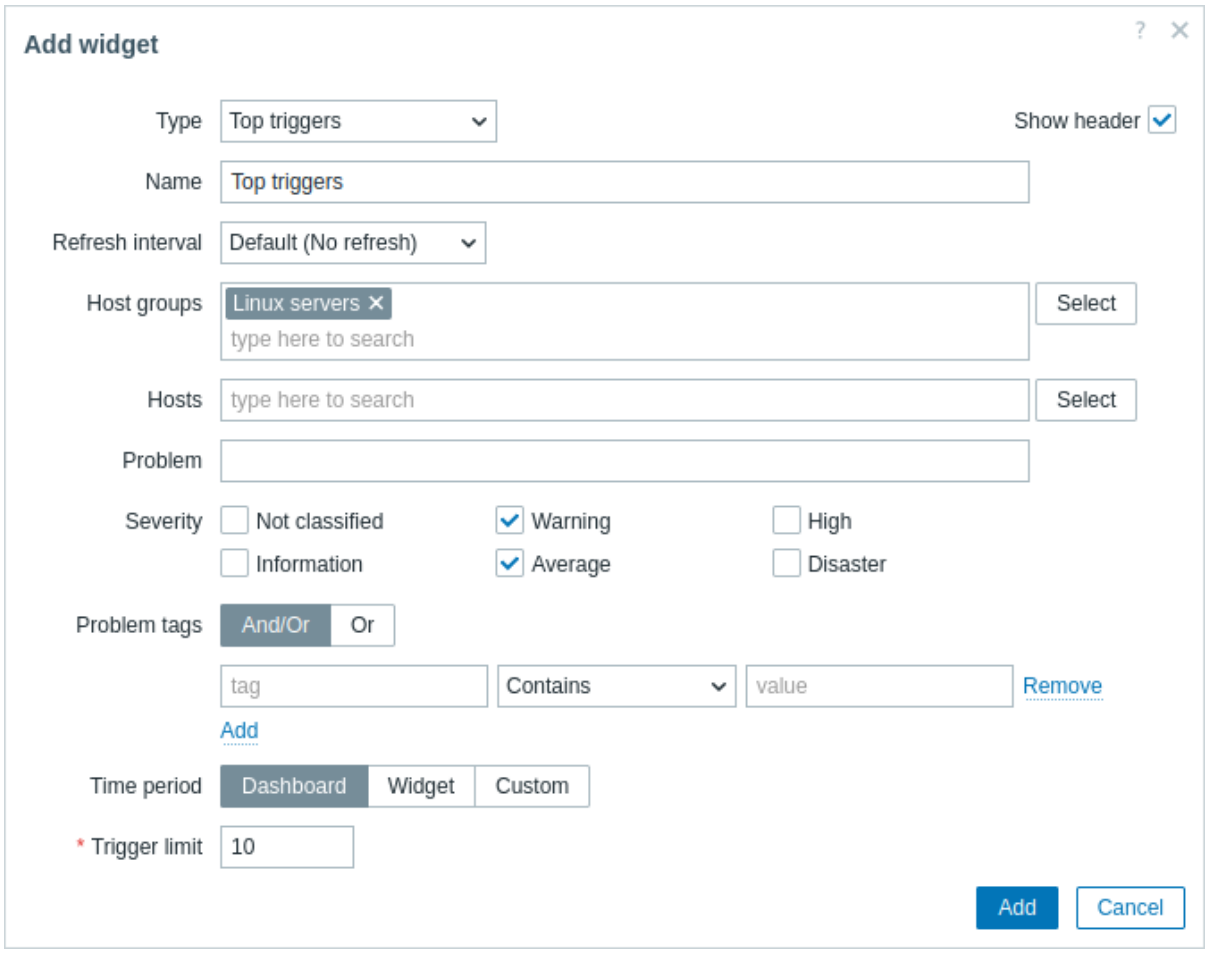
The *Top triggers* widget displays triggers with the highest number of problems.



It displays the same data as *Reports > Top 100 triggers* and can display up to 1000 records.

Configuration

To configure, select *Top triggers* as widget type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

Host groups	Select host groups to display problems of in the widget. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Specifying a parent host group implicitly selects all nested host groups. Problems from these host groups will be displayed in the widget; if no host groups are entered, problems from all host groups will be displayed. This parameter is not available when configuring the widget on a template dashboard.
Hosts	Select hosts to display problems of in the widget. This field is auto-complete, so starting to type the name of a host will offer a dropdown of matching hosts. If no hosts are entered, problems of all hosts will be displayed. This parameter is not available when configuring the widget on a template dashboard.
Problem	You can view the triggers for particular problems only. For this, enter the string to be matched in the problem name. Macros are not expanded.
Severity	Mark trigger severities to filter triggers to be displayed in the widget. If no severities are marked, all triggers will be displayed.

Problem tags	Specify the tags of the problems to be displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the *Or* condition; **Or** - enough if one condition is met.
Time period	Select the data source for the time period: **Dashboard** - use the dashboard time period selector; **Widget** - use a compatible widget (set in the *Widget* parameter); **Custom** - use a custom time period set in the *From* and *To* parameters; if set, a clock icon will be displayed in the upper-right corner of the widget, indicating the set time on mouseover. Note that regardless of the widget's *Time period* configuration, compatible widgets can still use it as a data source for the time period.
Widget	Enter or select a compatible widget as the data source for the time period. This parameter is available if *Time period* is set to "Widget".
From	Enter or select the start of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
To	Enter or select the end of the time period. Relative time syntax (`now`, `now/d`, `now/w-1w`, etc.) is supported. This parameter is available if *Time period* is set to "Custom".
Trigger limit	Set the number of triggers to be displayed. Possible value range: 1-1000.

Trigger overview

Overview

The *Trigger overview* widget displays the current states of host triggers as a table of colored blocks. It provides a quick visual snapshot of trigger health and activity on different hosts.



The widget uses multiple visual elements to convey trigger details:

- Block colors for PROBLEM triggers match problem severities.
- Blinking blocks indicate recent trigger state changes (within the last 2 minutes).
- Gray arrows indicate triggers with dependencies (hover over them to see details).
- Checkbox icons indicate acknowledged problems; all active or resolved problems for the trigger must be acknowledged for this icon to appear.

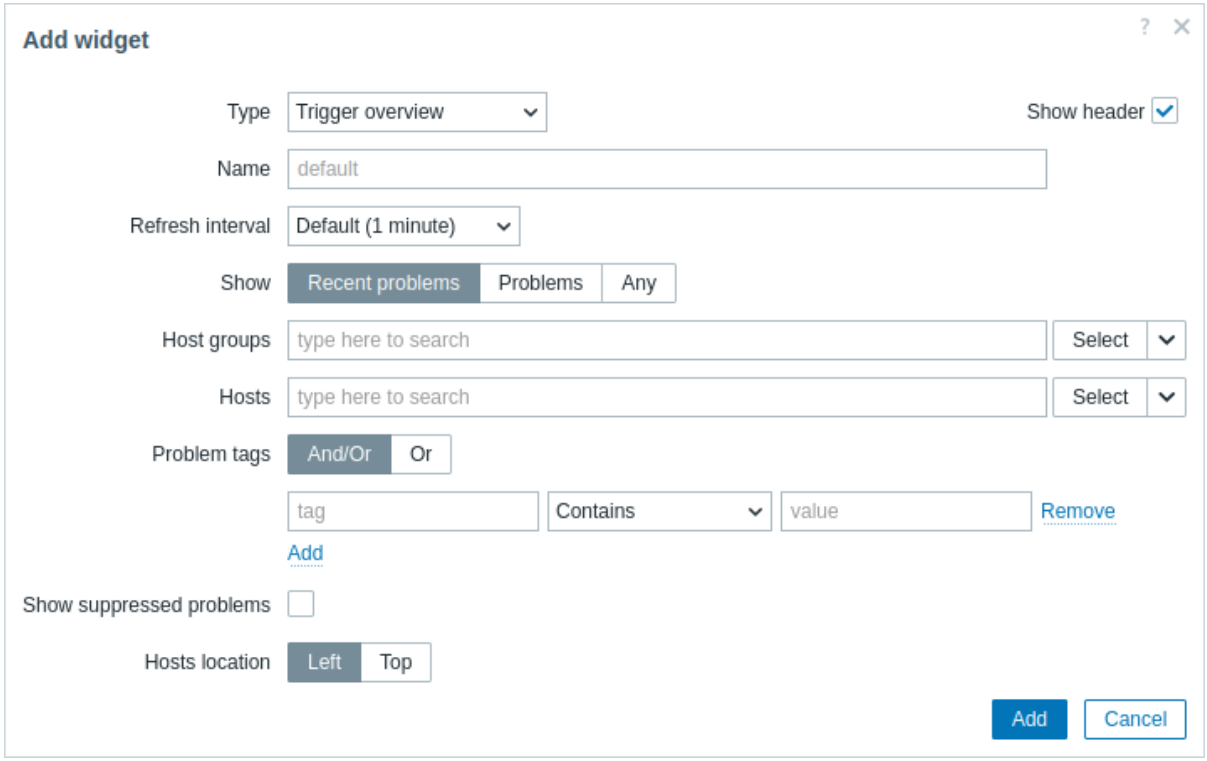
Note:
Trigger severity colors can be adjusted in *Administration* > *General*. Trigger severities can be modified when updating problems.

Clicking a block provides links to trigger-related information (problems, trigger configuration, etc.).

By default, up to 50 records are displayed (configurable in *Administration* > *General* > *GUI* under *Max number of columns and rows in overview tables*). If more triggers exist, a message prompts you to refine the filtering criteria. This limit is applied before any additional filtering (for example, by tags).

Configuration

To configure, select *Trigger overview* as type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

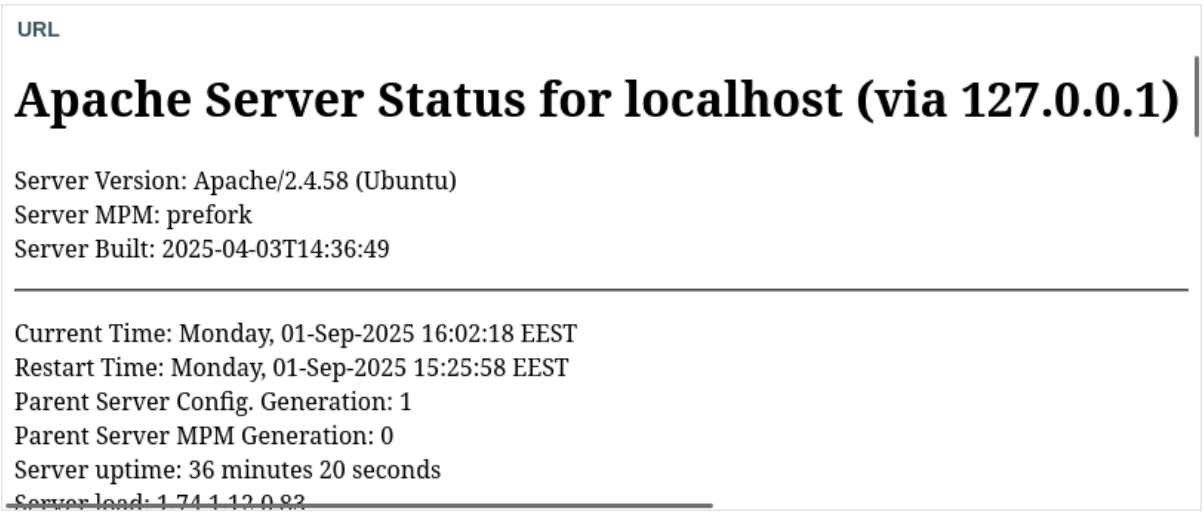
Show	Filter triggers by trigger state: **Recent problems** - *(default)* show triggers that recently have been or still are in a PROBLEM state (resolved and unresolved); **Problems** - show triggers that are in a PROBLEM state (unresolved); **Any** - show all triggers.
Host groups	Select host groups. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. This parameter is not available when configuring the widget on a template dashboard.
Hosts	Select hosts. Alternatively, select a compatible widget or the dashboard as the data source for hosts. This field is auto-complete, so starting to type the name of a host will offer a dropdown of matching hosts. This parameter is not available when configuring the widget on a template dashboard.

Problem tags	Specify tags to filter the triggers displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. **Note:** If the parameter *Show* is set to 'Any', all triggers will be displayed even if tags are specified. However, while recent trigger state changes (displayed as blinking blocks) will update for all triggers, the trigger state details (problem severity color and whether the problem is acknowledged) will only update for triggers that match the specified tags. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the *Or* condition; **Or** - enough if one condition is met.
Show suppressed problems	Mark the checkbox to display problems that would otherwise be suppressed (not shown) because of host maintenance.
Layout	Select the layout option: **Horizontal** - host names will be displayed on the left; **Vertical** - host names will be displayed at the top.

URL

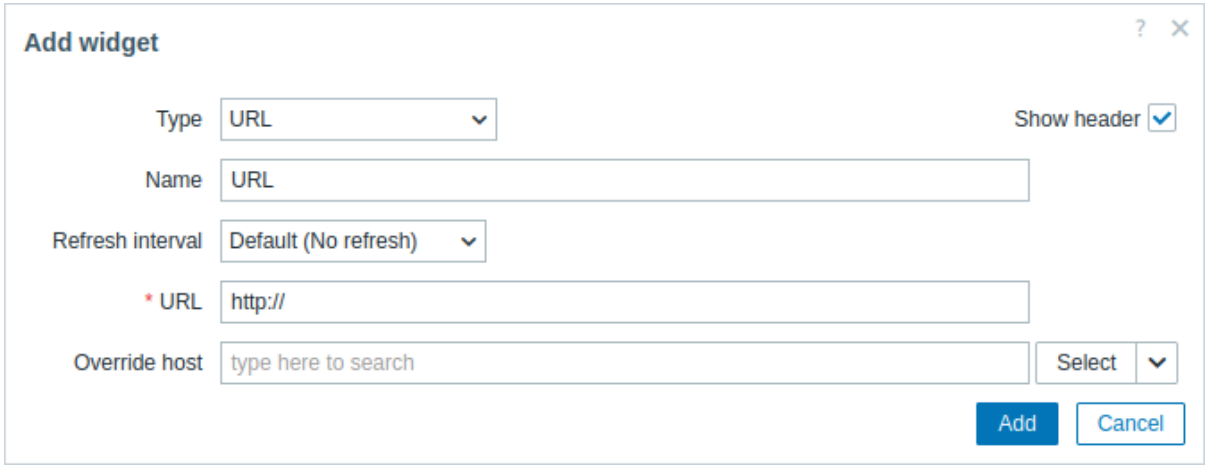
Overview

The *URL* widget displays the content retrieved from a URL.



Configuration

To configure, select *URL* as type:



In addition to the parameters that are common for all widgets, you may set the following specific options:

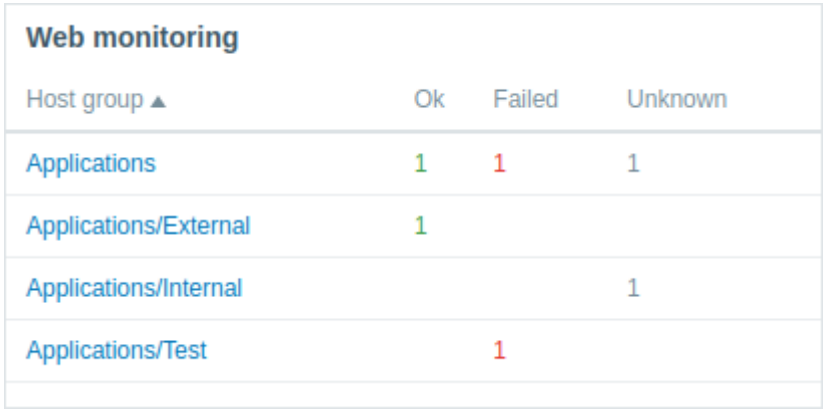
URL	Enter the URL to display (up to 2048 characters). External URLs must start with `http://` or `https://`. Internal URLs support relative paths (for example, `zabbix.php?action=report.status`). `{HOST.*}` macros are supported.
Override host	Select a compatible widget or the dashboard host selector as the data source for hosts. This parameter is not available when configuring the widget on a template dashboard.

Attention:
Browsers might not load an HTTP page configured in the widget if Zabbix frontend is accessed over HTTPS. Also, you may need to modify the X-Frame-Options HTTP response header to display content external to Zabbix frontend.

Web monitoring

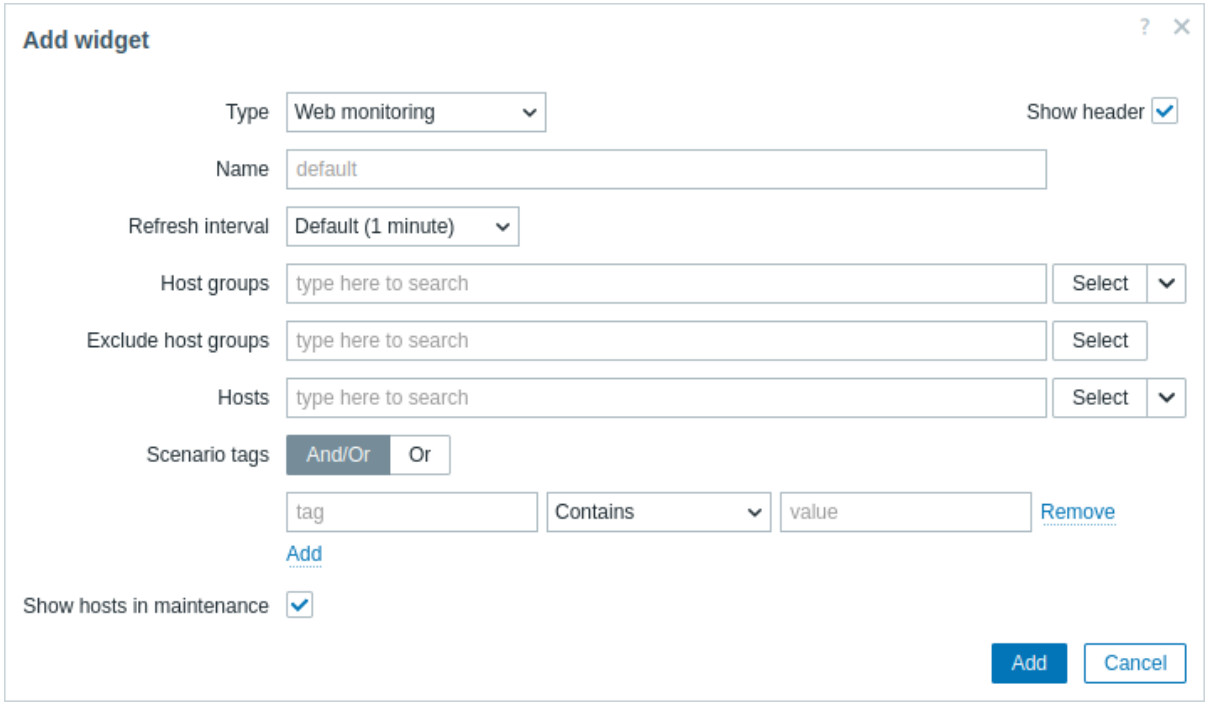
Overview

The *Web monitoring* widget displays the status of active web scenarios.



Clicking a host group name opens *Monitoring* > *Hosts* > *Web scenarios*, where you can view scenario details.

Configuration



In addition to the parameters that are common for all widgets, you may set the following specific options:

Host groups	Select host groups to display in the widget. Alternatively, select a compatible widget as the data source for host groups. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Specifying a parent host group implicitly selects all nested host groups. Host data from these host groups will be displayed in the widget; if no host groups are entered, all host groups will be displayed. This parameter is not available when configuring the widget on a template dashboard.
Exclude host groups	Select host groups to hide from the widget. This field is auto-complete, so starting to type the name of a group will offer a dropdown of matching groups. Specifying a parent host group implicitly selects all nested host groups. Host data from these host groups will not be displayed in the widget. For example, hosts 001, 002, 003 may be in Group A and hosts 002, 003 in Group B as well. If we select to *show* Group A and *exclude* Group B at the same time, only data from host 001 will be displayed in the dashboard. This parameter is not available when configuring the widget on a template dashboard.
Hosts	Select hosts to display in the widget. Alternatively, select a compatible widget or the dashboard as the data source for hosts. This field is auto-complete, so starting to type the name of a host will offer a dropdown of matching hosts. If no hosts are entered, all hosts will be displayed. This parameter is not available when configuring the widget on a template dashboard.

Scenario tags	Specify tags to limit the number of web scenarios displayed in the widget. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the *Or* condition; **Or** - enough if one condition is met.
Show hosts in maintenance	Include hosts that are in maintenance in the statistics. This parameter is labeled *Show data in maintenance* when configuring the widget on a template dashboard.

Monitoring

Overview

The Monitoring menu is all about displaying data. Whatever information Zabbix is configured to gather, visualize and act upon, it will be displayed in the various sections of the *Monitoring* menu.

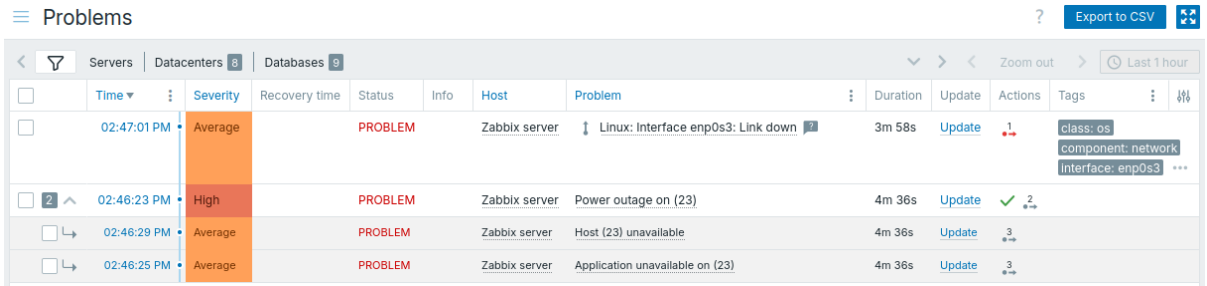
Problems

Overview

In *Monitoring* > *Problems* you can see what problems you currently have. Problems are those triggers that are in the "Problem" state.

By default, all new problems are classified as cause problems. It is possible to reclassify certain problems as symptom problems of the cause problem:

- In complex event processing
- Manually (see cause and symptom events).

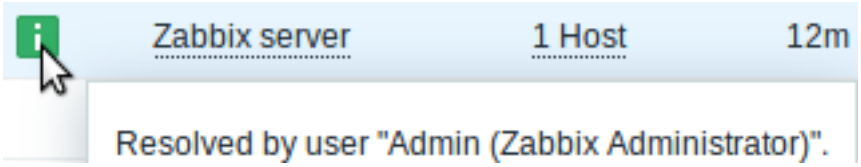
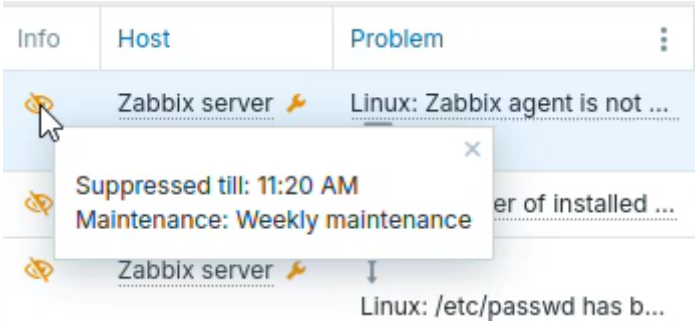


Problem details are displayed in a customizable table. The *Customize table* button in the table header lets you show, hide, or reorder columns, and reset the layout to default. These options are also available:

- *Compact view* - reduces row height to fit more problems on screen.
- *Highlight full row* - uses the problem severity color to highlight unresolved problems (not available in high-contrast themes).

You can also resize each column by dragging its right border; double-click the border to reset. All customizations are applied immediately and always remembered for the current user within the active filter tab (default or saved).

The following columns are available:

Column	Description
Checkbox	Checkboxes for problem selection are displayed. Icons next to the checkboxes have the following meaning:  - the number of symptom events for the cause problem;  - expand to show symptom events;  - collapse to hide symptom events;  - this is a symptom event.
Time	Problem start time is displayed. Clicking the *Options* button in the column header gives you access to the *Show timeline* option, which toggles the visual timeline and problem grouping (not available when *Compact view* is enabled).
Severity	Problem severity is displayed. Problem severity is originally based on the severity of the underlying problem trigger, however, after the event has happened it can be updated using the *Update problem* screen. Color of the problem severity is used as cell background during problem time.
Recovery time	Problem resolution time is displayed This column is not available when *Show* in the filter is set to *Problems*.
Status	Problem status is displayed: **Problem** - unresolved problem **Resolved** - recently resolved problem. You can hide recently resolved problems using the filter. New and recently resolved problems blink for 2 minutes. Resolved problems are displayed for 5 minutes in total. Both of these values are configurable in *Administration* > *General* > *Trigger displaying options*. This column is not available when *Show* in the filter is set to *Problems*.
Info	A green information icon is displayed if a problem is closed by global correlation or manually when updating the problem. Hovering over the icon displays more details:  If a suppressed problem is shown, the following icon is displayed:  . When the event is suppressed by host maintenance, the icon is shown in orange:  . If both maintenance and manual suppression apply, manual suppression takes priority in the icon color. Hovering over the icon displays more details: 
Host	Problem host is displayed. Clicking the host name brings up the host menu.

Column	Description
Problem	Problem name is displayed. Problem name is based on the name of the underlying problem trigger. Macros in the trigger name are resolved at the time of the problem happening and the resolved values do not update any more. Clicking the problem name brings up the event menu. Hovering on the  icon after the problem name will bring up the trigger description (for those problems that have it). Clicking the *Options* button in the column header gives you access to two options: - *Show operational data* that displays operational data after the problem name in parentheses (only if the *Operational data* field is not empty in the trigger configuration). - *Show trigger expression* that displays the problem's trigger expression below the problem name (not available when *Compact view* is enabled).
Operational data	Dynamic, real-time data for the problem is displayed. The content depends on the *Operational data* field in trigger configuration: if a string with macros is defined, it is shown here; if the field is left empty, the latest values of all items in the trigger expression are shown instead.
Duration	Problem duration is displayed. See also: Negative problem duration
Update	Click the *Update* link to go to the problem update screen where various actions can be taken on the problem, including commenting and acknowledging the problem.
Actions	History of activities about the problem is displayed using symbolic icons:  - problem has been acknowledged. This icon is always displayed first.  - comments have been made. The number of comments is also displayed.  - problem severity has been increased (e.g., Information > Warning).  - problem severity has been decreased (e.g., Warning > Information).  - problem severity has been changed, but returned to the original level (e.g., Warning > Information > Warning).  - problem has been suppressed.  - problem has been unsuppressed.  - problem has been suppressed due to host maintenance.  - problem has been unsuppressed after host maintenance.  - actions have been taken. The number of actions is also displayed.  - actions have been taken, at least one is in progress. The number of actions is also displayed.  - actions have been taken, at least one has failed. The number of actions is also displayed. When rolling the mouse over the icons, popups with details about the activity are displayed. See viewing details to learn more about icons used in the popup for actions taken.
Tags	Problem tags are displayed. In addition, tags from an external ticketing system may also be displayed (see the *Process tags* option when configuring webhooks). Hovering or clicking the ••• three dots icon after the displayed tags opens a pop-up with all tags. Clicking the *Options* button in the column header lets you set the number of tags displayed (1-3), choose how tag names are displayed (full, shortened, or none), set tag display priority, or duplicate the column (duplicated columns can be renamed or deleted).
Tag value	Problem tag values are displayed. Clicking the *Options* button in the column header lets you rename the column, choose which tag values to display, or duplicate the column (duplicated columns can be renamed or deleted).

Column	Description
Custom text	Custom text, macro values, or a combination of both is displayed. Clicking the ⋮ *Options* button in the column header lets you rename the column, enter the custom text or macros, or duplicate the column (duplicated columns can be renamed or deleted). Supported macros: `{HOST.*}` (except `{HOST.METADATA}`), `{INVENTORY.*}`, `{ITEM.*}` (except `{ITEM.STATE.ERROR}`), `{TRIGGER.ID}`, `{EVENT.ID}`, and user macros.

Negative problem duration

It is actually possible in some common situations to have negative problem duration i.e. when the problem resolution time is earlier than problem creation time, e. g.:

- If some host is monitored by proxy and a network error happens, leading to no data received from the proxy for a while, the nodata(/host/key) trigger will be fired by the server. When the connection is restored, the server will receive item data from the proxy having a time from the past. Then, the nodata(/host/key) problem will be resolved and it will have a negative problem duration;
- When item data that resolve the problem event are sent by Zabbix sender and contain a timestamp earlier than the problem creation time, a negative problem duration will also be displayed.

Note:
Negative problem duration is not affecting SLA calculation or Availability report of a particular trigger in any way; it neither reduces nor expands problem time.

Mass editing options

Buttons below the list offer some mass-editing options:

- *Mass update* - update the selected problems by navigating to the problem update screen

To use this option, mark the checkboxes before the respective problems, then click the *Mass update* button.

Buttons

The button to the right offers the following option:

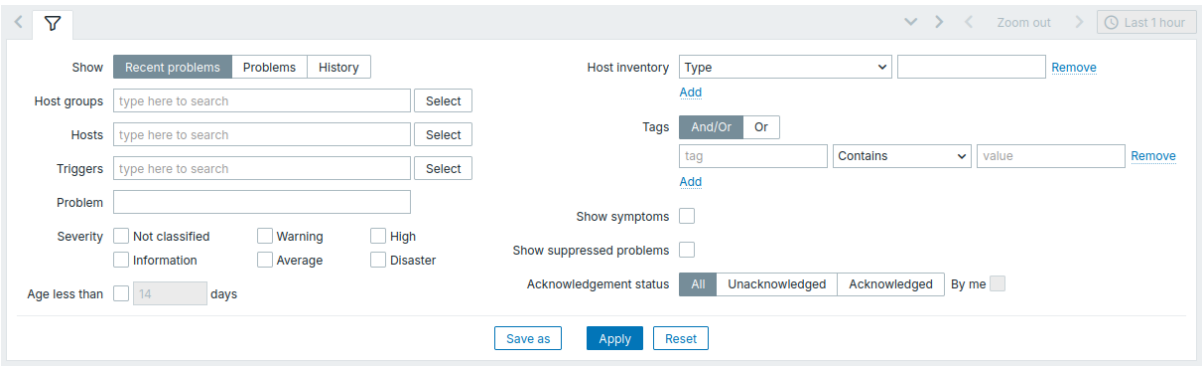
Export to CSV	Export content from all pages to a CSV file.

View mode buttons, being common for all sections, are described on the Monitoring page.

Using filter

You can use the filter to display only the problems you are interested in. For better search performance, data is searched with macros unresolved.

The filter is located above the table. Favorite filter settings can be saved as tabs and then quickly accessed by clicking the tabs above the filter.

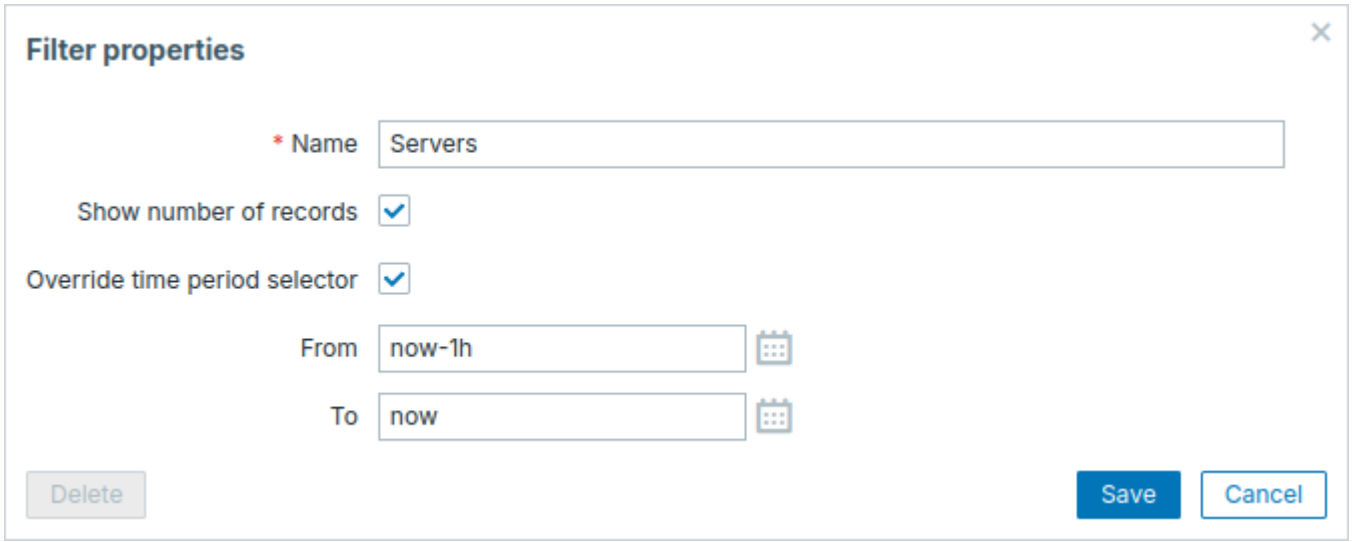


Parameter	Description
Show	Filter by problem status: **Recent problems** - unresolved and recently resolved problems are displayed (default) **Problems** - unresolved problems are displayed **History** - history of all events is displayed
Host groups	Filter by one or more host groups. Specifying a parent host group implicitly selects all nested host groups.
Hosts	Filter by one or more hosts.
Triggers	Filter by one or more triggers.
Problem	Filter by problem name.
Severity	Filter by trigger (problem) severity.
Age less than	Filter by how old the problem is.
Host inventory	Filter by inventory type and value.
Tags	Filter by event tag name and value. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names **Equals** - include the specified tag names and values (case-sensitive) **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive) **Does not exist** - exclude the specified tag names **Does not equal** - exclude the specified tag names and values (case-sensitive) **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive) There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition **Or** - enough if one condition is met When filtered, the tags specified here will be displayed first with the problem, unless overridden by the *Tag display priority* (see below) list.
Show symptoms	Mark the checkbox to display in its own line problems classified as symptoms.
Show suppressed problems	Mark the checkbox to display problems that would otherwise be suppressed and hidden because of host maintenance or suppressed manually by a user.
Acknowledgement status	Filter to display all problems, unacknowledged problems only, or acknowledged problems only. Mark the additional checkbox to filter out those problems ever acknowledged by you.

Tabs for favorite filters

Frequently used sets of filter parameters can be saved in tabs.

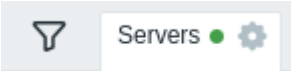
To save a new set of filter parameters, open the main tab, and configure the filter settings, then click the *Save as* button. In a new popup window, define *Filter properties*.



Parameter	Description
Name	The name of the filter to display in the tab list.
Show number of records	Check, if you want the number of problems to be displayed next to the tab name.
Override time period selector	Check to set specific default time period for this filter set. If set, you will only be able to change the time period for this tab by updating filter settings. For tabs without a custom time period, the time range can be changed by clicking the time selector button in the upper-right corner (button name depends on selected time interval: This week, Last 30 minutes, Yesterday, etc.). This option is available only for filters in *Monitoring* > *Problems*.
From/To	Time period start and end in absolute (`Y-m-d H:i:s`) or relative time syntax (`now-1d`). Available if *Set custom time period* is checked.

When saved, the filter is created as a named filter tab and immediately activated.

To edit the filter properties of an existing filter, click the gear symbol next to the active tab name.



Notes:

- To hide the filter area, click the name of the current tab. Click the active tab name again to open the filter area again.
- Keyboard navigation is supported: use arrows to switch between tabs, press *Enter* to open.
- The left/right buttons above the filter may be used to switch between saved filters. Alternatively, the downward pointing button opens a drop-down menu with all saved filters and you can click the one you need.
- If the settings of a saved filter have been changed (but not saved), a green dot is displayed after the filter name. To update the filter according to the new settings, click the *Update* button, which is displayed instead of the *Save as* button.
- Table column layout is saved per filter tab. To customize the layout of a saved filter, open that filter tab first, then customize the columns.
- Filter tabs, filter settings, and table column layout for each tab are remembered in the user profile and restored when the user opens the page again.

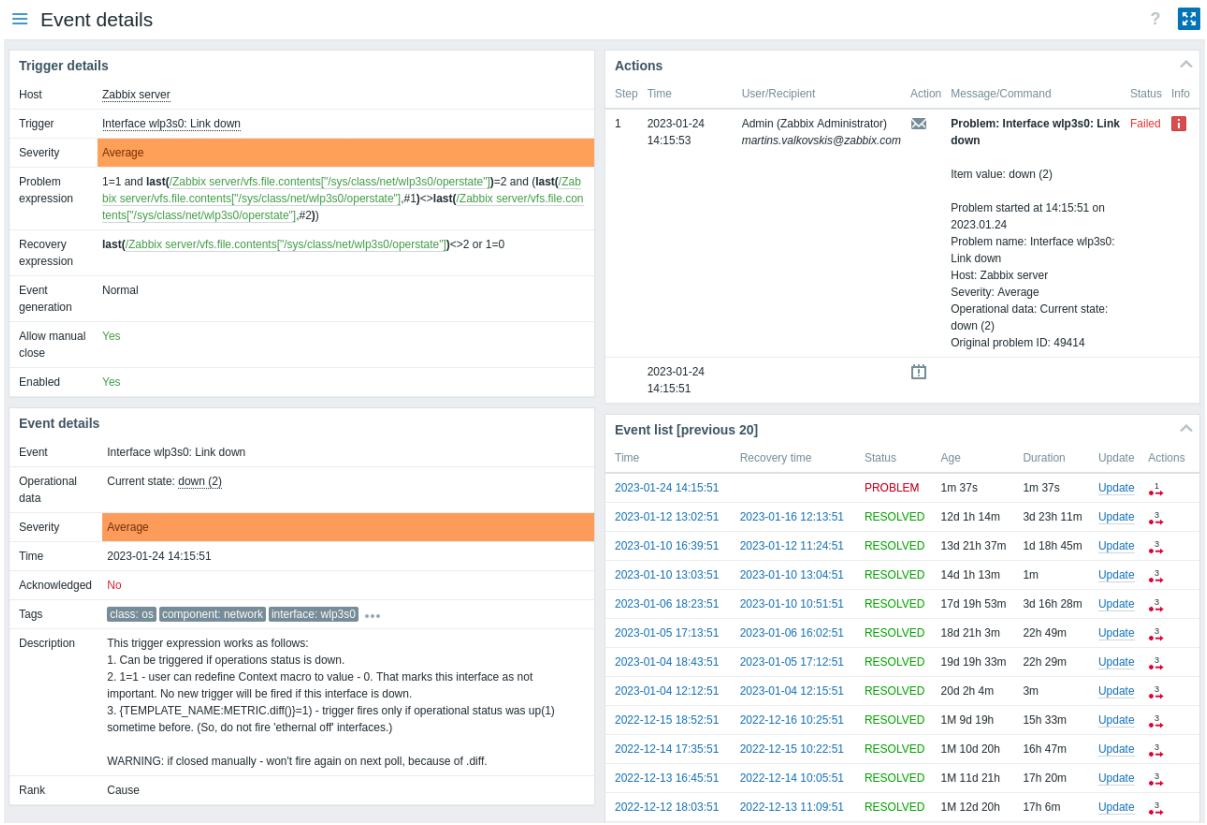
Note:
To share filters, copy and send to others a URL of an active filter. After opening this URL, other users will be able to save this set of parameters as a permanent filter in their Zabbix account.
 See also: Page parameters.

Filter buttons

Apply	Apply specified filtering criteria (without saving).
Reset	Reset current filter and return to saved parameters of the current tab. On the main tab, this will clear the filter.
Save as	Save current filter parameters in a new tab. Only available on the main tab.
Update	Replace tab parameters with currently specified parameters. Not available on the main tab.

Viewing details

The times for problem start and recovery in *Monitoring* > *Problems* are links. Clicking them opens more details of the event.



In event details, you can see whether the event was suppressed because of maintenance or suppressed manually by a user.

Note that the problem severity may differ for the trigger and the problem event - if it has been updated for the problem event using the *Update problem* screen.

In the action list, the following icons are used to denote the activity type:

- - problem event generated.
- - message has been sent.
- - problem event acknowledged.
- - problem event unacknowledged.
- - a comment has been added.
- - problem severity has been increased (e.g., Information > Warning).
- - problem severity has been decreased (e.g., Warning > Information).
- - problem severity has been changed, but returned to the original level (e.g., Warning > Information > Warning).
- - a remote command has been executed.
- - problem event has recovered.
- - the problem has been closed manually.
- - the problem has been suppressed.
- - the problem has been unsuppressed.
- - problem has been suppressed due to host maintenance.
- - problem has been unsuppressed after host maintenance.
- - the problem has been converted to a symptom problem.
- - the problem has been converted to a cause problem.
- - CEP operation has been performed (details in the description).

Cause and symptom problems

Overview

By default, all new problems are classified as cause problems. It is possible to manually reclassify certain problems as symptom problems of the cause problem.

For example, power outage may be the actual root cause why some host is unreachable or some service is down. In this case, "host is unreachable" and "service is down" problems must be classified as symptom problems of "power outage" - the cause problem.

The cause-symptom hierarchy supports only two levels. A problem that is already a symptom cannot be assigned "subordinate" symptom problems; any problems assigned as symptoms to a symptom problem will become symptoms of the same cause problem.

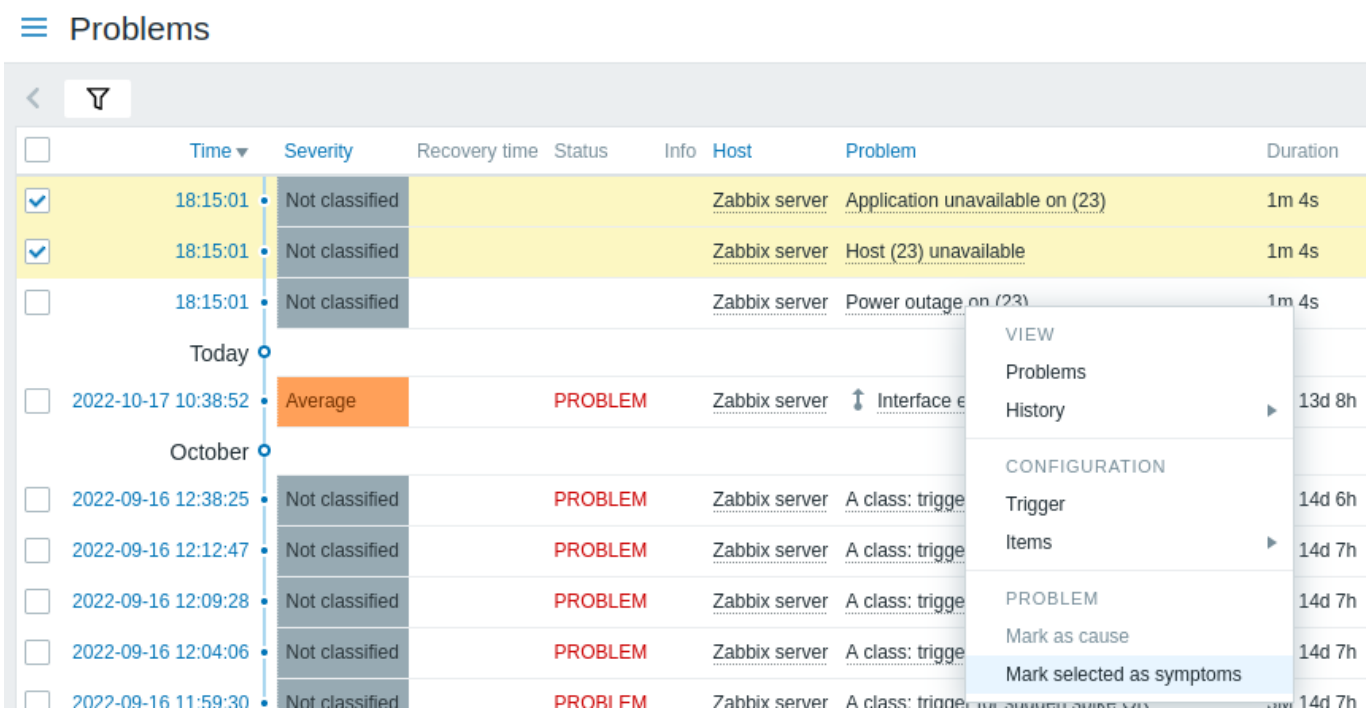
Only cause problems are counted in problem totals in maps, dashboard widgets such as *Problems by severity* or *Problem hosts*, etc. However, problem ranking does not affect services.

A symptom problem can be linked to only one cause problem. Symptom problems are not automatically resolved, if the cause problem is resolved or closed.

Configuration

To reclassify a problem as symptom problem, first select it in the list of problems. One or several problems can be selected.

Then go to the cause problem, and in its context menu click the *Mark selected as symptoms* option.



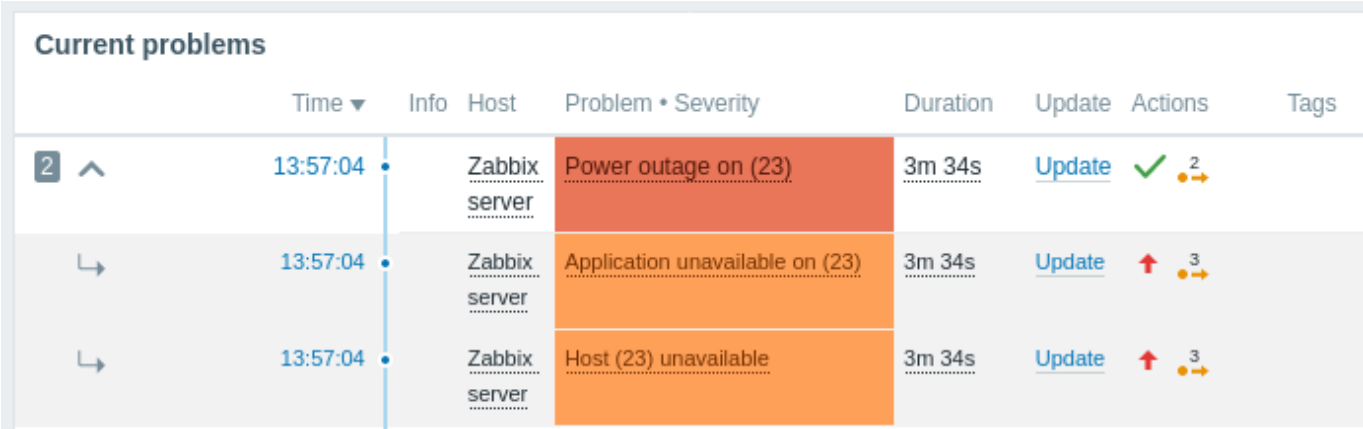
After that, the selected problems will be updated by the server to symptom problems of the cause problem.

While the status of the problem is being updated, it is displayed in one of two ways:

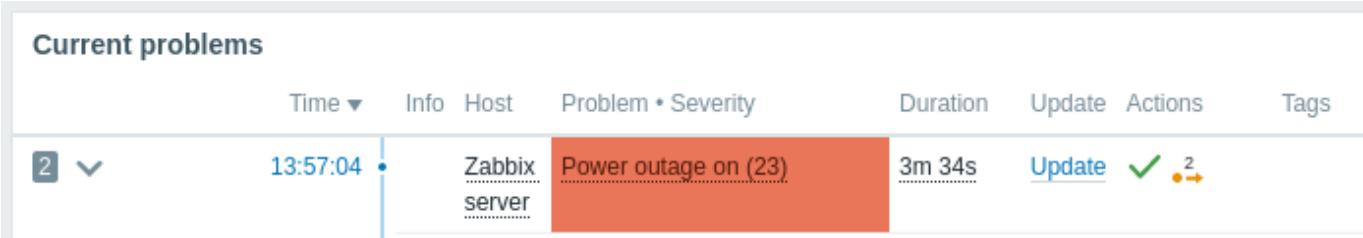
- A blinking "UPDATING" status is displayed in the Status column;
- A blinking or icon in the Info column (this is in effect if *Problems* only are selected in the filter and thus the Status column is not shown).

Display

Symptom problems are displayed below the cause problem and marked accordingly in *Monitoring* > *Problems* (and the *Problems* dashboard widget) - with an icon, smaller font and different background.



In collapsed view, only the cause problem is seen; the existence of symptom problems is indicated by the number in the beginning of the line and the icon for expanding the view.



It is also possible to additionally display symptom problems in normal font and in their own line. For that select *Show symptoms* in the filter settings or the widget configuration.

Reverting to cause problem

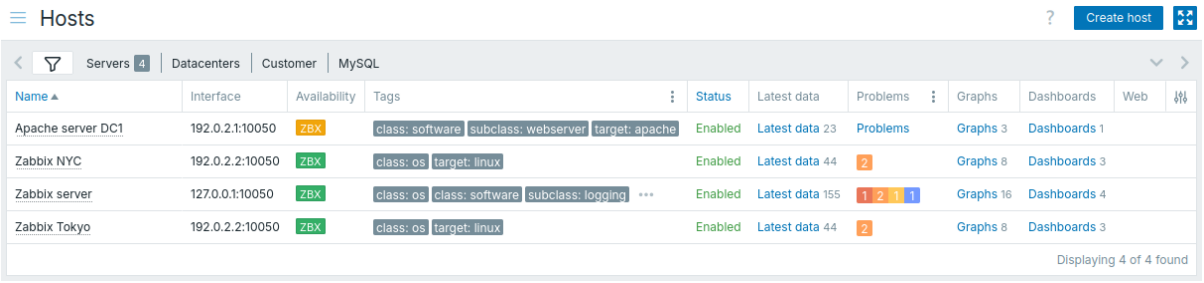
A symptom problem can be reverted back to a cause problem. To do that, either:

- click the *Mark as cause* option in the context menu of the symptom problem;
- mark the *Convert to cause* option in to the problem update screen and click *Update* (this option will also work if several problems are selected).

Hosts

Overview

The *Monitoring* > *Hosts* section displays a full list of monitored hosts with detailed information about host interface, availability, tags, current problems, status (enabled/disabled), and links to easily navigate to the host's latest data, problem history, graphs, dashboards and web scenarios.



Host details are displayed in a customizable table. The *Customize table* button in the table header lets you show, hide, or reorder columns, and reset the layout to default. You can also resize each column by dragging its right border; double-click the border to reset. All customizations are applied immediately and always remembered for the current user within the active filter tab (default or saved).

The following columns are available:

Column	Description
Name	The visible host name. Clicking the name brings up the host menu. An orange wrench icon after the name indicates that this host is in maintenance. Click the column header to sort hosts by name in ascending or descending order.

Column	Description
Interface	The main interface of the host is displayed.
Availability	Host availability per configured interface is displayed. Availability icons represent host interface current status on Zabbix server. Therefore, if you disable a host in the frontend, its availability will update after Zabbix server has synchronized the configuration changes. Similarly, if you enable a host, its availability will update after Zabbix server has synchronized the configuration changes and polled the host. Availability icons represent only those interface types (Agent, SNMP, IPMI, JMX) that are configured. Hovering over the icon displays a pop-up with a list of all interfaces of the same type with details, status, and errors. For Agent interface, the pop-up displays interfaces (passive) and active checks. If a host has active checks only, the Agent interface icon is displayed even if the host does not have an Agent interface configured. The column is empty for hosts with no interfaces. The status of a single host interface is determined by the connection between an enabled item using the interface and the host. The status can be: **Available** - the connection to the host was successful; **Not available** - the connection to the host was unsuccessful (timeout, firewall issues, etc.); **Unknown** - the connection to the host has not been attempted or the result is unknown. For additional details on how Zabbix server determines interface status, see Unknown interface status and Unreachable/unavailable host interface settings. The status of all host interfaces of a single type (Agent, SNMP, IPMI, JMX) is determined by those interfaces that are used by at least one enabled item. The status is indicated by the icon color: **Green** - all interfaces are available; **Yellow** - at least one interface is not available, and at least one is available or unknown; **Red** - all interfaces are not available; **Gray** - at least one interface is unknown, but none are not available. **Active check availability.** If at least one active check is enabled on the host, active checks also affect the total Agent interface availability as described above. To determine active check availability, heartbeat messages are sent in the agent active check thread. The frequency of heartbeat messages is controlled by the `HeartbeatFrequency` parameter in Zabbix agent or agent 2 configuration (default 60 seconds, range 0-3600). Active checks are considered not available when the active check heartbeat is older than 2 x HeartbeatFrequency seconds. **Note:** Zabbix agents older than version 6.2.x do not send active check heartbeats, so the availability of their hosts remains unknown.
Tags	Host tags are displayed. Hovering or clicking the three dots icon after the displayed tags opens a pop-up with all tags. Clicking the *Options* button in the column header lets you set the number of tags displayed (1–3), choose how tag names are displayed (full, shortened, or none), set tag display priority, or duplicate the column (duplicated columns can be renamed or deleted).
Tag value	Host tag values are displayed. Clicking the *Options* button in the column header lets you rename the column, choose which tag values to display, or duplicate the column (duplicated columns can be renamed or deleted).
Status	Host status - *Enabled* or *Disabled*. Click the column header to sort hosts by status in ascending or descending order.
Latest data	Clicking the link will open the *Monitoring* > *Latest data* section with all the latest data collected from the host. The number of items with latest data is displayed in gray.
Problems	The number of open problems for a host is displayed, sorted by severity. The color of each problem icon shows the severity, while the number shows the problem count for that severity. Clicking an icon opens *Monitoring* > *Problems* for that host; if a host has no problems, a *Problems* link is shown instead. Clicking the *Options* button in the column header gives you access to the *Show suppressed problems* option, which includes problems that would otherwise be suppressed (hidden) because of host maintenance or single problem suppression.

Column	Description
Graphs	Clicking the link will display graphs configured for the host. The number of graphs is displayed in gray. If a host has no graphs, no link is displayed.
Dashboards	Clicking the link will display dashboards configured for the host. The number of dashboards is displayed in gray. If a host has no dashboards, no link is displayed.
Web	Clicking the link will display web scenarios configured for the host. The number of web scenarios is displayed in gray. If a host has no web scenarios, no link is displayed.
Custom text	Custom text, macro values, or a combination of both is displayed. Clicking the *Options* button in the column header lets you rename the column, enter the custom text or macros, or duplicate the column (duplicated columns can be renamed or deleted). Supported macros: {HOST.*} (except {HOST.METADATA}), {INVENTORY.*}, and user macros.

Buttons

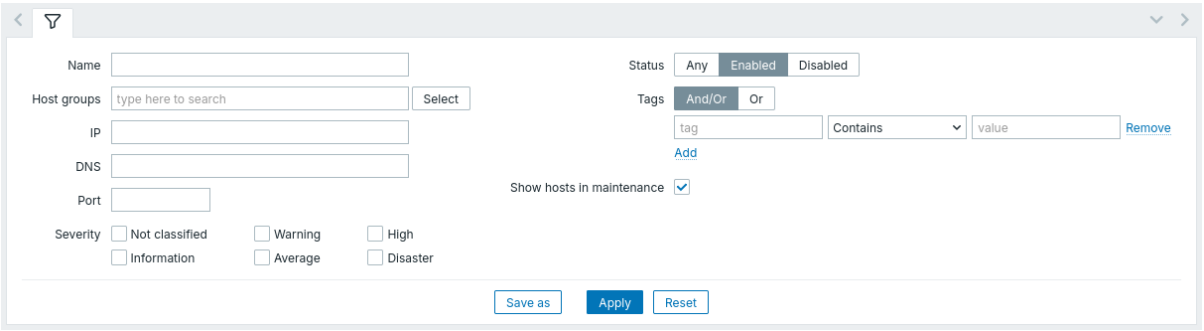
Create host allows you to create a new host. This button is available for Admin and Super Admin users only.

View mode buttons being common for all sections are described on the Monitoring page.

Using filter

You can use the filter to display only the hosts you are interested in. For better search performance, data is searched with macros unresolved.

The filter is located above the table. It is possible to filter hosts by name, host group, IP or DNS, interface port, tags, problem severity, status (enabled/disabled/any); you can also select whether to display hosts that are currently in maintenance.



Parameter	Description
Name	Filter by visible host name.
Host groups	Filter by one or more host groups. Specifying a parent host group implicitly selects all nested host groups.
IP	Filter by IP address.
DNS	Filter by DNS name.
Port	Filter by port number.
Severity	Filter by problem severity. By default, problems of all severities are displayed. Problems are displayed if not suppressed.
Status	Filter by host status.

Parameter	Description
Tags	Filter by host tag name and value. Hosts can be filtered by host-level tags as well as tags from all linked templates, including nested templates. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names **Equals** - include the specified tag names and values (case-sensitive) **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive) **Does not exist** - exclude the specified tag names **Does not equal** - exclude the specified tag names and values (case-sensitive) **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive) There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition **Or** - enough if one condition is met
Show hosts in maintenance	Mark the checkbox to display hosts that are in maintenance (displayed by default).

Saving filter

Favorite filter settings can be saved as tabs and then quickly accessed by clicking the respective tab above the filter.

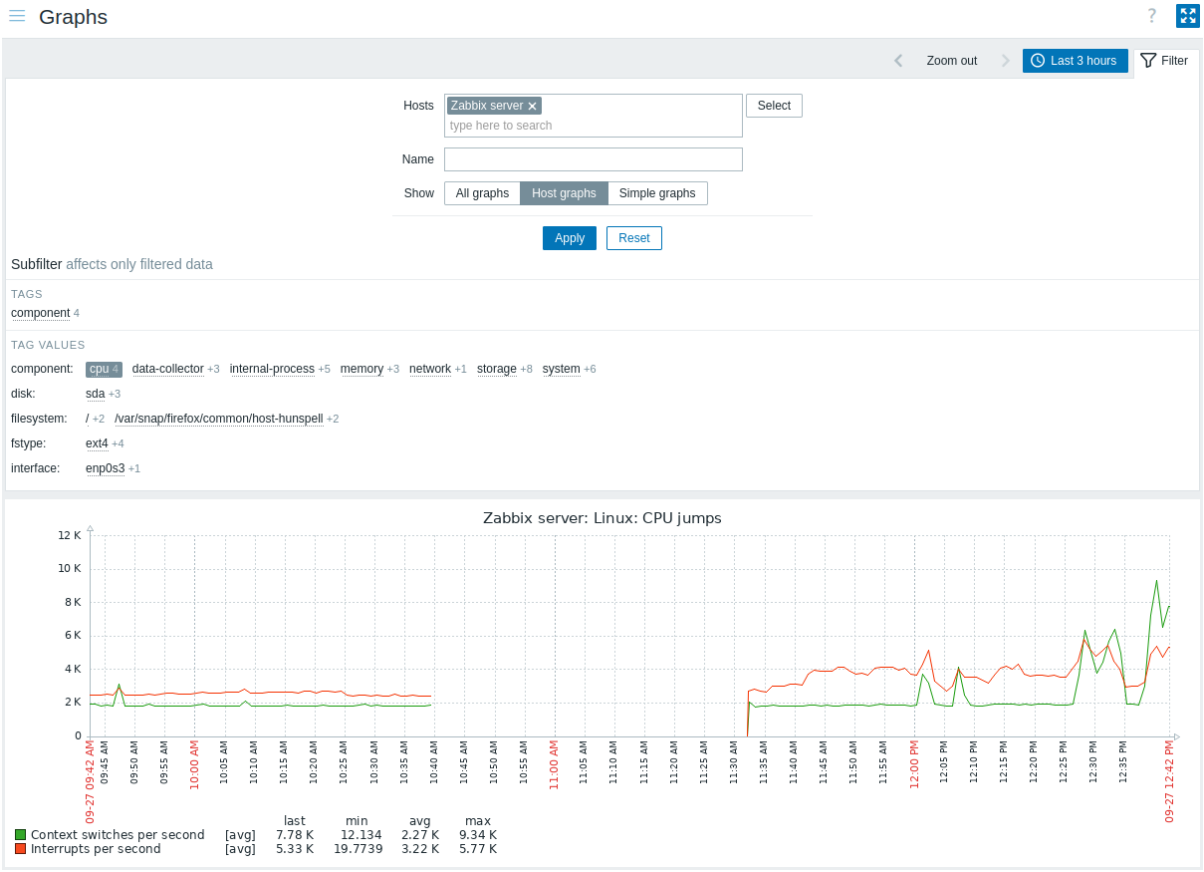
See more details about saving filters.

Graphs

Overview

Host graphs can be accessed from *Monitoring* > *Hosts* by clicking Graphs for the respective host.

Any custom graph that has been configured for the host can be displayed, as well as any simple graph.



Graphs are sorted by:

- graph name (custom graphs)
- item name (simple graphs)

Graphs for disabled hosts are also accessible.

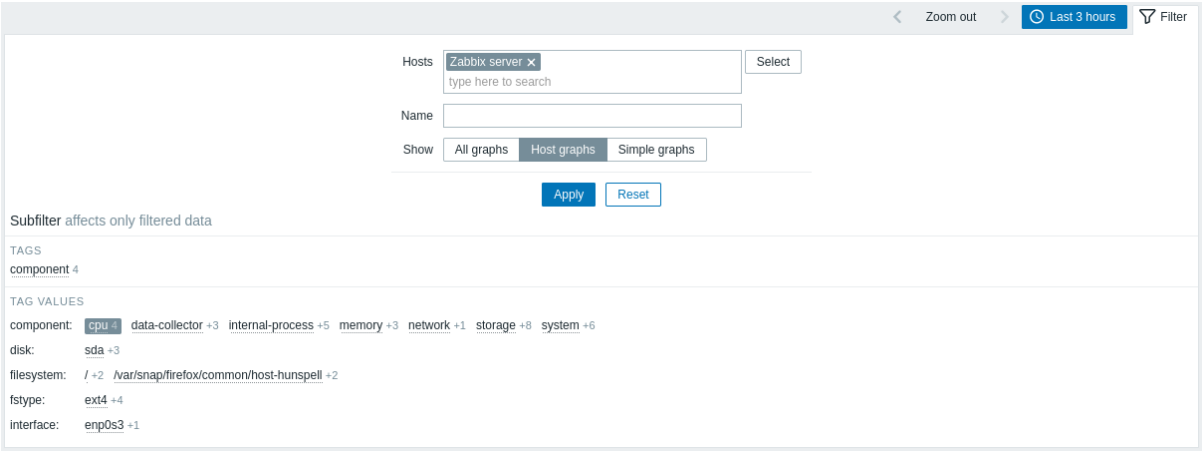
Time period selector

Take note of the time period selector above the graph. It allows selecting often required periods with one mouse click.

See also: *Time period selector*

Using filter

To view a specific graph, select it in the filter. The filter allows you to specify the host, the graph name and the *Show* option (all/host graphs/simple graphs).



If no host is selected in the filter, no graphs are displayed.

Using subfilter

The subfilter allows further modification of the filtering applied by the main filter.

It contains clickable links for a quick access to related graphs. Graphs are related by common entity - tag name or value. When a tag name/value is clicked, it is highlighted with a gray background, and graphs are immediately filtered (no need to click *Apply* in the main filter). Clicking another tag name/value adds it to the filtered results. Clicking the tag name/value again removes the filtering.

Subfilters are generated based on the filtered data, which is limited to 1000 records. If you want to see more records in the subfilter, you need to increase the value of *Limit for search and filter results* parameter (in *Administration* > *General* > *GUI*).

Unlike the main filter, the subfilter is updated together with each table refresh request to always get up-to-date information of available filtering options and their counter numbers.

The number of entities displayed is limited to 100 horizontally. If there are more, a three-dot icon is displayed at the end; it is not clickable. Vertical lists (such as tags with their values) are limited to 20 entries. If there are more, a three-dot icon is displayed; it is not clickable.

A number next to each clickable entity indicates the number of graphs it has in the results of the main filter.

Once one entity is selected, the numbers with other available entities are displayed with a plus sign indicating how many graphs may be added to the current selection.

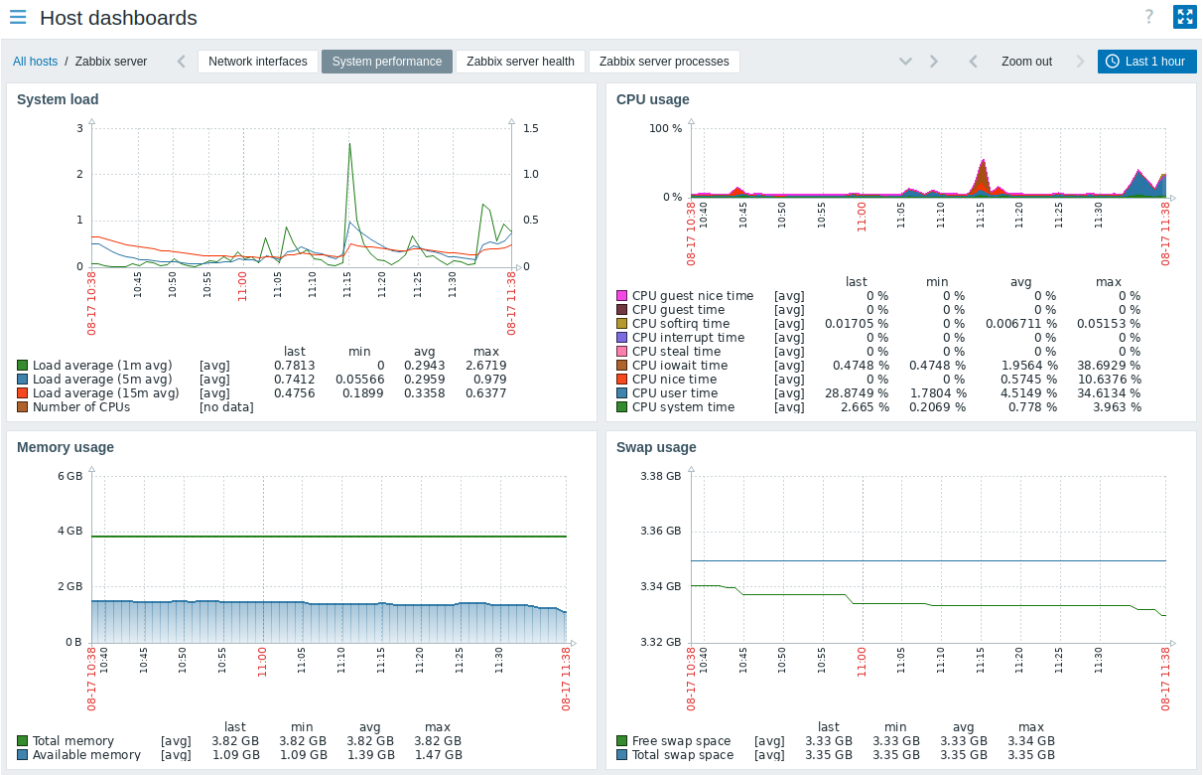
Buttons

View mode buttons, being common for all sections, are described on the Monitoring page.

Host dashboards

Overview

Host dashboards look similar to global dashboards; however, host dashboards lack an owner and display data only for the selected host.



When viewing host dashboards, you can switch between the configured dashboards by clicking:

- the dashboard tabs;
- the  arrow buttons under the header;
- the  arrow button under the header, which will display the full list of host dashboards available.

To switch to the *Monitoring* > *Hosts* section, click the *All hosts* navigation link under the header in the upper-left corner.

Configuration

Host dashboards are configured at the template level. Once a template is linked to a host, host dashboards are generated for that host. Note that host dashboards *cannot* be configured in the *Dashboards* section, which is reserved for global dashboards.

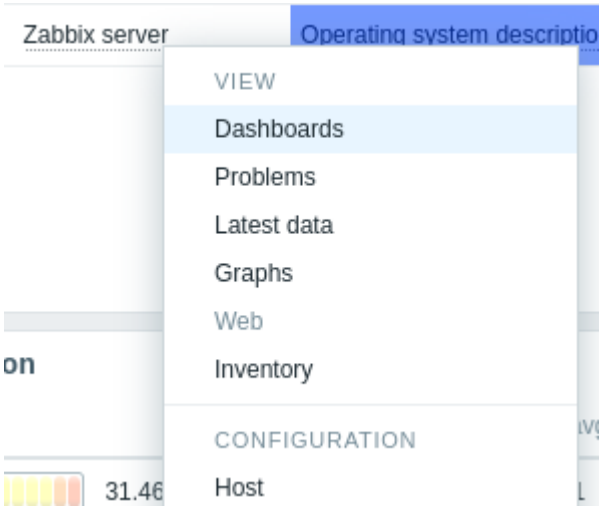
Widgets of host dashboards can also be configured only at the template level, except for changing the refresh interval. Moreover, widgets of host dashboards can only be copied to other host dashboards within the same template. Note that widgets from global dashboards *cannot* be copied to host dashboards.

Note:
Host dashboards used to be host screens before Zabbix 5.2. When importing an older template that contains screens, the screen import will be ignored.

Access

Host dashboards can be accessed:

- after searching for a host name in global search (click the *Dashboards* link provided in the search results);
- after clicking a host name in *Inventory* > *Hosts* (click the *Dashboards* link provided in the host overview);
- from the host menu by clicking *Dashboards*.

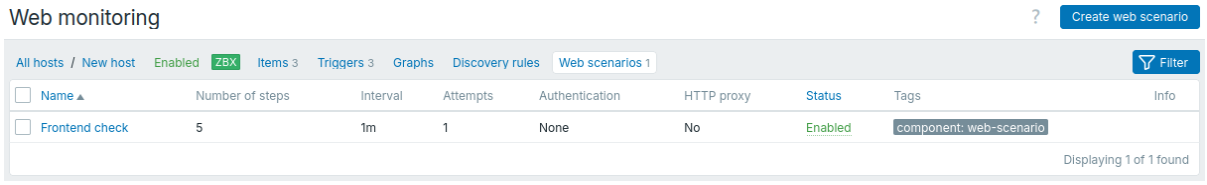


Note that host dashboards *cannot* be directly accessed in the *Dashboards* section, which is reserved for global dashboards.

Web scenarios

Overview

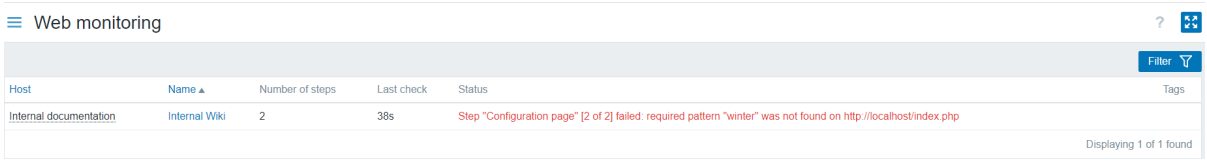
Host web scenario information can be accessed from *Monitoring* > *Hosts* by clicking Web for the respective host.



Clicking the host name brings up the host menu. Data of disabled hosts is also accessible. The name of a disabled host is listed in red.

The status of a web scenario can be:

- OK (green) - all steps have succeeded and Zabbix has collected data for the scenario.
- Error (red) - one or more steps have failed (see the error message for hints on which parameters may need correction).
- Unknown (gray) - no step status is available yet or no data has been collected.



The maximum number of scenarios displayed per page depends on the *Rows per page* user profile setting.

By default, only values that fall within the last 24 hours are displayed. This limit has been introduced with the aim of improving initial loading times for large pages of latest data. You can extend this time period by changing the value of *Max history display period* parameter in the *Administration > General > GUI* menu section.

The scenario name is link to more detailed statistics about it:



Using filter

The page shows a list of all web scenarios of the selected host. To view web scenarios for another host or host group without returning to the *Monitoring* > *Hosts* page, select that host or group in the filter. You may also filter scenarios based on tags.

Buttons

View mode buttons being common for all sections are described on the Monitoring page.

Latest data

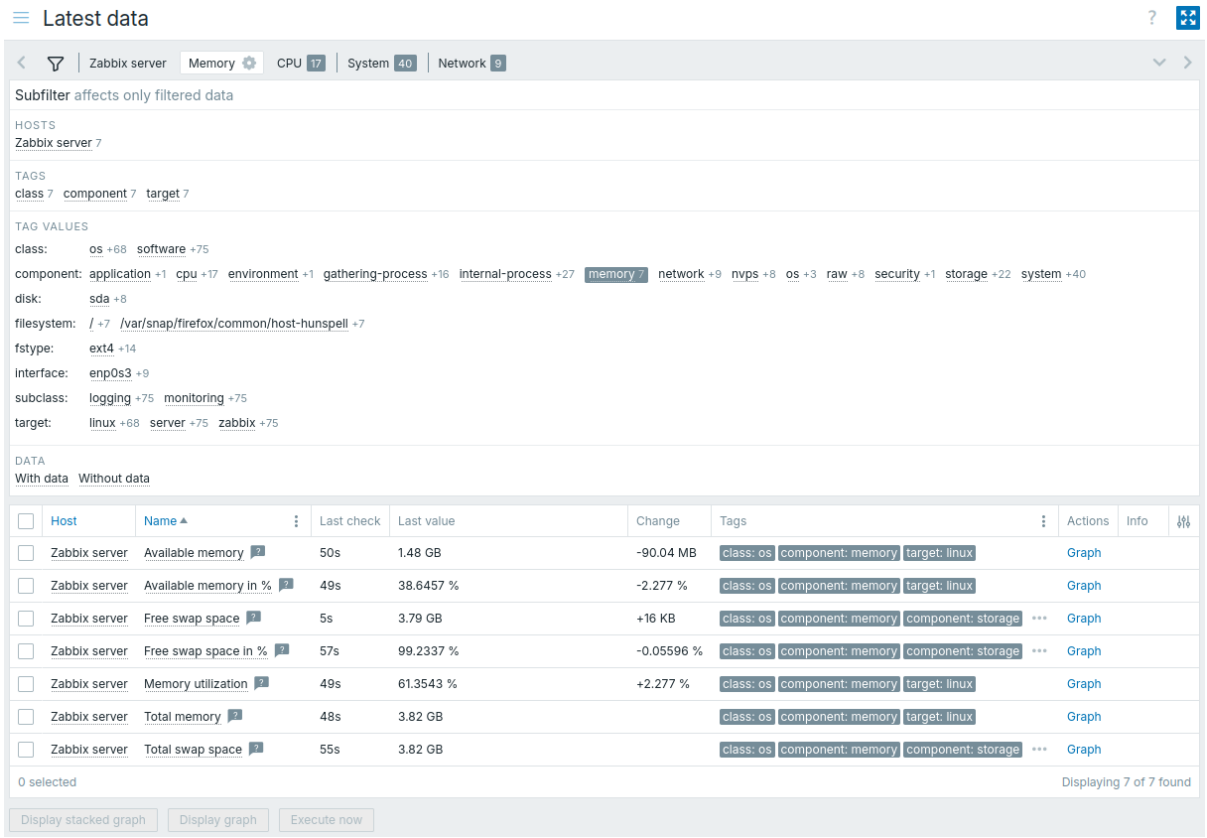
Overview

The *Monitoring* > *Latest data* section displays the latest values gathered by items.

This section contains the following elements:

- Filter
- Subfilter
- Item list

Note:
The subfilter and item list is displayed only if the filter is set and there are results to display.



Latest value details are displayed in a customizable table. The *Customize table* button in the table header lets you show, hide, or reorder columns, and reset the layout to default. You can also resize each column by dragging its right border; double-click the border to reset. All customizations are applied immediately and always remembered for the current user within the active filter tab (default or saved).

The following columns are available:

Column	Description
Host	Name of the host to which the item belongs. Clicking the name brings up the host context menu. If a host is in maintenance, an orange wrench icon is displayed after the host's name. If a host is disabled, the name of the host is displayed in red. Note that data of disabled hosts (including graphs and item value lists) is accessible in the *Latest data* section.
Name	Name of the item. Clicking the name brings up the item menu. A question mark icon is displayed next to the item name for all items that have a description. Hover over the icon to display a tooltip with the item description. Clicking the *Options* button in the column header gives you access to the *Show item key* option that displays the item key below each item's name.
Interval	Update interval of the item.
History	History storage period of the item.
Trends	Trends storage period of the item.
Type	Type of the item.
Last check	Time since the last item check.
Last value	Most recent value for the item. Values are displayed with unit conversion and value mapping applied. Hover over the value to display raw data (truncated to 8192 characters). Binary values are displayed as a placeholder (*binary value*) instead of the actual value. By default, only values received in the last 24 hours are displayed. This limit improves initial loading times for large pages of latest data; to extend it, update the value of the *Max history display period* parameter in *Administration* > *General* > *GUI*.
Change	Difference between the previous value and the most recent value. For items with an update frequency of 1 day or more, the change amount will never be displayed (with the default setting). In this case, the last value will not be displayed at all if it was received more than 24 hours ago.

Column	Description
Tags	Tags associated with the item. Clicking a tag enables it in the subfilter, thus filtering the item list by that tag (and any other tags selected in the subfilter). Once clicked, such tags are no longer clickable; to adjust filtering by tags, the subfilter must be used. Clicking the *Options* button in the column header lets you set the number of tags displayed (1-3), choose how tag names are displayed (full, shortened, or none), set tag display priority, or duplicate the column (duplicated columns can be renamed or deleted).
Tag value	Item tag values are displayed. Clicking the *Options* button in the column header lets you rename the column, choose which tag values to display, or duplicate the column (duplicated columns can be renamed or deleted).
Actions	Link to simple graph/history of item values.
Info	Additional information about the item. If an item has errors (for example, has become unsupported), an information icon is displayed. Hover over the icon for details.
Custom text	Custom text, macro values, or a combination of both is displayed. Clicking the *Options* button in the column header lets you rename the column, enter the custom text or macros, or duplicate the column (duplicated columns can be renamed or deleted). Supported macros: {HOST.*} (except {HOST.METADATA}), {INVENTORY.*}, {ITEM.*} (except {ITEM.STATE.ERROR}), and user macros.

Buttons

View mode buttons being common for all sections are described on the Monitoring page.

Mass actions

Buttons below the list offer mass actions with one or several selected items:

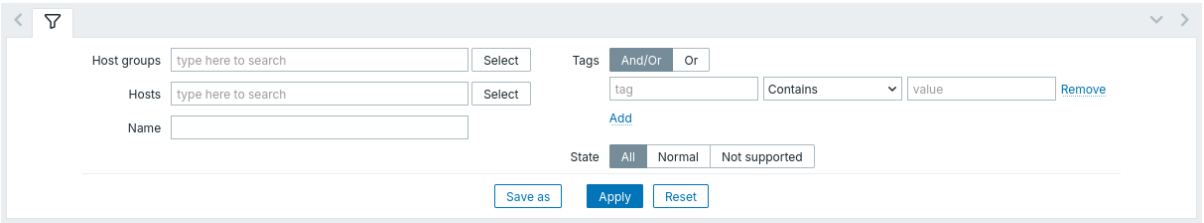
- *Display stacked graph* - display a stacked ad-hoc graph.
- *Display graph* - display a simple ad-hoc graph.
- *Execute now* - execute a check for new item values immediately. Supported for **passive** checks only (see more details). This option is available only for hosts with read-write access. Accessing this option for hosts with read-only permissions depends on the user role option called *Invoke "Execute now" on read-only hosts*.

To use these options, mark the checkboxes before the respective items, then click the required button.

Using filter

You can use the filter to display only the items you are interested in. For better search performance, data is searched with macros unresolved.

The filter icon is located above the item list and the subfilter. Click it to expand the filter.



The filter allows you to narrow the item list by host group, host, item name, tags, and state. Specifying a parent host group in the filter implicitly selects all nested host groups. See *Monitoring* > *Problems* for details on filtering by tags.

The *Show details* filter option allows the information displayed for the items to be expanded. Mark it to display such details as the item refresh interval, history and trends settings, item type, and item errors (fine/unsupported).

Saving filter

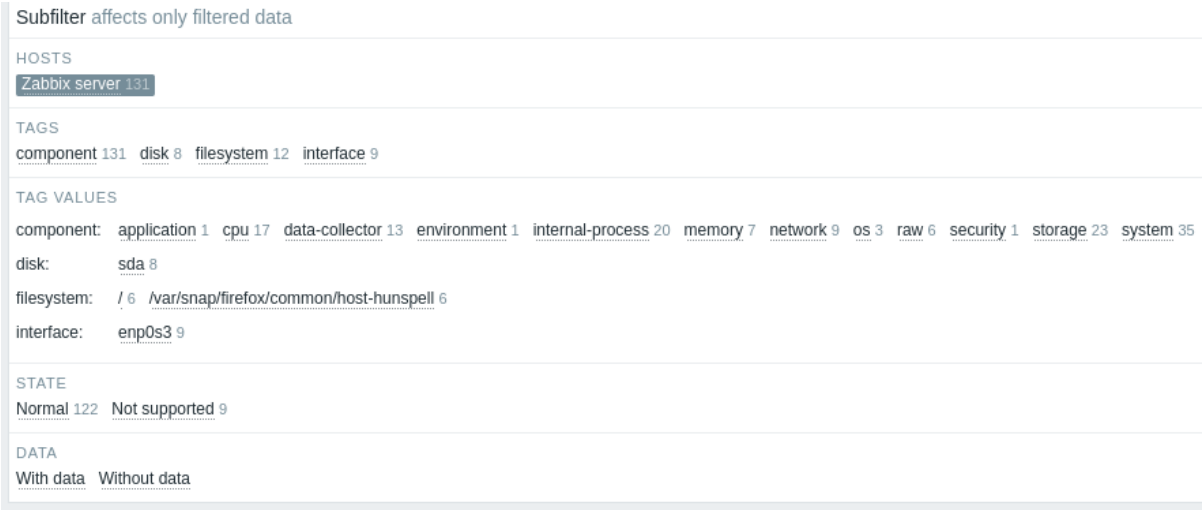
Favorite filter settings can be saved as tabs and then quickly accessed by clicking the respective tab above the filter.

See more details about saving filters.

Using subfilter

The subfilter allows you to further modify the filtering from the main filter.

It contains clickable links for a quick access to related items. Items are related by common entity - host, tag name or value, item state, or data status. When an entity is clicked, the entity is highlighted with a gray background, and items are immediately filtered (no need to click *Apply* in the main filter). Clicking another entity adds it to the filtered results. Clicking the entity again removes the filtering.



Subfilters are generated based on the filtered data, which is limited to 1000 records. If there are 20 hosts each having 100 items (so 2000 records in total), only half of the hosts will be visible in the subfilter. If you want to see more records in the subfilter, you need to increase the value of *Limit for search and filter results* parameter (in *Administration* > *General* > *GUI*).

Unlike the main filter, the subfilter is updated with each table refresh request to always have up-to-date information of available filtering options and their counter numbers.

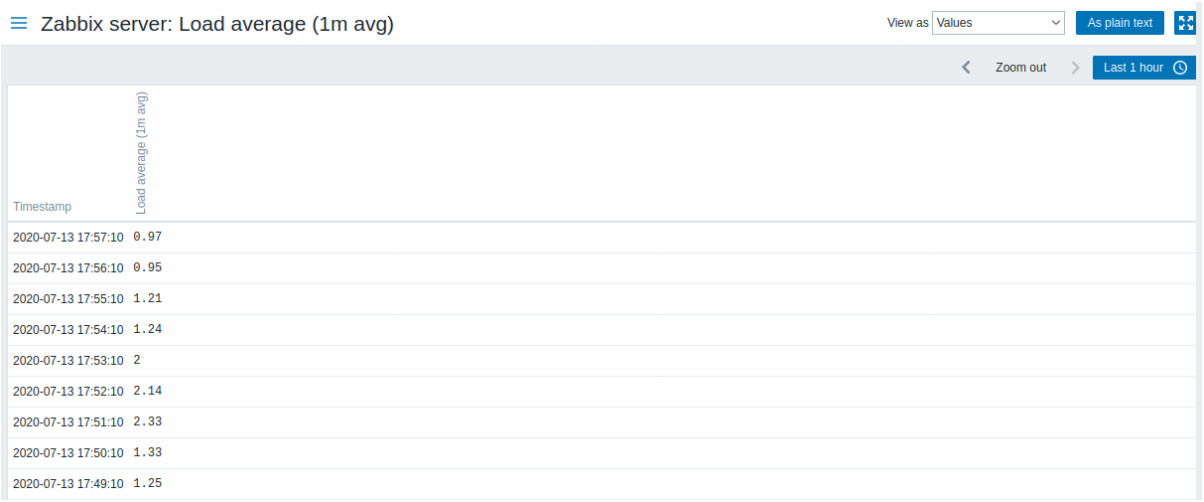
For each entity group (hosts, tags, tag values, etc.), up to 10 rows of entities are displayed. If there are more entities, this list can be expanded to display a maximum of 1000 entries (the value of SUBFILTER_VALUES_PER_GROUP in frontend definitions) by clicking the three-dot icon at the end of the list. For *Tag values*, the list can be expanded to display a maximum of 200 tag names with their corresponding values. Note that once fully expanded, the list cannot be collapsed.

A number next to each clickable entity indicates the number of items grouped in it (based on the results of the main filter). When an entity is clicked, the numbers with other available entities are displayed with a plus sign indicating how many items may be added to the current selection. Entities without items are not displayed unless selected in the subfilter before.

Graphs and history

The *Graph/History* column in the item list offers the following links:

- **History** - for all textual items, leading to listings (*Values/500 latest values*) displaying the history of previous item values.
- **Graph** - for all numeric items, leading to a simple graph. Note that when the graph is displayed, a dropdown on the upper right offers a possibility to switch to *Values/500 latest values* as well.



The values displayed in this list are raw, that is, no postprocessing is applied.

Note:
The total amount of values displayed is defined by the value of *Limit for search and filter results* parameter, set in *Administration > General > GUI*.

Maps

Overview

In the *Monitoring > Maps* section you can configure, manage and view network maps.

When you open this section, you will either see the last map you accessed or a listing of all maps you have access to.

All maps can be either public or private. Public maps are available to all users, while private maps are accessible only to their owner and the users the map is shared with.

Map listing



Displayed data:

Column	Description
Name	Name of the map. Click the name to view the map.
Width	Map width is displayed.
Height	Map height is displayed.
Actions	Two actions are available: **Properties** - set general map properties **Edit** - access the grid for adding map elements

To configure a new map, click the *Create map* button in the upper-right corner. To import a map from a YAML, XML, or JSON file, click the *Import* button in the upper-right corner. The user who imports the map will be set as its owner.

Two buttons below the list offer some mass-editing options:

- *Export* - export the maps to a YAML, XML, or JSON file
- *Delete* - delete the maps

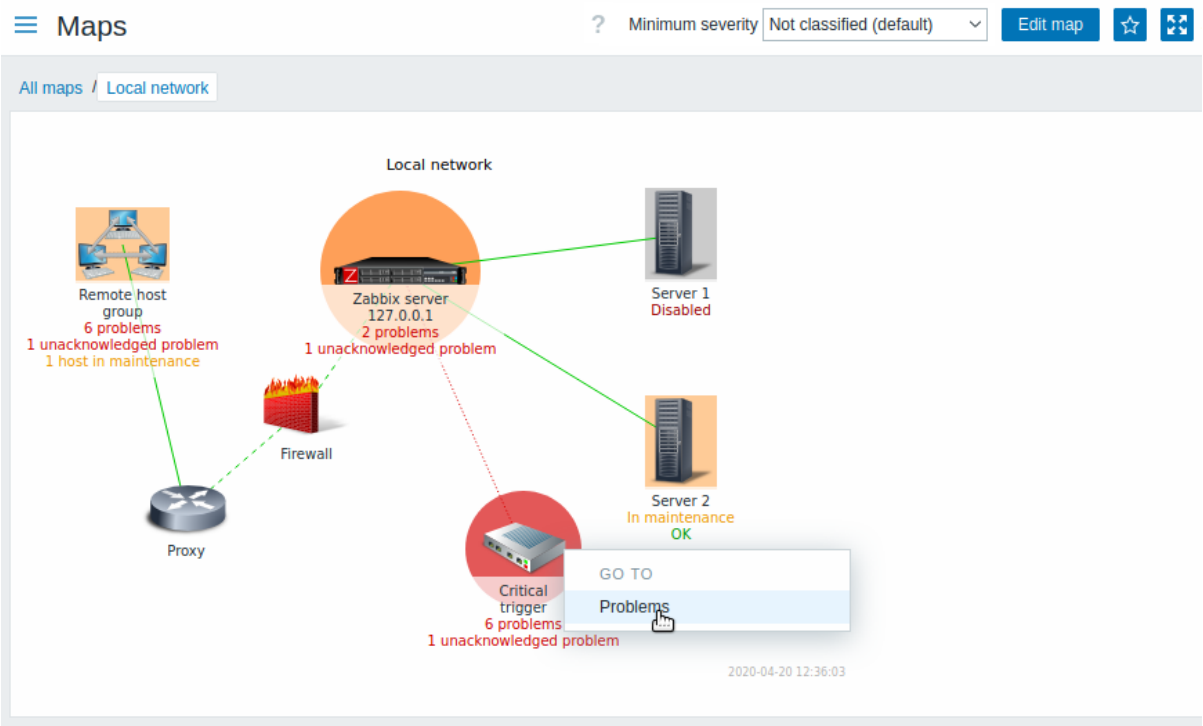
To use these options, mark the checkboxes before the respective maps, then click the required button.

Using filter

You can use the filter to display only the maps you are interested in. For better search performance, data is searched with macros unresolved.

Viewing maps

To view a map, click its name in the list of all maps.



You can use the drop-down in the map title bar to select the lowest severity level of the problem triggers to display. The severity marked as *default* is the level set in the map configuration. If the map contains a sub-map, navigating to the sub-map will retain the higher-level map severity (except if it is *Not classified*, in this case, it will not be passed to the sub-map).

Icon highlighting

If a map element is in problem state, it is highlighted with a round circle. The fill color of the circle corresponds to the severity color of the problem. Only problems on or above the selected severity level will be displayed with the element. If all problems are acknowledged, a thick green border around the circle is displayed.

Additionally:

- a host in maintenance is highlighted with an orange, filled square. Note that maintenance highlighting has priority over the problem severity highlighting, if the map element is host.
- a disabled (not-monitored) host is highlighted with a gray, filled square.

Highlighting is displayed if the *Icon highlighting* check-box is marked in the map configuration.

Recent change markers

Inward pointing red triangles around an element indicate a recent change of trigger state - one that's happened within the last 30 minutes. These triangles are shown if the *Mark elements on trigger status change* check-box is marked in the map configuration.

Links

Clicking a map element opens a menu with some available links. Clicking the host name brings up the host menu.

Buttons

Buttons to the right offer the following options:

Edit map	Go to editing of the map content.
	Add map to the favorites widget in Dashboards.
	The map is in the favorites widget in Dashboards. Click to remove map from the favorites widget.

View mode buttons being common for all sections are described on the Monitoring page.

Readable summary in maps

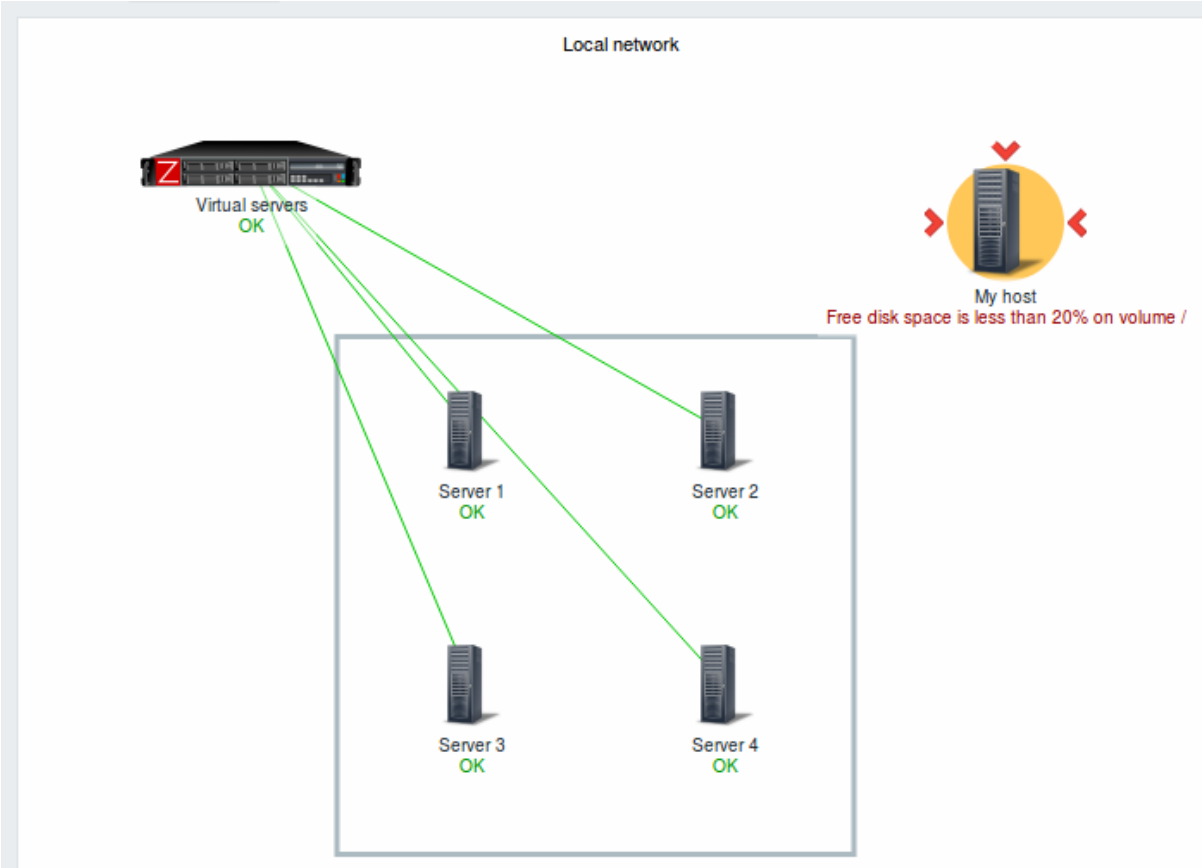
A hidden "aria-label" property is available, allowing map information to be read with a screen reader. Both general map description and individual element description is available, in the following format:

- for map description: `<Map name>, <* of * items in problem state>, <* problems in total>.`
- for describing one element with one problem: `<Element type>, Status <Element status>, <Element name>, <Problem description>.`
- for describing one element with multiple problems: `<Element type>, Status <Element status>, <Element name>, <* problems>.`
- for describing one element without problems: `<Element type>, Status <Element status>, <Element name>.`

For example, this description is available:

```
'Local network, 1 of 6 elements in problem state, 1 problem in total. Host, Status problem, My host, Free
```

for the following map:



Referencing a network map

Network maps can be referenced by both `sysmapid` and `mapname` GET parameters. For example,

```
http://zabbix/zabbix/zabbix.php?action=map.view&mapname=Local%20network
```

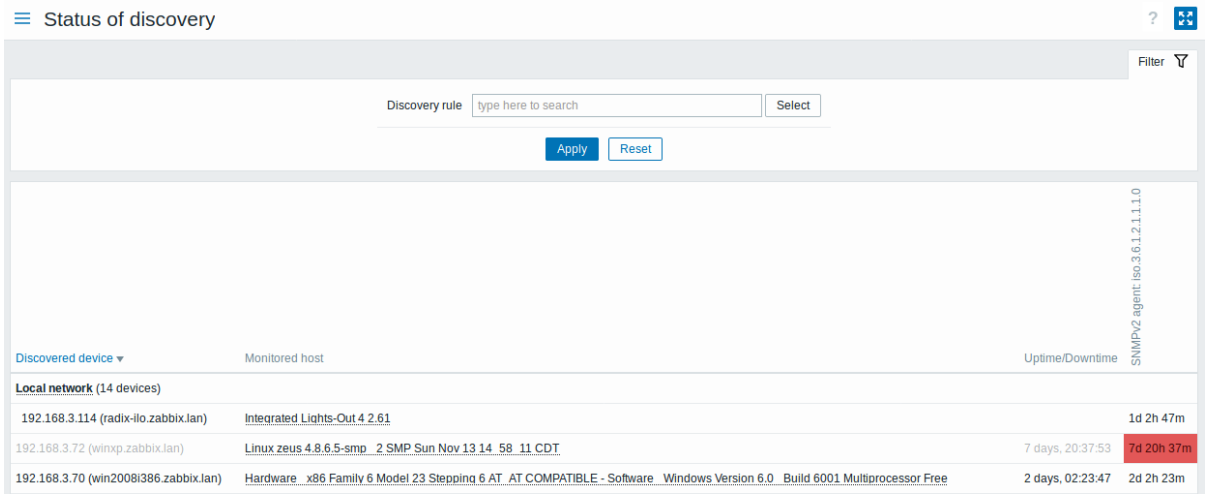
will open the map with that name (Local network).

If both `sysmapid` (map ID) and `mapname` (map name) are specified, `mapname` has higher priority.

Discovery

Overview

In the *Monitoring* > *Discovery* section results of network discovery are shown. Discovered devices are sorted by the discovery rule.



Displayed data:

Column	Description
Discovered device	Discovered devices are listed, grouped by the discovery rule. Clicking the discovery rule brings up the rule menu containing the link to the discovery rule configuration form. If a device has several IP addresses (as determined by the *Device uniqueness criteria* of the discovery rule), they are listed together: the primary IP address of the device first, followed by its other IP addresses (indented). The primary IP address is the lowest IP address that has a host listed in the *Monitored host* column or, if there is no such address, the lowest IP address of the device. Devices can be sorted by their primary IP address in ascending or descending order; the other IP addresses of a device are always listed below its primary IP address, in ascending order.
Monitored host	If a device is already monitored, the host name will be listed in this column. Clicking the host name brings up the host menu.
Uptime/Downtime	The duration of the device being discovered or lost after previous discovery is displayed in this column. For a device with several IP addresses, the duration is displayed only for its primary IP address.
Discovery check	The state of the individual service (discovery check) for each discovered device is displayed. A red cell shows that the service is down. Service uptime or downtime is included within the cell. This column is displayed only if the service has been found on at least one discovered device.

Note:
The number of discovery rules and the total number of discovered devices displayed in this section are limited by the *Limit for search and filter results* parameter in *Administration > General > GUI*. Rules or devices exceeding the limit are not displayed, and no indication of this is shown.

Buttons

View mode buttons being common for all sections are described on the Monitoring page.

Using filter

You can use the filter to display only the discovery rules you are interested in. For better search performance, data is searched with macros unresolved.

With nothing selected in the filter, all enabled discovery rules are displayed. To select a specific discovery rule for display, start typing its name in the filter. All matching enabled discovery rules will be listed for selection. More than one discovery rule can be selected.

APM

Overview

The *APM* menu is for application performance monitoring (APM) features of Zabbix. It displays telemetry data that your applications and services send using OpenTelemetry, an open standard for collecting telemetry data. This data helps you see how requests flow through your applications and services, measure their performance, and find the cause of slow or failed operations.

The menu contains the following sections that correspond to OpenTelemetry signal types:

- Traces
- Metrics
- Logs

Unlike *Monitoring* > *Latest data*, which displays values collected by Zabbix items, the *APM* sections display telemetry data directly from the data source, so you don't need to create any items. To display data in the *APM* sections:

1. Specify a telemetry data source in *Administration* > *Data source*.
2. Make sure that telemetry data is written to the data source by Zabbix proxy with APM enabled, or by an OpenTelemetry collector. Note that Zabbix server cannot receive telemetry data; only Zabbix proxy can.

To use telemetry data in triggers or graphs, use the telemetry query item.

Traces

Overview

The *APM* > *Traces* section displays traces collected from your applications and services.

A **trace** represents the path that a request takes through your applications and services. For example, when a user clicks the *Place order* button in an online shop, the request can pass from the frontend service, through the checkout and payment services, to the database.

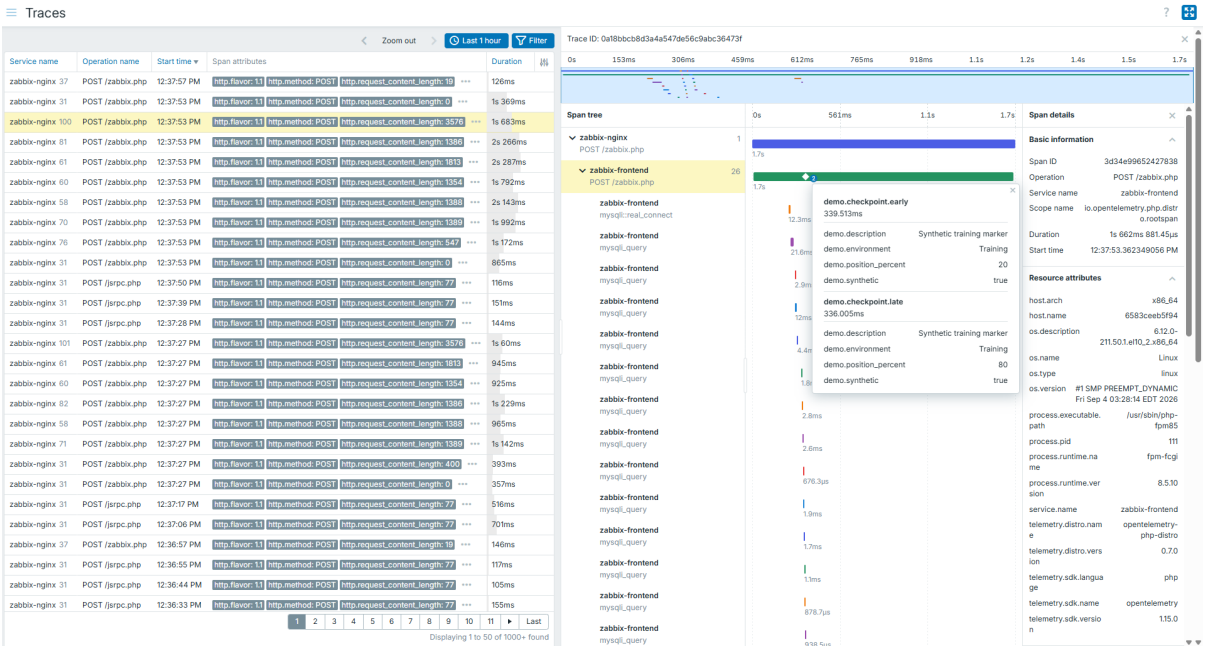
A trace consists of **spans**. Each span represents a single operation, such as an HTTP request that the frontend service sends to the checkout service, a database query that saves the order, or a function call that calculates the shipping cost.

Spans form a hierarchy. The **root span** represents the first operation of the request. Every other span is a **child span** of the span during which its operation was started. For example, the frontend service receiving the *Place order* request is the root span, and the checkout service processing the order is its child span.

Each span is produced by a **service** (e.g., `zabbix-frontend`; named by the OpenTelemetry `service.name` resource attribute), and by an instrumentation scope, the part of the code that created it (e.g., a library that traces HTTP requests).

To provide context about each operation, spans also have key-value **attributes**: span attributes describe the operation itself (e.g., `http.method: POST`), and resource attributes describe the service and the environment it runs in (e.g., `host.name: eabf8c2a5ddc`).

For more details on traces, see the OpenTelemetry documentation.

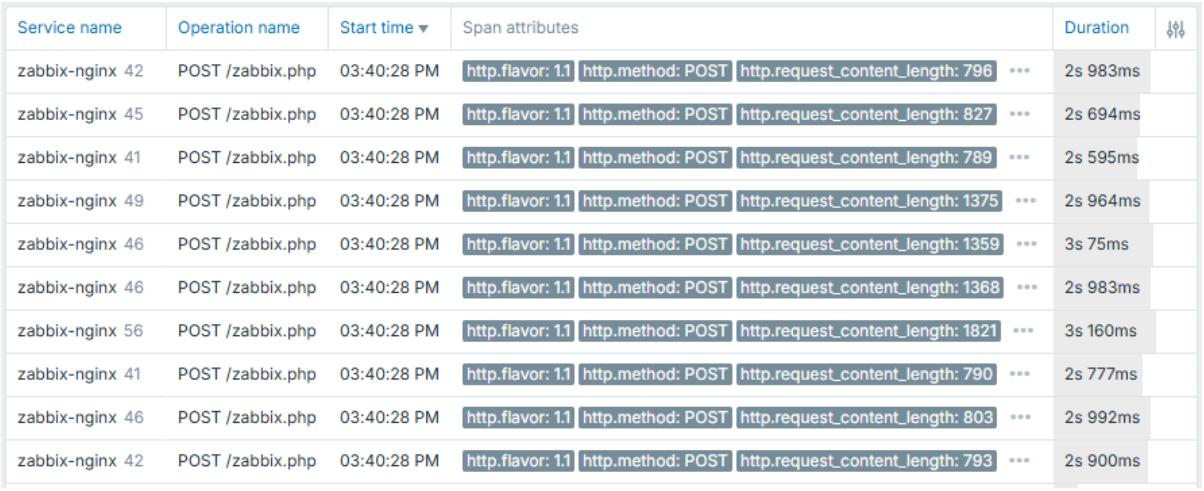


The *APM* > *Traces* section contains the following elements:

- Kiosk mode button - display only the page content, without the menu and header.
- Time period selector - select the time period for which to display traces.
- Filter - display only the traces you are interested in.
- Trace list - view the traces, described by their root spans.
- Trace panel - view all spans of the selected trace on a timeline.

Trace list

Each row in the trace list represents a trace; the columns show the data of its root span.



Service name	Operation name	Start time ▾	Span attributes	Duration
zabbix-nginx 42	POST /zabbix.php	03:40:28 PM	http.flavor: 1.1 http.method: POST http.request_content_length: 796 •••	2s 983ms
zabbix-nginx 45	POST /zabbix.php	03:40:28 PM	http.flavor: 1.1 http.method: POST http.request_content_length: 827 •••	2s 694ms
zabbix-nginx 41	POST /zabbix.php	03:40:28 PM	http.flavor: 1.1 http.method: POST http.request_content_length: 789 •••	2s 595ms
zabbix-nginx 49	POST /zabbix.php	03:40:28 PM	http.flavor: 1.1 http.method: POST http.request_content_length: 1375 •••	2s 964ms
zabbix-nginx 46	POST /zabbix.php	03:40:28 PM	http.flavor: 1.1 http.method: POST http.request_content_length: 1359 •••	3s 75ms
zabbix-nginx 46	POST /zabbix.php	03:40:28 PM	http.flavor: 1.1 http.method: POST http.request_content_length: 1368 •••	2s 983ms
zabbix-nginx 56	POST /zabbix.php	03:40:28 PM	http.flavor: 1.1 http.method: POST http.request_content_length: 1821 •••	3s 160ms
zabbix-nginx 41	POST /zabbix.php	03:40:28 PM	http.flavor: 1.1 http.method: POST http.request_content_length: 790 •••	2s 777ms
zabbix-nginx 46	POST /zabbix.php	03:40:28 PM	http.flavor: 1.1 http.method: POST http.request_content_length: 803 •••	2s 992ms
zabbix-nginx 42	POST /zabbix.php	03:40:28 PM	http.flavor: 1.1 http.method: POST http.request_content_length: 793 •••	2s 900ms

Traces are displayed in a customizable table. The *Customize table* button in the table header lets you show, hide, or reorder columns, and reset the layout to default. You can also resize each column by dragging its right border; double-click the border to reset. All customizations are applied immediately and remembered for the current user.

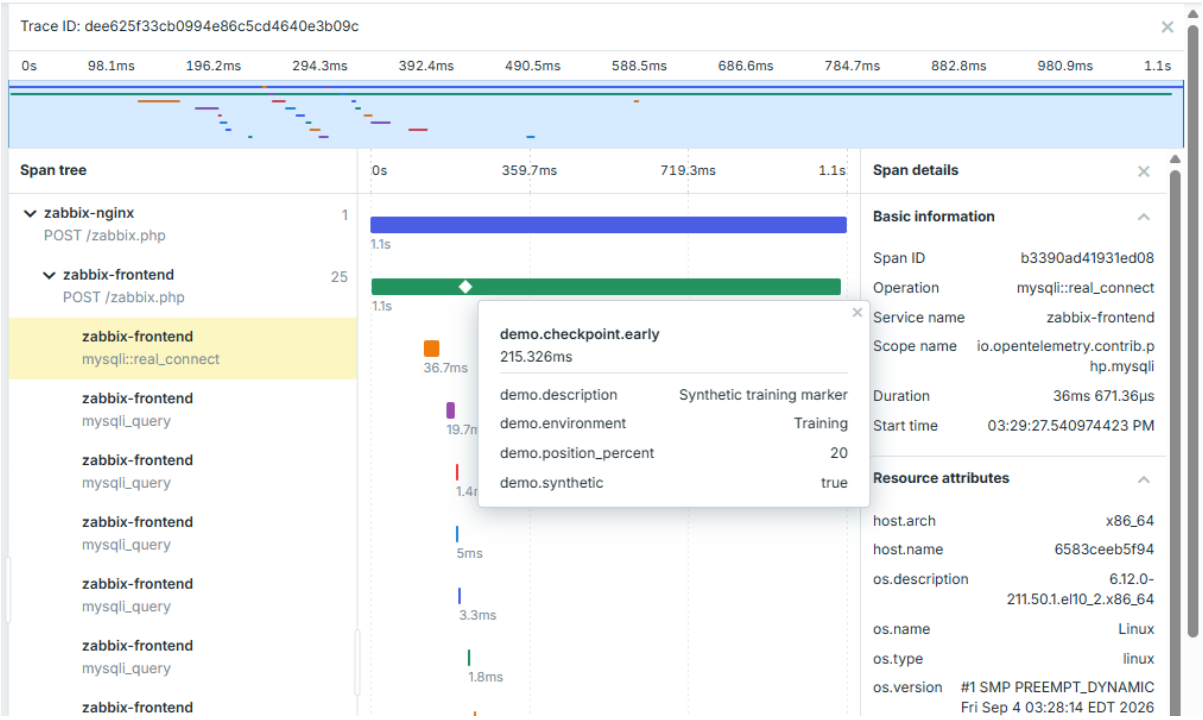
The trace list can have the following columns:

Note:
Each column corresponds to a column in the `otel_traces` table of the telemetry data source; for the list of columns, see Supported telemetry data columns. Most columns have the same name as in the data source (e.g., *Trace ID* corresponds to the `TraceId` column). Where the names differ, the data source column is specified in the table below.

Column	Description
Trace ID	ID of the trace, shared by all spans that belong to the trace.
Span ID	ID of the root span.
Trace state	Additional data, if any, that tracing tools attach to the trace (see W3C Trace Context).
Service name	Service that produced the root span, followed by the number of spans in the trace. If any spans have the *Error* status code (see below), a red icon shows the number of them.
Operation name	Operation that the root span represents (e.g., `POST /zabbix.php`). Data source column: `SpanName`.
Span kind	Role of the root span: **Unspecified** - not specified. **Internal** - operation performed within an application. **Server** - handling of an incoming request. **Client** - outgoing request. **Producer** - creation of an asynchronous operation (e.g., sending a message). **Consumer** - processing of an asynchronous operation.
Scope name	Instrumentation scope that produced the root span.
Scope version	Version of the instrumentation scope.
Status code	Status of the operation that the root span represents: **Unset** - no explicit status has been assigned to the root span. **Ok** - the operation was completed successfully. **Error** - the operation was completed with an error.
Status message	Additional information, if any, about the status.
Start time	Date and time when the root span started. Data source column: `Timestamp`.
Span attributes	Span attributes of the root span. Click the ••• three-dot button to display all attributes.
Resource attributes	Resource attributes of the root span. Click the ••• three-dot button to display all attributes.
Duration	Duration of the trace; the gray bar shows it relative to the other traces.

Trace panel

Click a trace in the trace list to open the trace panel, which shows all spans of the trace on a timeline: which operations were performed, in what order, and how long each took.

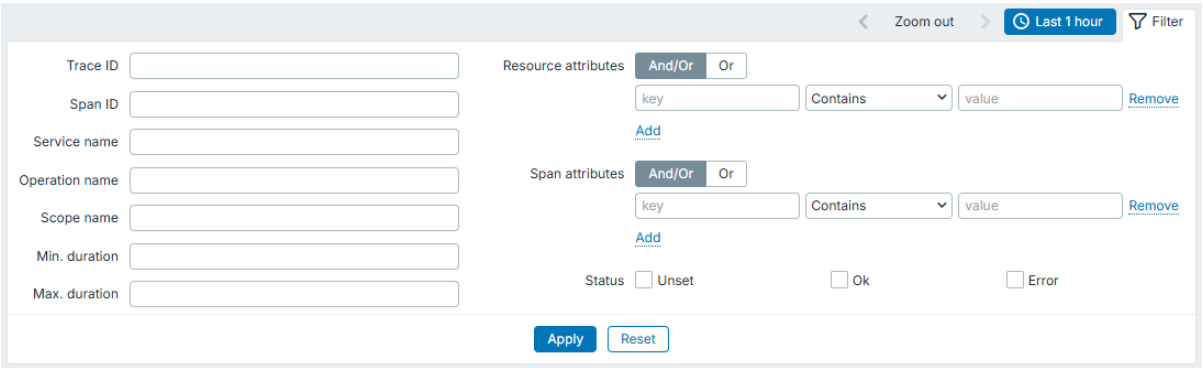


The trace panel can have the following elements:

Element	Description
Trace ID	ID of the selected trace.
Trace timeline overview	The whole trace, with spans in up to 8 rows (several spans can share a row). Drag the borders of the selected range, or the range itself, to zoom the span timeline to that part of the trace.
Span tree	Spans in their hierarchy; each shows its service (in bold), operation name, and number of direct child spans.
Span timeline	Spans as bars: the position and length show when each span started and how long it took. Markers show events (e.g., exceptions); hover over or click a marker for details.
Span details	Click a span in the span tree to see its details and attributes.

Using filter

You can use the filter to display only the traces you are interested in.



To filter by attributes, enter the attribute key, select the operator, and enter the attribute value.

The following operators are supported:

- **Contains** - *(default)* the attribute value contains the entered string.
- **Equals** - the attribute value equals the entered string.
- **Does not contain** - the attribute value does not contain the entered string.
- **Does not equal** - the attribute value does not equal the entered string.
- **Exists** - the attribute with the entered key exists.

- **Does not exist** - the attribute with the entered key does not exist.

If you add several conditions, select how to evaluate them:

- **And/Or** - *(default)* all conditions must be met; conditions with the same attribute key are combined with *Or*.
- **Or** - at least one condition must be met.

Metrics

Overview

The *APM* > *Metrics* section displays metric data points collected from your applications and services.

A **metric** is a measurement (e.g., request duration, number of processed orders, or memory usage). Applications send their values as **data points** that contain the value (or aggregated values) measured at a specific time. A metric is similar to a Zabbix item, and a data point is similar to a value of that item.

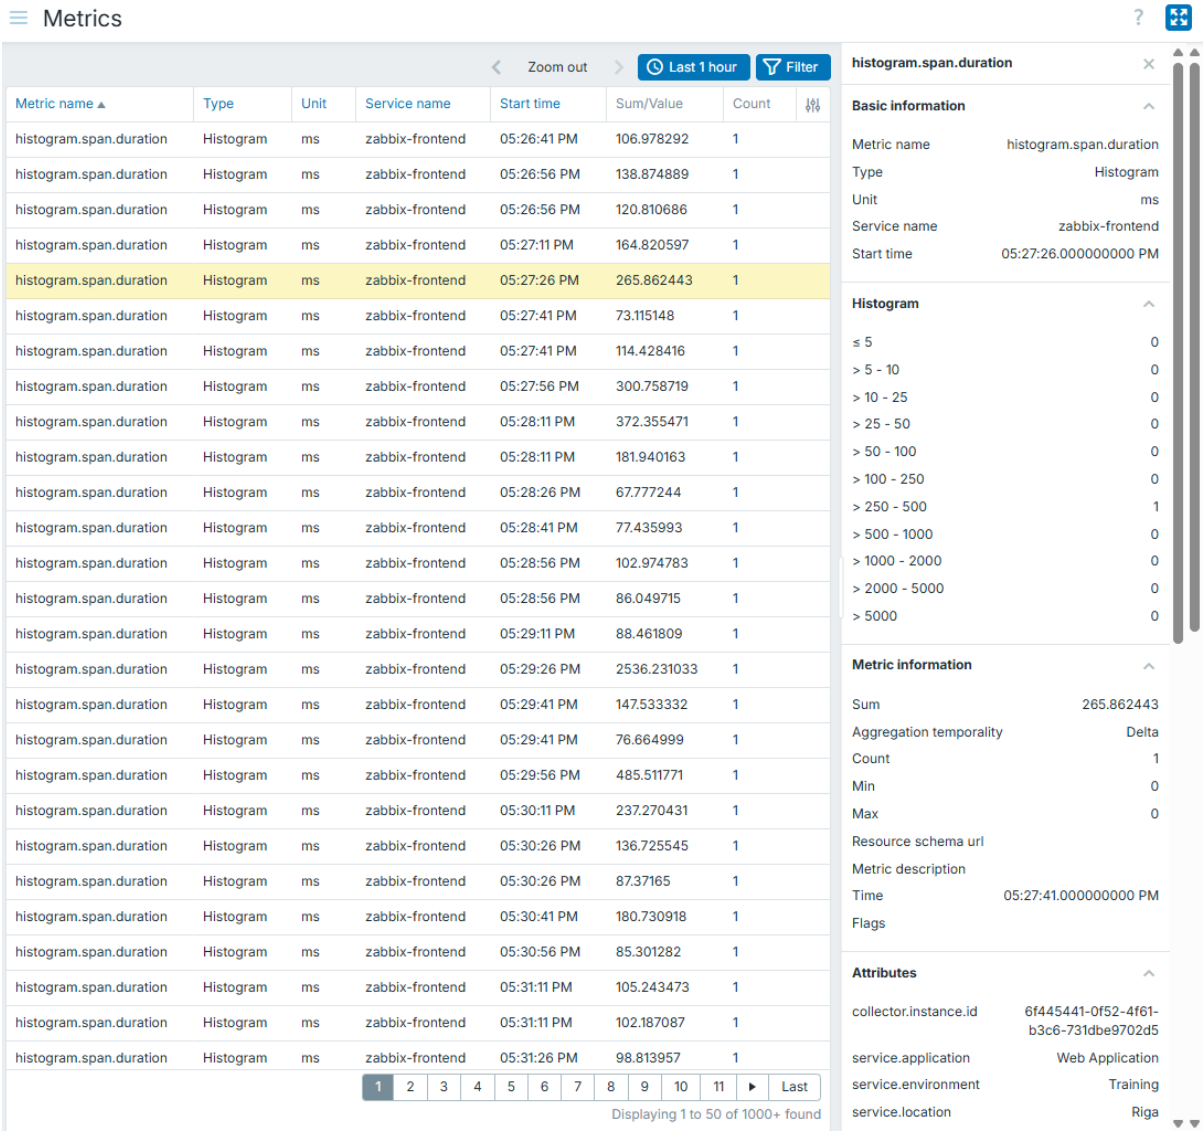
A data point can be one of the following types:

- **Gauge** - a value measured at a specific time (e.g., current memory usage).
- **Sum** - a sum of measurements over a time period (e.g., total number of requests).
- **Histogram** - the number of measurements over a time period that fall into each of several fixed value ranges, called buckets (e.g., how many requests took 0-100 ms, 100-250 ms, or 250-500 ms).
- **Exponential histogram** - like a histogram, but each bucket is wider than the previous one by the same factor (e.g., 1-2 ms, 2-4 ms, 4-8 ms).

Each data point is produced by a **service** (e.g., `zabbix-frontend`; named by the OpenTelemetry `service.name` resource attribute), and by an instrumentation scope, the part of the code that created it (e.g., a library that measures HTTP request durations).

To provide context about each measurement, data points also have key-value **attributes**: data point attributes describe what was measured and distinguish data points of the same metric (e.g., `http.request.method: POST`), and resource attributes describe the service and the environment it runs in (e.g., `host.name: eabf8c2a5ddc`).

For more details on metrics, see the OpenTelemetry documentation.



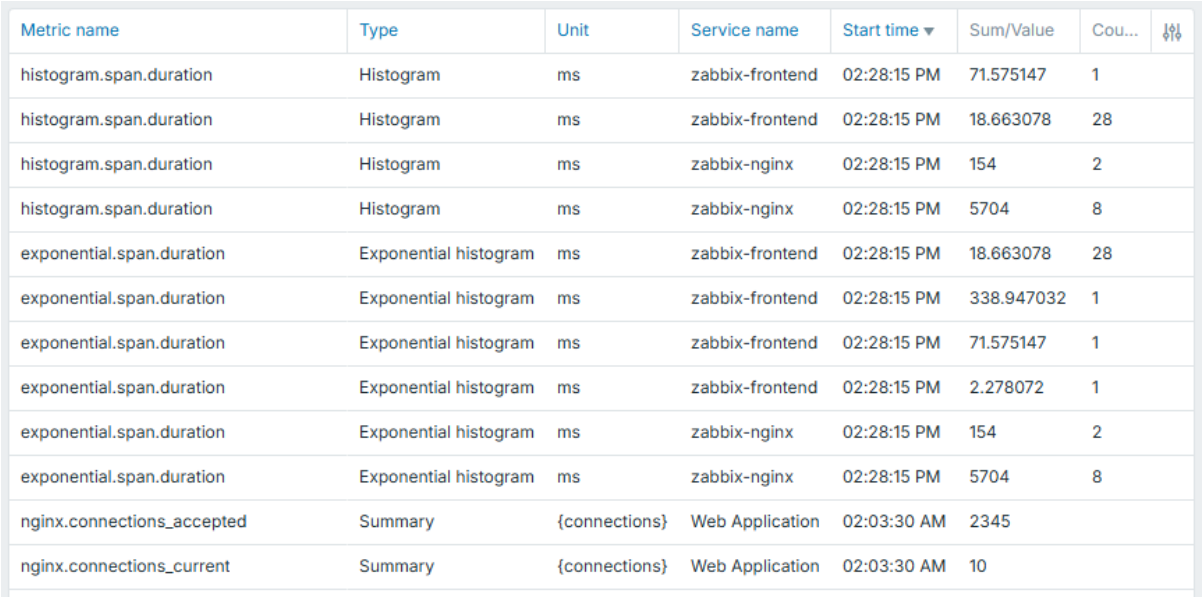
Metric name	Type	Unit	Service name	Start time	Sum/Value	Count
histogram.span.duration	Histogram	ms	zabbix-frontend	05:26:41 PM	106.978292	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:26:56 PM	138.874889	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:26:56 PM	120.810686	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:27:11 PM	164.820597	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:27:26 PM	265.862443	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:27:41 PM	73.115148	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:27:41 PM	114.428416	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:27:56 PM	300.758719	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:28:11 PM	372.355471	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:28:11 PM	181.940163	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:28:26 PM	67.777244	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:28:41 PM	77.435993	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:28:56 PM	102.974783	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:28:56 PM	86.049715	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:29:11 PM	88.461809	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:29:26 PM	2536.231033	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:29:41 PM	147.533332	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:29:41 PM	76.664999	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:29:56 PM	485.511771	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:30:11 PM	237.270431	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:30:26 PM	136.725545	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:30:26 PM	87.37165	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:30:41 PM	180.730918	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:30:56 PM	85.301282	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:31:11 PM	105.243473	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:31:11 PM	102.187087	1
histogram.span.duration	Histogram	ms	zabbix-frontend	05:31:26 PM	98.813957	1

The *APM* > *Metrics* section contains the following elements:

- Kiosk mode button - display only the page content, without the menu and header.
- Time period selector - select the time period for which to display data points.
- Filter - display only the data points you are interested in.
- Metric list - view the data points.
- Metric details - view the details of the selected data point.

Metric list

Each row in the metric list represents a data point.



Metric name	Type	Unit	Service name	Start time ▾	Sum/Value	Cou...
histogram.span.duration	Histogram	ms	zabbix-frontend	02:28:15 PM	71.575147	1
histogram.span.duration	Histogram	ms	zabbix-frontend	02:28:15 PM	18.663078	28
histogram.span.duration	Histogram	ms	zabbix-nginx	02:28:15 PM	154	2
histogram.span.duration	Histogram	ms	zabbix-nginx	02:28:15 PM	5704	8
exponential.span.duration	Exponential histogram	ms	zabbix-frontend	02:28:15 PM	18.663078	28
exponential.span.duration	Exponential histogram	ms	zabbix-frontend	02:28:15 PM	338.947032	1
exponential.span.duration	Exponential histogram	ms	zabbix-frontend	02:28:15 PM	71.575147	1
exponential.span.duration	Exponential histogram	ms	zabbix-frontend	02:28:15 PM	2.278072	1
exponential.span.duration	Exponential histogram	ms	zabbix-nginx	02:28:15 PM	154	2
exponential.span.duration	Exponential histogram	ms	zabbix-nginx	02:28:15 PM	5704	8
nginx.connections_accepted	Summary	{connections}	Web Application	02:03:30 AM	2345	
nginx.connections_current	Summary	{connections}	Web Application	02:03:30 AM	10	

Data points are displayed in a customizable table. The *Customize table* button in the table header lets you show, hide, or reorder columns, and reset the layout to default. You can also resize each column by dragging its right border; double-click the border to reset. All customizations are applied immediately and remembered for the current user.

The metric list can have the following columns:

Note:
Each column, except *Type*, corresponds to a column in the metric tables (`otel_metrics_gauge`, `otel_metrics_sum`, `otel_metrics_histogram`, and `otel_metrics_exponential_histogram`) of the telemetry data source; for the list of columns, see Supported telemetry data columns. Most columns have the same name as in the data source (e.g., *Metric name* corresponds to the `MetricName` column). Where the names differ, the data source column is specified in the table below.

Column	Description
Metric name	Metric that the data point belongs to (e.g., `http.server.duration`).
Type	Type of the data point: *Gauge*, *Sum*, *Histogram*, or *Exponential histogram*. Determined by the metric table that holds the data point.
Unit	Unit, if any, of the metric values (e.g., `s` or `By`; see UCUM). Data source column: `MetricUnit`.
Service name	Service that produced the data point.
Scope name	Instrumentation scope that produced the data point.
Start time	Date and time when the measurement period of the data point started. Data source column: `StartTimeUnix`.
Sum/Value	Sum of all measurements (*Histogram*, *Exponential histogram*), or the value (*Gauge*, *Sum*) of the data point.
Count	Number of measurements in a *Histogram* or *Exponential histogram* data point.
Time	Date and time of the data point; for *Sum*, *Histogram*, and *Exponential histogram* data points, the end of the measurement period. Data source column: `TimeUnix`.
Flags	Flags of the data point: **0** - no flags are set. **1** - no value was recorded for the data point.
Aggregation temporality	How the values of a *Sum*, *Histogram*, or *Exponential histogram* data point are aggregated over time: **Unspecified** - not specified. **Delta** - aggregated over the period since the previous data point. **Cumulative** - aggregated over the period since the start time.
Metric description	Description, if any, of the metric.
Attributes	Data point attributes. Click the ••• three-dot button to display all attributes.
Resource attributes	Resource attributes of the data point. Click the ••• three-dot button to display all attributes.

Column	Description
Scope attributes	Attributes of the instrumentation scope. Click the ••• three-dot button to display all attributes.

Metric details

Click a data point in the metric list to open the metric details panel, which shows all information about the data point.

exponential.span.calls

Basic information

Metric name	exponential.span.calls
Type	Summary
Unit	
Service name	zabbix-frontend
Start time	02:22:15.000000000 PM

Metric information

Value	1
Aggregation temporality	Delta
Is monotonic	true
Resource schema url	
Metric description	
Time	02:22:30.000000000 PM
Flags	

Attributes

Resource attributes

Scope attributes

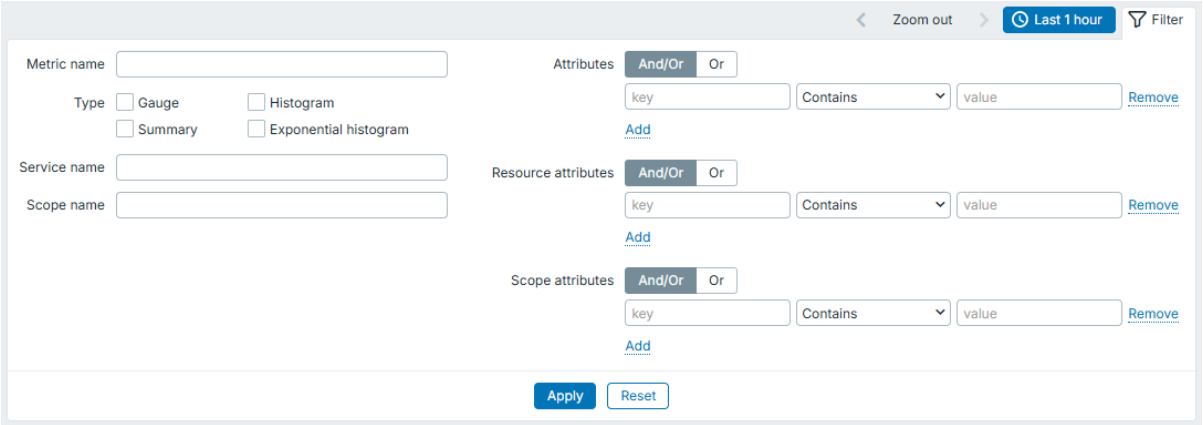
Scope

The metric details panel can have the following sections:

Section	Description
Basic information	Metric name, type, unit, service name, and start time.
Histogram	Buckets (e.g., `> 5 – 10`), each with the number of measurements in it. Displayed only for *Histogram* data points.
Exponential histogram	Buckets (e.g., `>= 249.7885 – 249.7887`), each with the number of measurements in it. Displayed only for *Exponential histogram* data points.
Metric information	Values of the data point, which depend on its type, and the metric description, time, flags, and resource schema URL.
Attributes	Data point attributes.
Resource attributes	Resource attributes.
Scope attributes	Attributes, if any, of the instrumentation scope.
Scope	Name, version, and schema URL of the instrumentation scope.

Using filter

You can use the filter to display only the data points you are interested in.



To filter by attributes, enter the attribute key, select the operator, and enter the attribute value.

The following operators are supported:

- **Contains** - *(default)* the attribute value contains the entered string.
- **Equals** - the attribute value equals the entered string.
- **Does not contain** - the attribute value does not contain the entered string.
- **Does not equal** - the attribute value does not equal the entered string.
- **Exists** - the attribute with the entered key exists.
- **Does not exist** - the attribute with the entered key does not exist.

If you add several conditions, select how to evaluate them:

- **And/Or** - *(default)* all conditions must be met; conditions with the same attribute key are combined with *Or*.
- **Or** - at least one condition must be met.

Logs

Overview

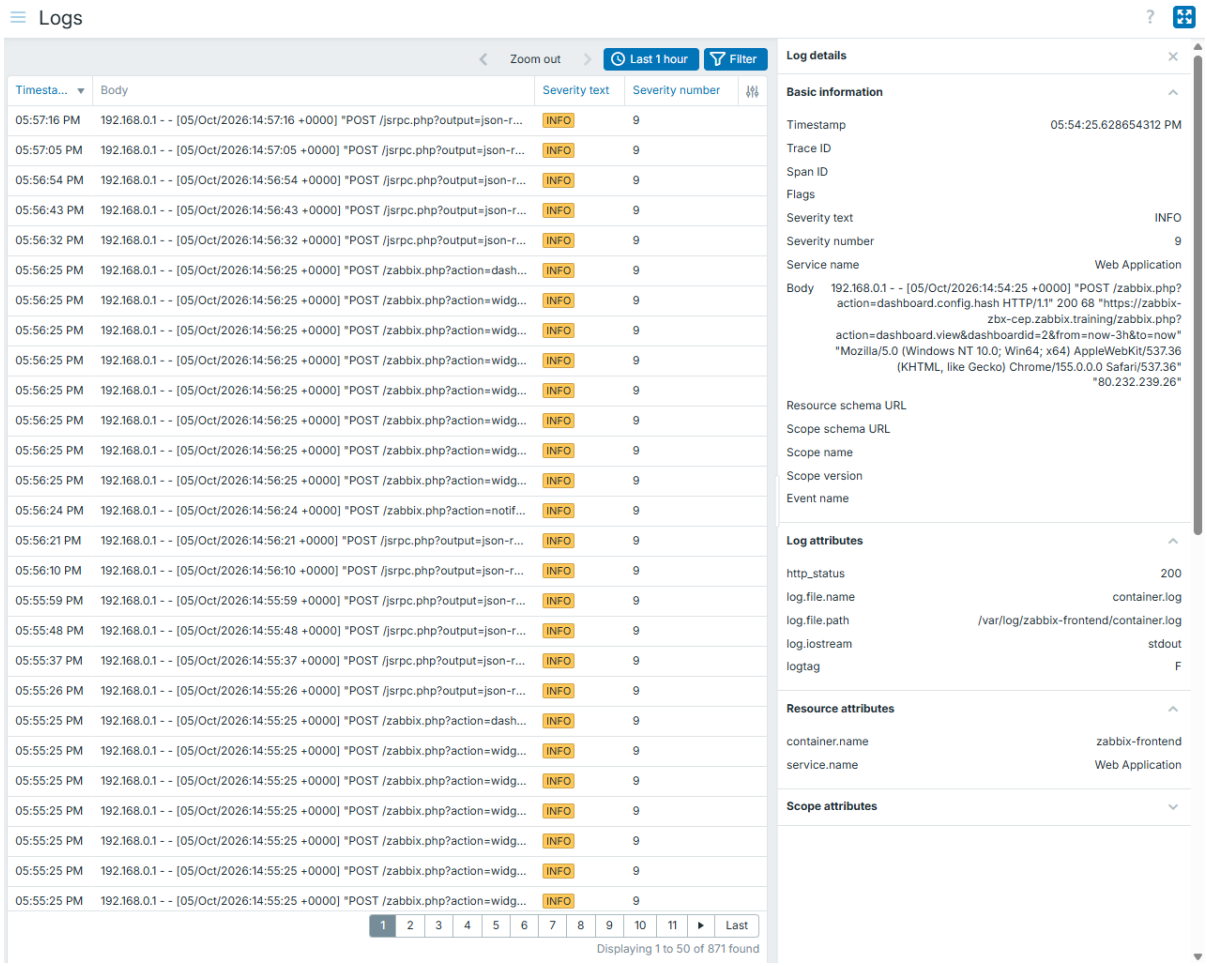
The *APM* > *Logs* section displays log records collected from your applications and services.

A **log record** is a text record of an event (e.g., a failed payment or a database timeout). Each log record has a severity and a message (body). If a log record was produced during an operation that is part of a trace, it also contains the IDs of that trace and span.

Each log record is produced by a **service** (e.g., `zabbix-frontend`; named by the OpenTelemetry `service.name` resource attribute), and by an instrumentation scope, the part of the code that created it (e.g., a logging library).

To provide context about each event, log records also have key-value **attributes**: log attributes describe the event itself (e.g., `log.file.path: /var/log/zabbix/zabbix_frontend.log`), and resource attributes describe the service and the environment it runs in (e.g., `host.name: eabf8c2a5ddc`).

For more details on logs, see the OpenTelemetry documentation.



The *APM* > *Logs* section contains the following elements:

- Kiosk mode button - display only the page content, without the menu and header.
- Time period selector - select the time period for which to display log records.
- Filter - display only the log records you are interested in.
- Log list - view the log records.
- Log details - view the details of the selected log record.

Log list

Each row in the log list represents a log record.

Timestamp	Body	Severity text	Severity number
06:04:03 PM	192.168.0.1 - - [05/Oct/2026:15:04:03 +0000] "POST /jsrpc.php?output=json-rpc HTTP/1.1" 200 76 "htt...	INFO	9
06:03:52 PM	192.168.0.1 - - [05/Oct/2026:15:03:52 +0000] "POST /jsrpc.php?output=json-rpc HTTP/1.1" 200 76 "htt...	INFO	9
06:03:41 PM	192.168.0.1 - - [05/Oct/2026:15:03:41 +0000] "POST /jsrpc.php?output=json-rpc HTTP/1.1" 200 76 "htt...	INFO	9
06:03:30 PM	192.168.0.1 - - [05/Oct/2026:15:03:30 +0000] "POST /jsrpc.php?output=json-rpc HTTP/1.1" 200 76 "htt...	INFO	9
06:03:27 PM	192.168.0.1 - - [05/Oct/2026:15:03:27 +0000] "POST /zabbix.php?action=widget.svggraph.view HTTP/...	INFO	9
06:03:27 PM	192.168.0.1 - - [05/Oct/2026:15:03:27 +0000] "POST /zabbix.php?action=widget.svggraph.view HTTP/...	INFO	9
06:03:26 PM	192.168.0.1 - - [05/Oct/2026:15:03:26 +0000] "POST /zabbix.php?action=widget.svggraph.view HTTP/...	INFO	9
06:03:26 PM	192.168.0.1 - - [05/Oct/2026:15:03:26 +0000] "POST /zabbix.php?action=widget.svggraph.view HTTP/...	INFO	9
06:03:26 PM	192.168.0.1 - - [05/Oct/2026:15:03:26 +0000] "POST /zabbix.php?action=dashboard.config.hash HTTP...	INFO	9
06:03:26 PM	192.168.0.1 - - [05/Oct/2026:15:03:26 +0000] "POST /zabbix.php?action=widget.svggraph.view HTTP/...	INFO	9
06:03:26 PM	192.168.0.1 - - [05/Oct/2026:15:03:26 +0000] "POST /zabbix.php?action=widget.problems.view HTTP/...	INFO	9
06:03:25 PM	192.168.0.1 - - [05/Oct/2026:15:03:25 +0000] "POST /zabbix.php?action=notifications.get&output=ajax...	INFO	9
06:03:25 PM	192.168.0.1 - - [05/Oct/2026:15:03:25 +0000] "POST /zabbix.php?action=widget.svggraph.view HTTP/...	INFO	9
06:03:19 PM	192.168.0.1 - - [05/Oct/2026:15:03:19 +0000] "POST /jsrpc.php?output=json-rpc HTTP/1.1" 200 76 "htt...	INFO	9

Log records are displayed in a customizable table. The *Customize table* button in the table header lets you show, hide, or reorder columns, and reset the layout to default. You can also resize each column by dragging its right border; double-click the

border to reset. All customizations are applied immediately and remembered for the current user.

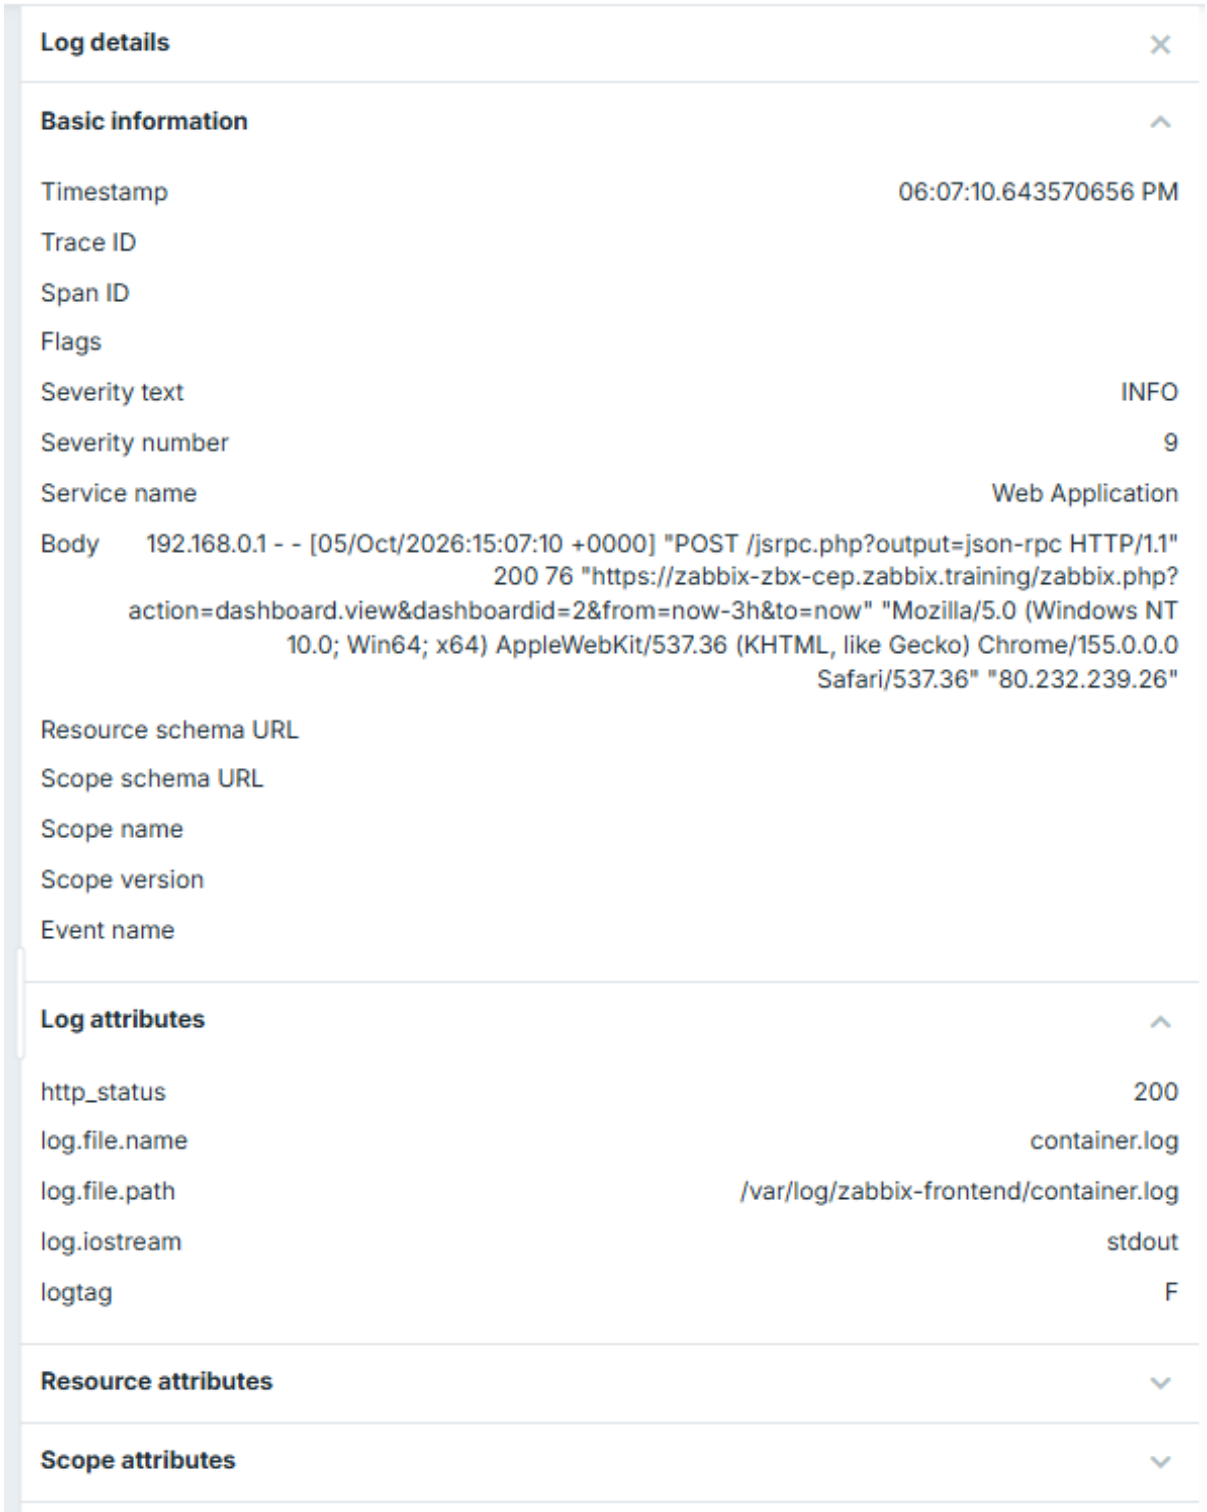
The log list can have the following columns:

> **Note:**
> Each column corresponds to a column in the `otel_logs` table of the telemetry data source; for the list of columns, see Supported telemetry data columns. Most columns have the same name as in the data source (e.g., *Severity text* corresponds to the `SeverityText` column). Where the names differ, the data source column is specified in the table below.

Column	Description
Timestamp	Date and time when the event occurred.
Body	Message of the log record.
Severity text	Severity level as provided by the application (e.g., `INFO`, `WARN`, or `ERROR`). The label color depends on the severity.
Severity number	Severity level as a number: **0** - UNSPECIFIED. **1-4** - TRACE. **5-8** - DEBUG. **9-12** - INFO. **13-16** - WARN. **17-20** - ERROR. **21-24** - FATAL.
Trace ID	ID, if any, of the trace during which the log record was produced.
Span ID	ID, if any, of the span during which the log record was produced.
Flags	Trace flags of the log record: **0** - the trace is not sampled. **1** - the trace is sampled (recorded). Data source column: `TraceFlags`.
Service name	Service that produced the log record.
Scope name	Instrumentation scope that produced the log record.
Scope version	Version of the instrumentation scope.
Event name	Name, if any, of the event type that the log record represents.
Log attributes	Log attributes of the log record. Click the ••• three-dot button to display all attributes.
Resource attributes	Resource attributes of the log record. Click the ••• three-dot button to display all attributes.
Scope attributes	Attributes of the instrumentation scope. Click the ••• three-dot button to display all attributes.

Log details

Click a log record in the log list to open the *Log details* panel, which shows all information about the log record.

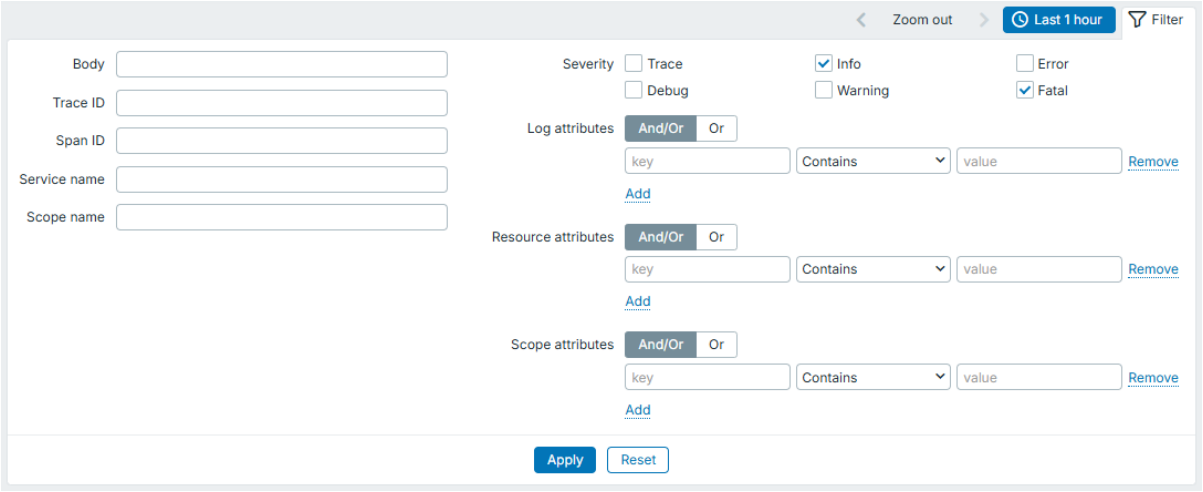


The *Log details* panel can have the following sections:

Section	Description
Basic information	Trace ID, span ID, and trace flags of the log record.
Log attributes	Log attributes of the log record.
Resource attributes	Schema URL and resource attributes.
Scope attributes	Schema URL, name, version, and attributes of the instrumentation scope.

Using filter

You can use the filter to display only the log records you are interested in.



To filter by attributes, enter the attribute key, select the operator, and enter the attribute value.

The following operators are supported:

- **Contains** - *(default)* the attribute value contains the entered string.
- **Equals** - the attribute value equals the entered string.
- **Does not contain** - the attribute value does not contain the entered string.
- **Does not equal** - the attribute value does not equal the entered string.
- **Exists** - the attribute with the entered key exists.
- **Does not exist** - the attribute with the entered key does not exist.

If you add several conditions, select how to evaluate them:

- **And/Or** - *(default)* all conditions must be met; conditions with the same attribute key are combined with *Or*.
- **Or** - at least one condition must be met.

Services

Overview

The Services menu is for the service monitoring functions of Zabbix.

Services

Overview

In this section you can see a high-level status of whole services that have been configured in Zabbix, based on your infrastructure.

A service may be a hierarchy consisting of several levels of other services, called ”child” services, which are attributes to the overall status of the service (see also an overview of the service monitoring functionality.)

The main categories of service state are *OK* or *Problem*, where the *Problem* state is expressed by the corresponding problem severity name and color.

While view mode allows services to be monitored along with their status and other details, you can also configure the service hierarchy in this section (add/edit services, child services) by switching to edit mode.

To switch from view to edit mode (and back), click the respective button in the upper-right corner:

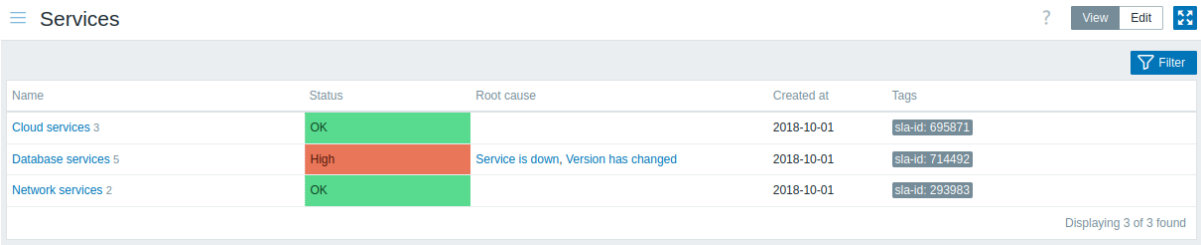
- 

- view services
- 

- add/edit services, and child services

Note that access to editing depends on user role settings.

Viewing services



A list of the existing services is displayed.

Displayed data:

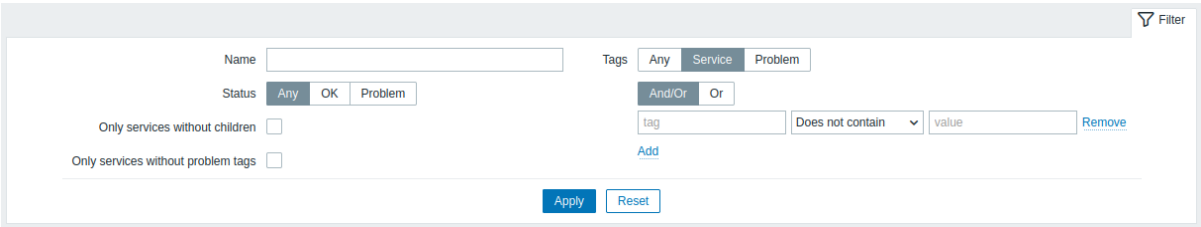
Parameter	Description
Name	Service name. The service name is a link to service details. The number after the name indicates how many child services the service has.
Status	Service state: **OK** - no problems **<problem color and severity>** - indicates a problem and its severity. In case of multiple problems, the color and severity of the most critical problem are displayed.
Root cause	Underlying problems that directly or indirectly affect the service state are listed. The same problems are listed as returned by the {SERVICE.ROOTCAUSE} macro. Click the problem name to see more details about it in *Monitoring* > *Problems*. Problems that do not affect the service state are not in the list.
Created at	The time when the service was created is displayed.
Tags	Tags of the service are displayed. Tags are used to identify a service in service actions and SLAs.

Buttons

View mode buttons being common for all sections are described on the Monitoring page.

Using filter

You can use the filter to display only the services you are interested in.



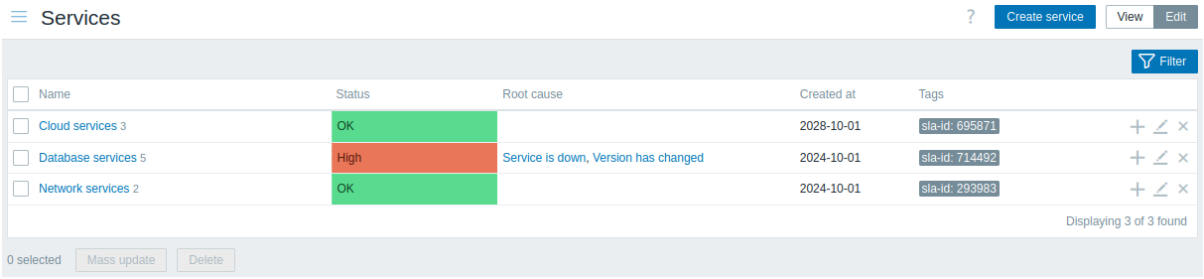
Parameter	Description
Name	Filter by service name.
Status	Filter by service state.
Only services without children	Mark the checkbox to display only services without child services. This parameter is available only in service edit mode.
Only services without problem tags	Mark the checkbox to display only services without problem tags. This parameter is available only in service edit mode.

Parameter	Description
Tags	Filter by service tag name and value, or service problem tag name and value (in service edit mode). It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met.

Editing services

Click the *Edit* button to access edit mode. When in edit mode, the listing is complemented with checkboxes before the entries and also these additional options:

-  - add a child service to this service
-  - edit this service
-  - delete this service

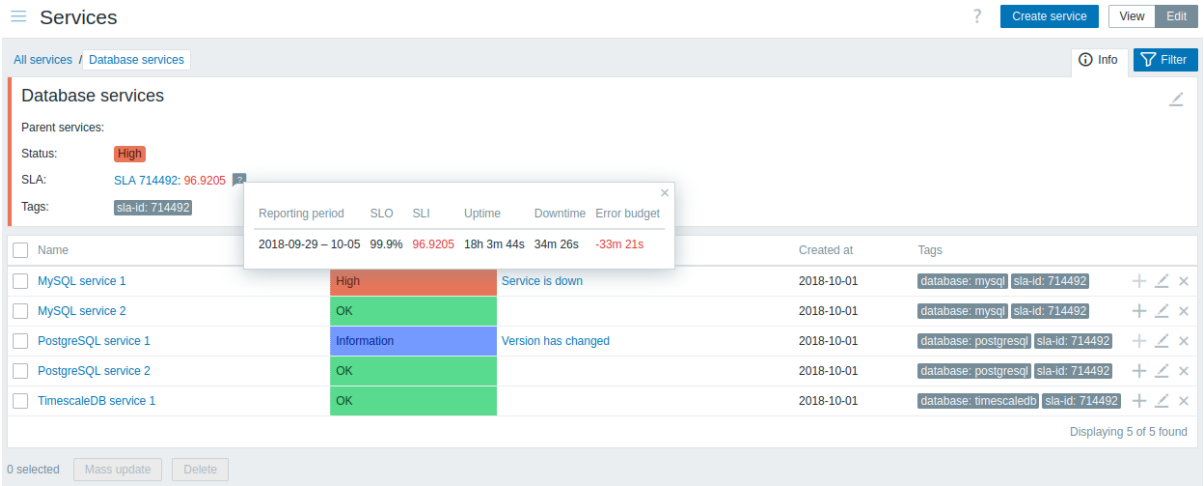


To configure a new service, click the *Create service* button in the upper-right corner.

Service details

To access service details, click the service name. To return to the list of all services, click *All services*.

Service details include the info box and the list of child services.



To access the info box, click the *Info* tab. The info box contains the following entries:

- Names of parent services (if any)
- Current status of this service
- Current SLA(s) of this service, in the format `SLA name:service level indicator`. 'SLA name' is also a link to the SLA report for this service. If you position the mouse on the info box next to the service-level indicator (SLI), a pop-up info list is displayed with SLI details. The service-level indicator displays the current service level, in percentage.
- Service tags

The info box also contains a link to the service configuration.

To use the filter for child services, click the *Filter* tab.

When in edit mode, the child service listing is complemented with additional editing options:

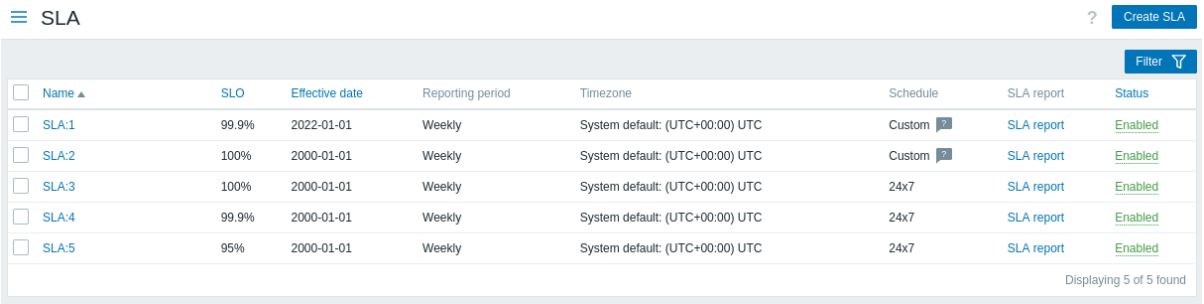
-  - add a child service to this service
-  - edit this service
-  - delete this service

SLA

Overview

This section allows SLAs to be viewed and configured.

SLAs



A list of the configured SLAs is displayed. *Note* that only the SLAs related to services accessible to the user will be displayed (as read-only, unless *Manage SLA* is enabled for the user role).

Displayed data:

Parameter	Description
Name	The SLA name is displayed. The name is a link to SLA configuration.
SLO	The service level objective (SLO) is displayed.
Effective date	The date of starting SLA calculation is displayed.
Reporting period	The period used in the SLA report is displayed - *daily*, *weekly*, *monthly*, *quarterly*, or *annually*.
Time zone	The SLA time zone is displayed.
Schedule	The SLA schedule is displayed - 24x7 or custom.
SLA report	Click the link to see the SLA report for this SLA.
Status	The SLA status is displayed - enabled or disabled.

SLA report

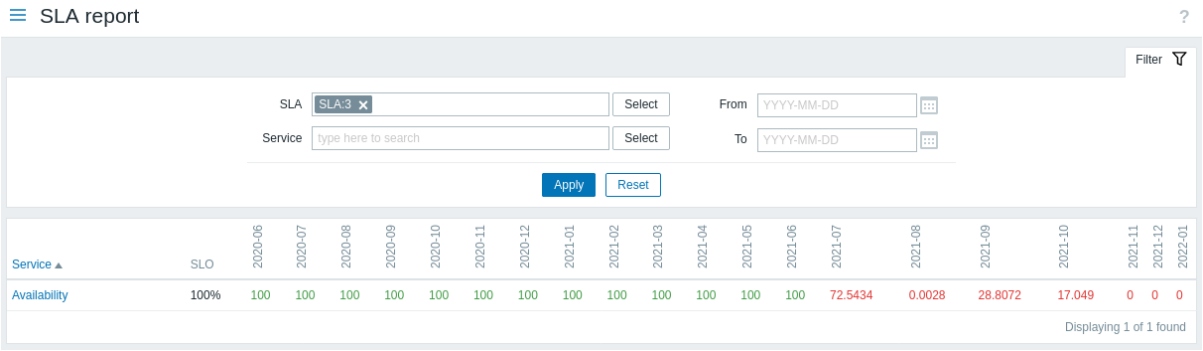
Overview

This section allows SLA reports to be viewed based on the criteria selected in the filter.

SLA reports can also be displayed as a dashboard widget.

Report

The filter allows you to select the report based on the SLA name as well as the service name. It is also possible to limit the displayed period.

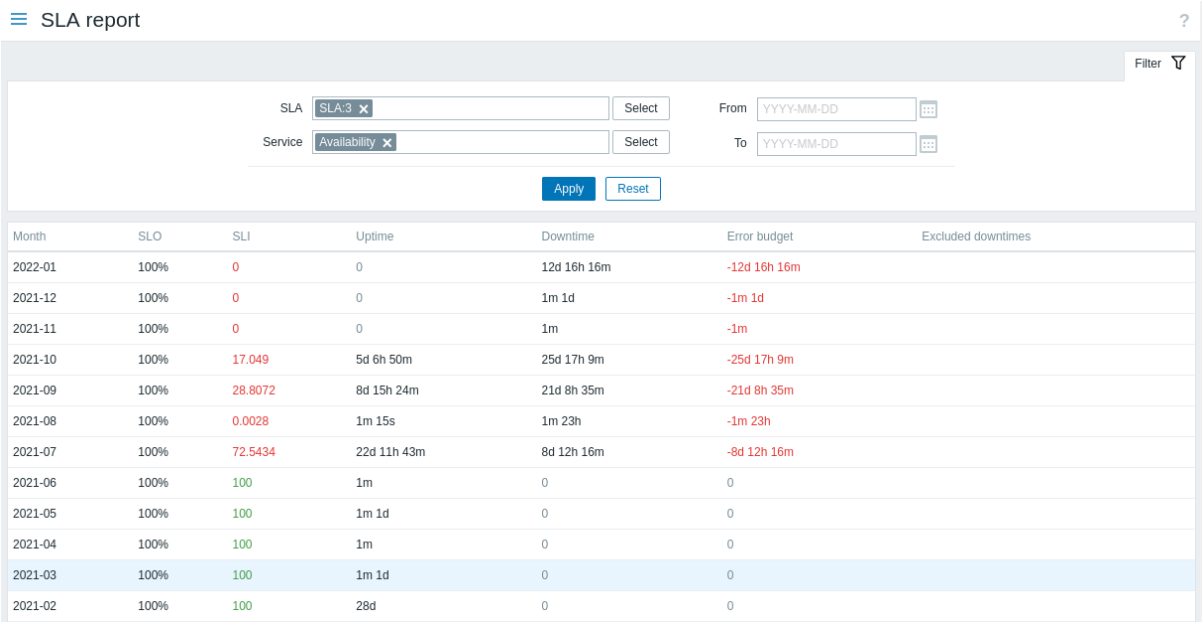


Each column (period) displays the SLI for that period. SLIs that are in breach of the set SLO are highlighted in red.

20 periods are displayed in the report. A maximum of 100 periods can be displayed, if both the *From* date and *To* date are specified.

Report details

If you click the service name in the report, you can access another report that displays a more detailed view.



Note that negative problem duration does not affect SLA calculation or reporting.

Inventory

Overview

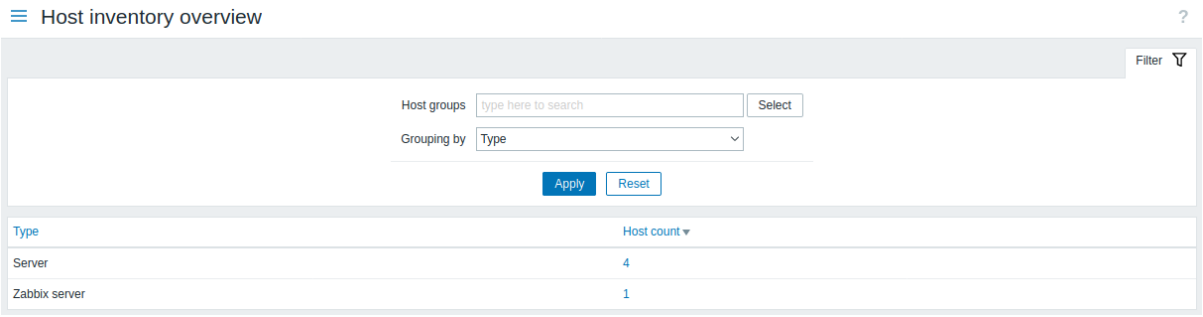
The Inventory menu features sections providing an overview of host inventory data by a chosen parameter as well as the ability to view host inventory details.

Overview

Overview

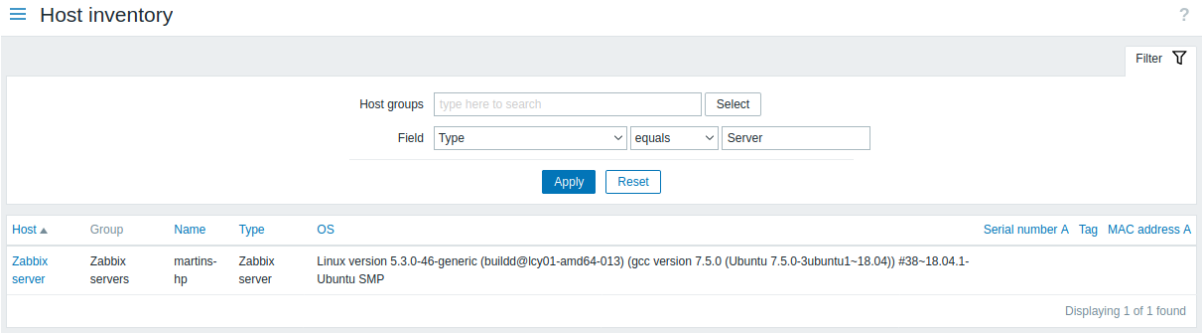
The *Inventory* > *Overview* section provides ways of having an overview of host inventory data.

For an overview to be displayed, choose host groups (or none) and the inventory field by which to display data. The number of hosts corresponding to each entry of the chosen field will be displayed.



The completeness of an overview depends on how much inventory information is maintained with the hosts.

Numbers in the *Host count* column are links; they lead to these hosts being filtered out in the *Host Inventories* table.

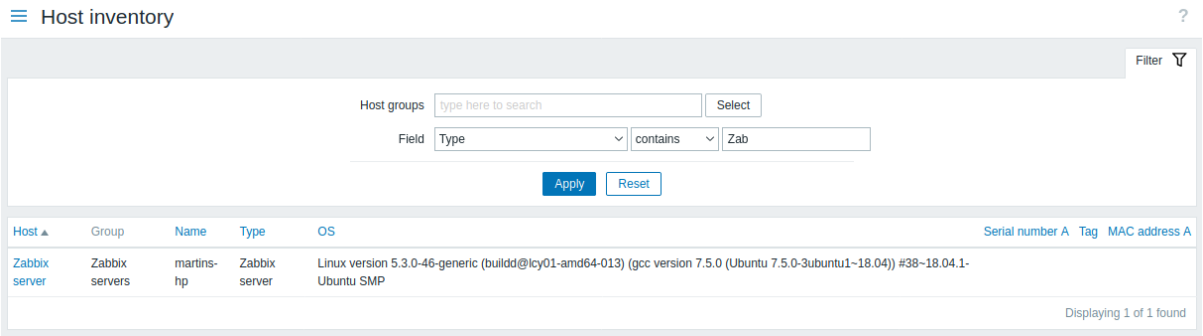


Hosts

Overview

In the *Inventory* > *Hosts* section inventory data of hosts are displayed.

You can filter the hosts by host group(s) and by any inventory field to display only the hosts you are interested in.

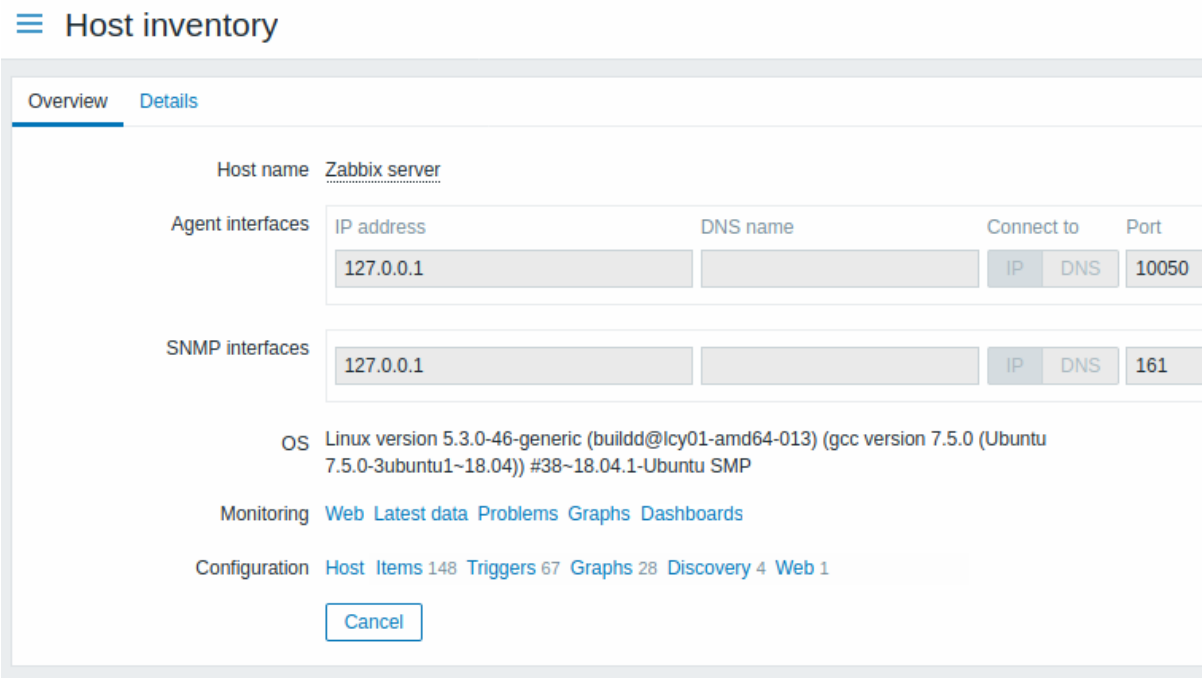


To display all host inventories, click the *Reset* button.

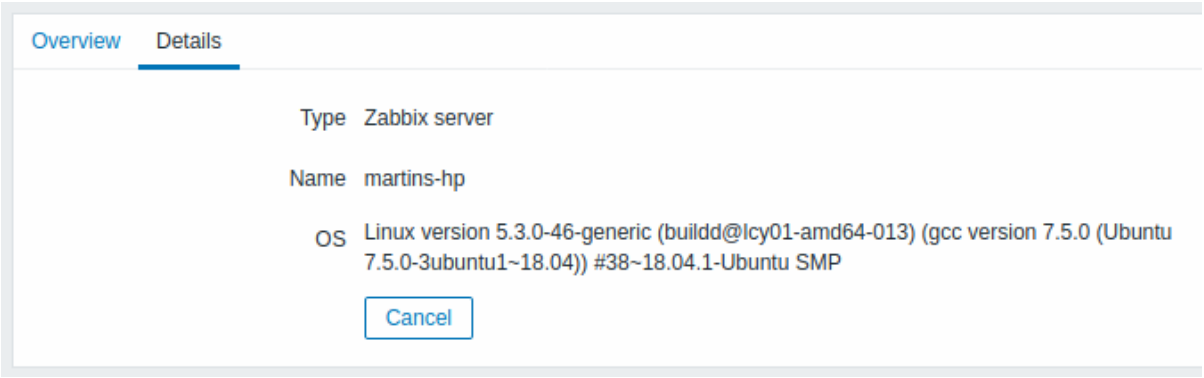
While only some key inventory fields are displayed in the table, you can also view all available inventory information for that host. To do that, click the host name in the first column.

Inventory details

The **Overview** tab contains some general information about the host, latest monitoring data, and host configuration options:



The **Details** tab contains all available inventory details for the host:



The completeness of inventory data depends on how much inventory information is maintained with the host. If no information is maintained, the *Details* tab is disabled.

Reports

Overview

The Reports menu features several sections that contain a variety of predefined and user-customizable reports focused on displaying an overview of such parameters as system information, triggers and gathered data.

System information

Overview

In *Reports* > *System information*, a summary of key Zabbix server and system data is displayed. System data is collected using internal items.

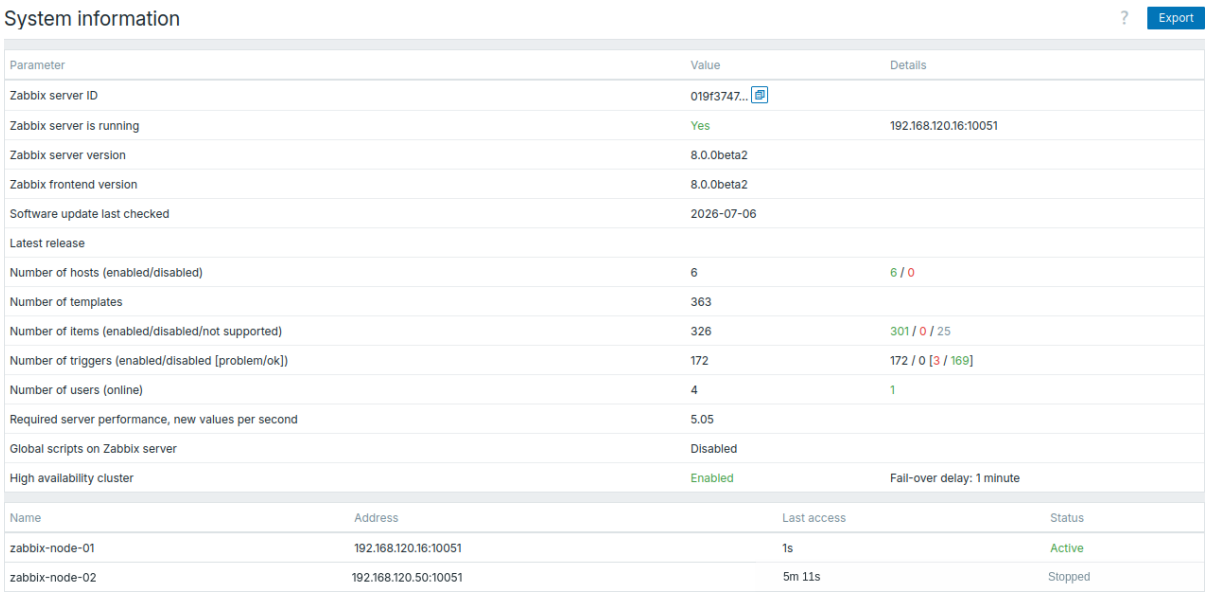
The system information report is visible only to Zabbix *Super Admin* and Zabbix *Admin* users.

Note that in a high availability setup, it is possible to redirect the system information source (server instance). To do this, edit the frontend configuration file file - uncomment and set `$ZBX_SERVER` or both `$ZBX_SERVER` and `$ZBX_SERVER_PORT` to a server other than the one shown active. Note that when setting `$ZBX_SERVER` only, a default value (10051) for `$ZBX_SERVER_PORT` will be used.

With the high availability setup enabled, a separate block is displayed below the system stats with details of high availability nodes.

System information is also available as a dashboard widget.

System stats



Displayed data:

Parameter	Value	Details
Zabbix server ID	Zabbix server unique identifier (UUIDv7).	Server ID is generated on the first server start. The first 8 characters are shown, followed by an ellipsis. When hovering over or clicking a server ID value, a tooltip appears showing the full value. Clicking the icon next to the value copies the ID to the clipboard. If high-availability is enabled, each node has its own node ID for command-line addressing, but all nodes are grouped under a single server ID generated by the node that was launched first (or was previously a standalone server), regardless of which node is currently active.
Zabbix server is running	Status of Zabbix server: **Yes** - server is running **No** - server is not running *Note:* To display the rest of the information the web, frontend needs the server to be running and there must be at least one trapper process started on the server (StartTrappers parameter in `zabbix_server.conf` file > 0).	Location and port of Zabbix server.
Zabbix server version	Current server version number is displayed. *Note:* It is only displayed when Zabbix server is running.	Server version status is displayed: **Up to date** - using the latest version; **New update available** - a more up-to-date version is available; **Outdated** - the full support period for this version has expired. This information is only available if software update check is enabled in the Zabbix server configuration. Nothing is displayed if the last software update check was performed more than a week ago or no data exist about the current version.

Parameter	Value	Details
Zabbix frontend version	Zabbix frontend version number is displayed.	Zabbix frontend version status is displayed: **Up to date** - using the latest version; **New update available** - a more up-to-date version is available; **Outdated** - the full support period for this version has expired. This information is only available if software update check is enabled in the Zabbix server configuration. Nothing is displayed if the last software update check was performed more than a week ago or no data exist about the current version.
Software update last checked	The date of the last Zabbix software update check is displayed. This information is only available if software update check is enabled in the Zabbix server configuration.	
Latest release	The number of a newer release (if available) for the current Zabbix version is displayed. This information is only available if software update check is enabled in the Zabbix server configuration. Nothing is displayed if the last software update check was performed more than a week ago or no data exist about the current version.	A link to the release notes of the latest available Zabbix release is displayed.
Number of hosts	Total number of hosts configured is displayed.	Number of monitored hosts/not monitored hosts.
Number of templates	Total number of templates is displayed.	
Number of items	Total number of items is displayed.	Number of monitored/disabled/unsupported host-level items. Items on disabled hosts are counted as disabled.
Number of triggers	Total number of triggers is displayed.	Number of enabled/disabled host-level triggers; split of the enabled triggers according to ”Problem”/”OK” states. Triggers listed under the ”OK” state include triggers with the status ”Unknown”. Triggers dependent on disabled items or assigned to disabled hosts are counted as disabled.
Number of users	Total number of users configured is displayed.	Number of users online.
Required server performance, new values per second	The expected number of new values processed by Zabbix server per second is displayed.	*Required server performance* is an estimate and can be useful as a guideline. For precise numbers of values processed, use the `zabbix[wcache,values,all]` internal item. Enabled items from monitored hosts are included in the calculation. Log items are counted as one value per item update interval. Regular interval values are counted; flexible and scheduling interval values are not. The calculation is not adjusted during a ”nodata” maintenance period. Trapper items are not counted.
Global scripts on Zabbix server	**Disabled** will be displayed in this field if global scripts are disabled on Zabbix server by setting EnableGlobalScripts=0 in server configuration.	
High availability cluster	Status of high availability cluster for Zabbix server: **Disabled** - standalone server **Enabled** - at least one high availability node exists	If enabled, the failover delay is displayed.

System information will also display an error message in the following conditions:

- The database used does not have the required character set or collation (UTF-8).
- The version of the database is below or above the supported range (available only to users with the *Super admin role* type).
- Housekeeping for TimescaleDB is incorrectly configured (history or trend tables contain compressed chunks, but *Override item history period* or *Override item trend period* options are disabled).

Buttons

The button at the upper-right corner of the page offers the following option:

Export	Export system information to a JSON file. The file is saved as `Zabbix-system-information-{time}.json`, where `{time}` is the Unix timestamp indicating when the data was prepared on the server. If high availability is enabled, details of the high-availability nodes are also exported.

High availability nodes

If high availability cluster is enabled, an additional block of data is displayed showing the status of each high availability node.

Name	Address	Last access	Status
node-active	192.168.1.13:10051	12s	Active
node6	192.168.1.10:10053	1h 2m 40s	Unavailable
node7	192.168.1.11:10053	3m 40s	Unavailable
node4	192.168.1.8:10052	1h 34m 29s	Stopped
node5	192.168.1.9:10053	1h 9m 51s	Stopped
node8	192.168.1.12:10051	21m 16s	Stopped
node1	192.168.1.5:10051	17s	Standby
node2	192.168.1.6:10051	16s	Standby
node3	192.168.1.7:10052	16 2021-10-20 17:58:47	Standby

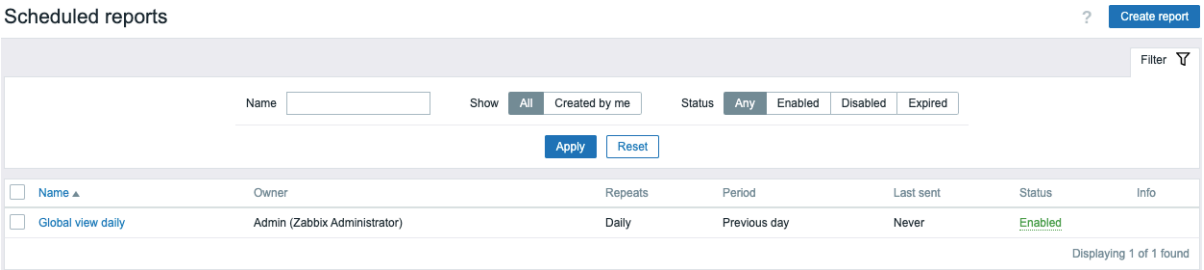
Displayed data:

Column	Description
Name	Node name, as defined in server configuration.
Address	Node IP address and port.
Last access	Time of node last access. Hovering over the cell shows the timestamp of last access in long format.
Status	Node status. Table rows are sorted by these statuses in this priority order: **Active** - node is up and working **Unavailable** - node hasn't been seen for more than failover delay (you may want to find out why) **Stopped** - node has been stopped or couldn't start (you may want to start it or delete it) **Standby** - node is up and waiting

Scheduled reports

Overview

In *Reports* > *Scheduled reports*, users with sufficient permissions can configure scheduled generation of PDF versions of the dashboards, which will be sent by email to specified recipients.



The opening screen displays information about scheduled reports, which can be filtered out for easy navigation - see Using filter section below.

Displayed data:

Column	Description
Name	Name of the report. Clicking it opens the report configuration form.
Owner	User who created the report.
Repeats	Report generation frequency (daily/weekly/monthly/yearly).
Period	Time period for which the report is prepared (previous day, week, month, or year).
Last sent	The date and time when the latest report has been sent.
Status	Current status of the report (enabled/disabled/expired). Users with sufficient permissions can change the status by clicking it - from ”Enabled” to ”Disabled” (and back); from ”Expired” to ”Disabled” (and back). For users with insufficient rights, the status is not clickable.
Info	Displays informative icons: A red icon indicates that report generation has failed; hovering over it will display a tooltip with the error information. A yellow icon indicates that a report was generated, but sending to some (or all) recipients has failed or that a report is expired; hovering over it will display a tooltip with additional information.

Using filter

You may use the filter to narrow down the list of reports. For better search performance, data is searched with macros unresolved.

The following filtering options are available:

- *Name* - partial name match is allowed
- *Show* - created by current user or all reports
- *Status* - select between ”Any” (show all reports), ”Enabled”, ”Disabled”, or ”Expired”

The filter is located below the *Scheduled reports* section name. It can be opened and collapsed by clicking the *Filter* tab in the upper-right corner.

Mass update

Sometimes you may want to delete or change the status of a number of reports at once. Instead of opening each individual report for editing, you may use the mass update function for that.

To mass-update some reports, do the following:

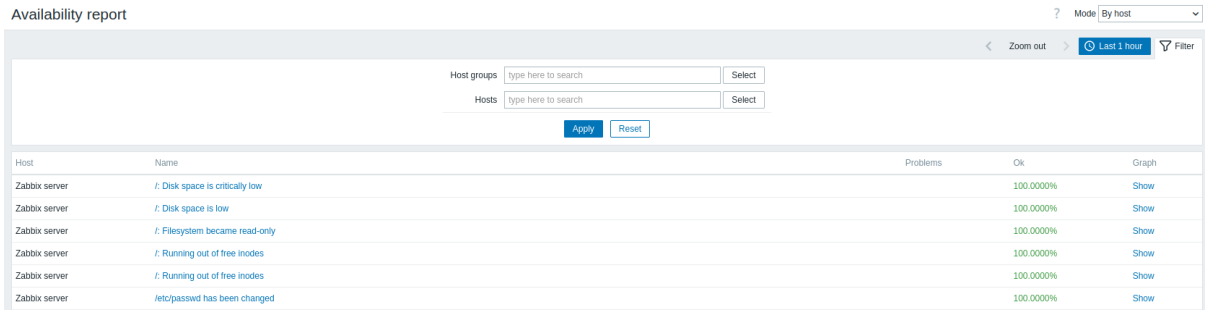
- Mark the checkboxes of the reports to update in the list
- Click the required button below the list to make the changes (*Enable*, *Disable*, or *Delete*)

Availability report

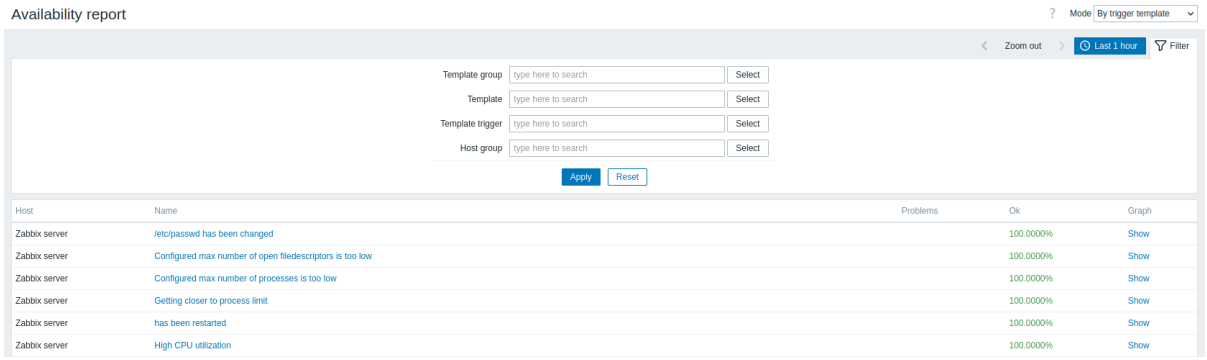
Overview

In *Reports* > *Availability report*, you can see what proportion of time each trigger has been in the problem/OK state.

For each state, a percentage of time is displayed, making it easy to determine the availability of various elements on your system.



In the dropdown in the upper-right corner, you can choose the selection mode - whether to display triggers by hosts or by triggers belonging to a template.



The name of the trigger is a link to the latest events of that trigger.

Using filter

The filter can help narrow down the number of hosts and/or triggers displayed. For better search performance, data is searched with macros unresolved.

The filter is located below the *Availability report* section name. It can be opened and collapsed by clicking the *Filter* tab on the right.

Filtering by trigger template

In *By trigger template* mode, results can be filtered by one or several of the parameters listed below.

Parameter	Description
Template group	Filter hosts by triggers that are inherited from templates belonging to the selected template group. Specifying a parent template group implicitly selects all nested template groups.
Template	Filter hosts by triggers that are inherited from the selected template, including nested templates. If a nested template has its own triggers, those triggers will not be displayed.
Template trigger	Filter hosts by the selected trigger. Other triggers of the filtered hosts will not be displayed.
Host group	Filter hosts belonging to the selected host group.

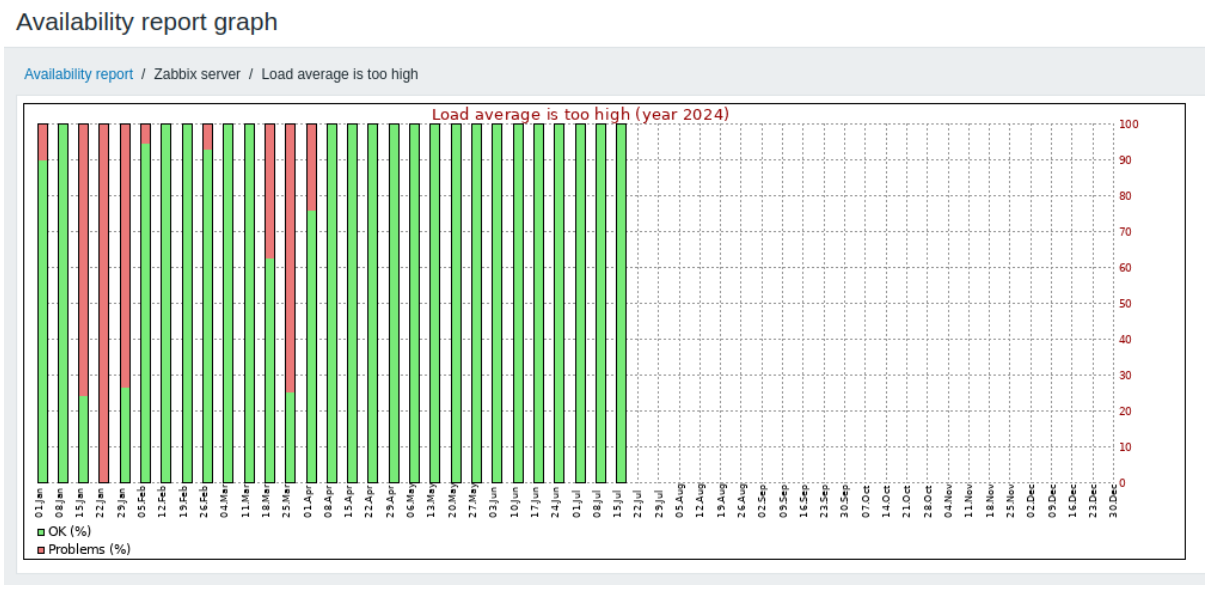
Filtering by host

In *By host* mode, results can be filtered by host or host group. Specifying a parent host group implicitly selects all nested host groups.

Time period selector

The time period selector allows you to select commonly used time periods with one click. The selector can be expanded and collapsed by clicking the *Time period* tab next to the filter.

Clicking *Show* in the *Graph* column displays availability information in a bar graph where each bar represents a passed week of the current year.



The green of a bar stands for OK time and red - for problem time.

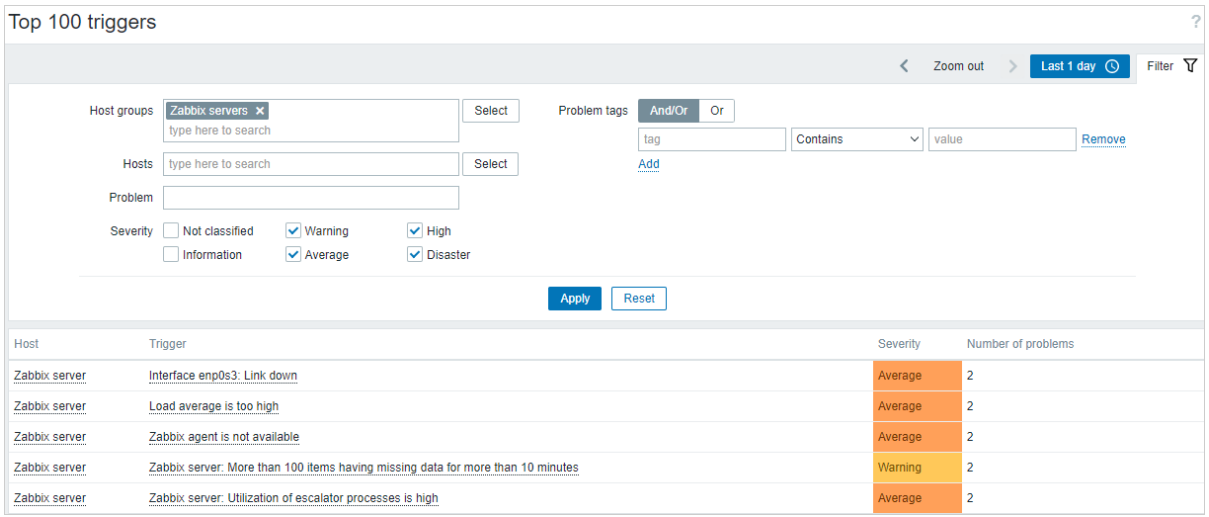
Effect of maintenance periods

Maintenance does not automatically exclude time from the availability report. Only maintenance configured as *No data collection* stops polling (so no problems are generated) and therefore leaves the availability percentage unchanged for affected triggers.

Top 100 triggers

Overview

In *Reports* > *Top 100 triggers*, you can see the triggers with the highest number of problems detected during the selected period.



Both host and trigger column entries are links that offer some useful options:

- for host - clicking the host name brings up the host menu
- for trigger - clicking the trigger name brings up links to the latest events, simple graph for each trigger item, and the configuration forms of the trigger itself and each trigger item

Using filter

You may use the filter to display triggers by host group, host, problem name, tags, or trigger severity. Specifying a parent host group implicitly selects all nested host groups. For better search performance, data is searched with macros unresolved.

The filter is located below the *Top 100 triggers* section name. It can be opened and collapsed by clicking the *Filter* tab on the right.

Time period selector

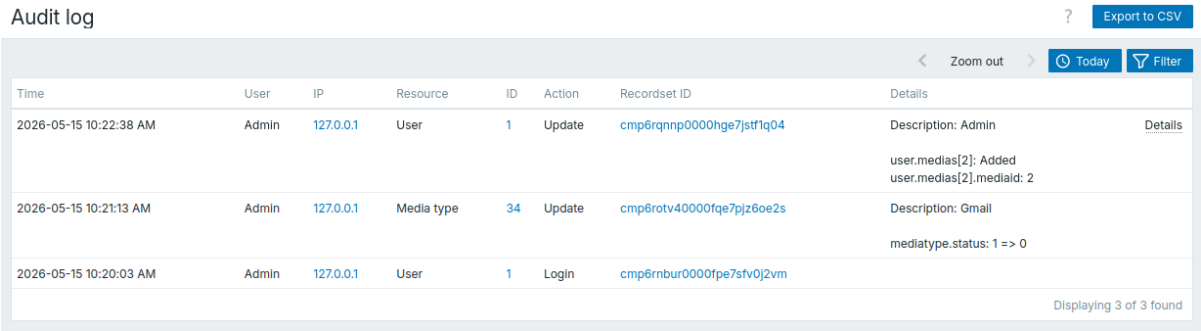
The time period selector next to the filter allows you to select often required periods with one mouse click. For more information, see Time period and host selectors.

Audit log

Overview

In the *Reports* > *Audit log* section, the records of user and system activity can be viewed.

> **Note:**
> For audit records to be collected and displayed, the *Enable audit logging* checkbox has to be marked in the *Administration* > *Audit log* section. Without this setting enabled, the history of activities will not be recorded in the database and will not be shown in the audit log.



Audit log displays the following data:

Column	Description
Time	Timestamp of the audit record.
User	User who performed the activity.
IP	IP from which the activity was initiated. Clicking the hyperlink will result in filtering audit log records by this IP.
Resource	Type of the affected resource (*API token*, *Action*, *Authentication*, *Autoregistration*, etc.).
ID	ID of the affected resource. Clicking the hyperlink will result in filtering audit log records by this resource ID.
Action	Type of the activity (*Add*, *Configuration refresh*, *Delete*, *Execute*, *Failed login*, *History clear*, *Login*, *Logout*, *Push*, *Update*).
Recordset ID	Shared ID for all audit log records created as a result of the same operation. For example, when linking a template to a host, a separate audit log record is created for each inherited template entity (item, trigger, etc.) - all these records will have the same *Recordset ID*. Clicking the hyperlink will result in filtering audit log records by this *Recordset ID*.
Details	Description of the resource and detailed information about the performed activity. If a record contains more than two rows, an additional *Details* link will be displayed. Click this link to view the full list of changes.

Note:
When a trapper item or an HTTP agent item (with trapping enabled) has received some data, an entry in the audit log will be added only if the data was sent using the `history.push` API method, and not the Zabbix sender utility.

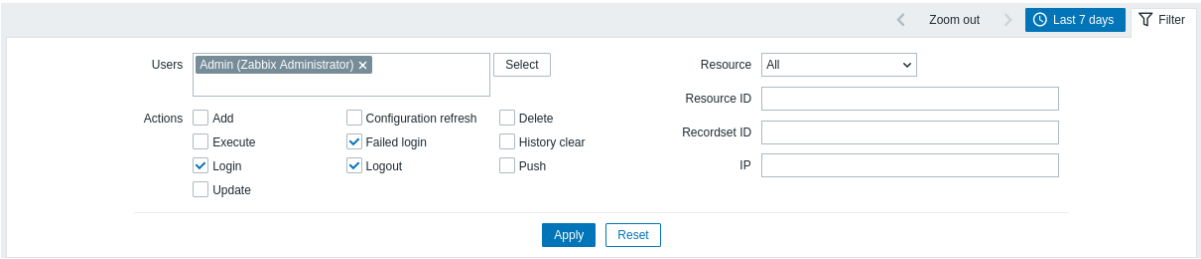
Buttons

The button at the upper-right corner of the page offers the following option:

Export to CSV	Export audit log records from all pages to a CSV file. If a filter is applied, only the filtered records will be exported.

Using filter

The filter is located below the *Audit log* bar. It can be opened and collapsed by clicking the *Filter* tab in the upper-right corner.



You may use the filter to narrow the records by user, affected resource, resource ID, performed operation (*Recordset ID*), and IP. Depending on the resource, one or more specific actions can be selected in the filter.

For better search performance, all data is searched with macros unresolved.

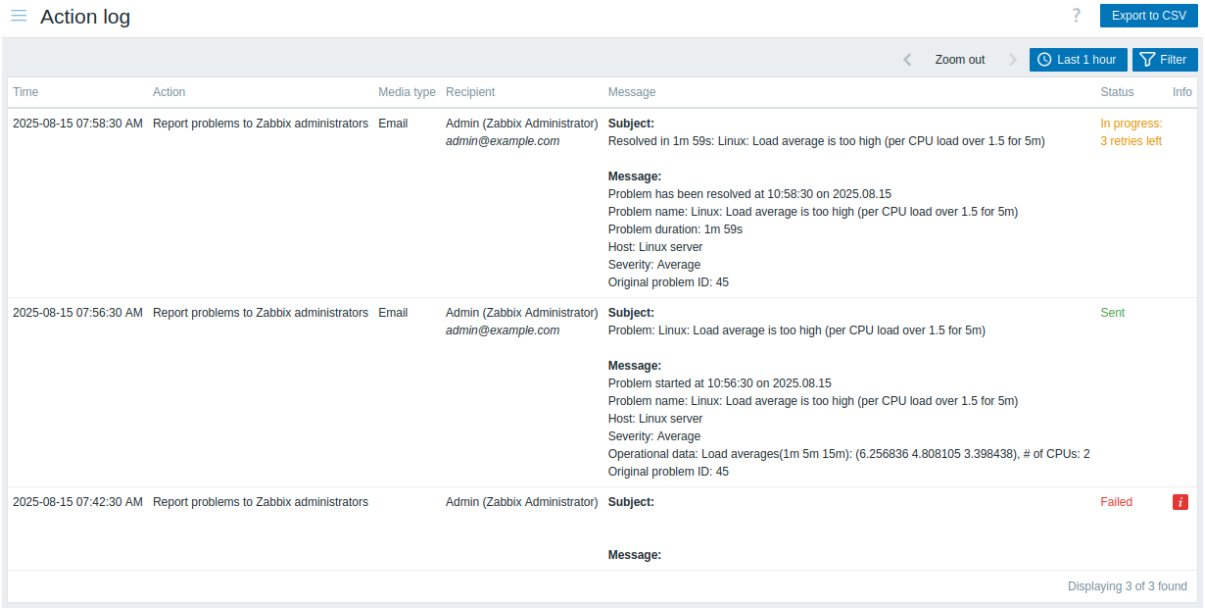
Time period selector

The time period selector next to the filter allows you to select often required periods with one mouse click. For more information, see Time period and host selectors.

Action log

Overview

In *Reports* > *Action log*, you can view details of operations (notifications, remote commands) executed within an action.



Displayed data:

Column	Description
Time	Timestamp of the operation.
Action	Name of the action causing operations.
Media type	Media type (e.g., Email, Jira, etc.) used for sending a notification. For operations that executed remote commands, this column will be empty.
Recipient	Information about the notification recipient - username, name and surname (in parentheses), and additional information depending on the media type (email, username, etc.). For operations that executed remote commands, this column will be empty.
Message	The content of the message, remote command or global script name. A remote command is separated from the target host with a colon symbol: `<host>:<command>`. For example, if the remote command was executed on Zabbix server, then the information will have the following format: `Zabbix server:<command>`.
Status	Operation status: *In progress* - operation for sending a notification is in progress (the remaining number of times the server will try to send the notification is also displayed) *Sent* - notification has been sent *Executed* - remote command has been executed *Failed* - operation has not been completed
Info	Error information (if any) regarding the operation execution.

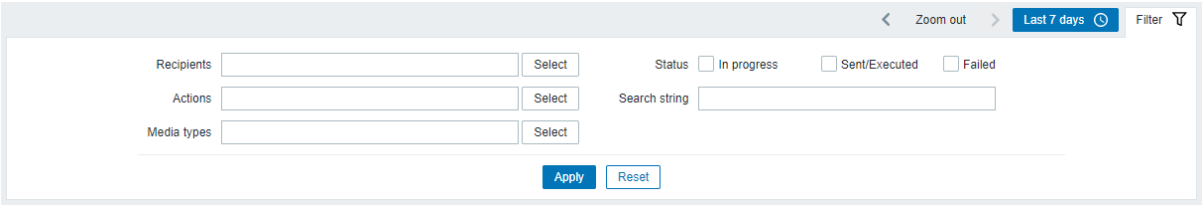
Buttons

The button at the upper-right corner of the page offers the following option:

Export to CSV	Export action log records from the current page to a CSV file. If a filter is applied, only the filtered records will be exported. In the exported CSV file the columns "Recipient" and "Message" are divided into several columns - "Recipient's Zabbix username", "Recipient's name", "Recipient's surname", "Recipient", and "Subject", "Message", "Command".

Using filter

The filter is located below the *Action log* bar. It can be opened and collapsed by clicking the *Filter* tab at the upper-right corner of the page.



You may use the filter to narrow down the records by notification recipients, actions, media types, status, or by the message/remote command content (*Search string*). For better search performance, data is searched with macros unresolved.

Time period selector

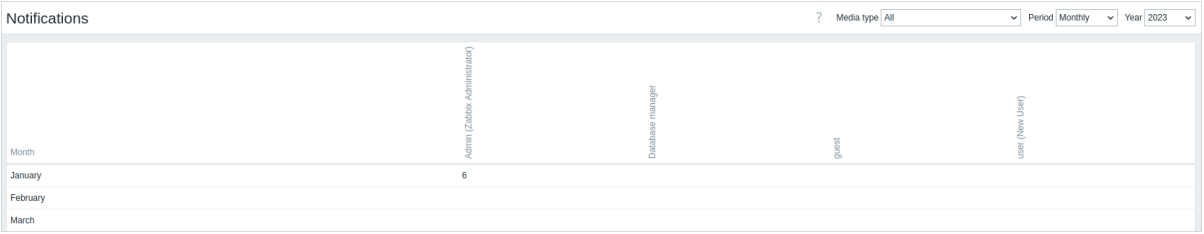
The time period selector next to the filter allows you to select often required periods with one mouse click. For more information, see Time period and host selectors.

Notifications

Overview

In the *Reports* > *Notifications* section a report on the number of notifications sent to each user is displayed.

From the dropdowns in the upper-right corner you can choose the media type (or all), period (data for each day/week/month/year) and year for the notifications sent.



Each column displays totals per one system user.

Data collection

Overview

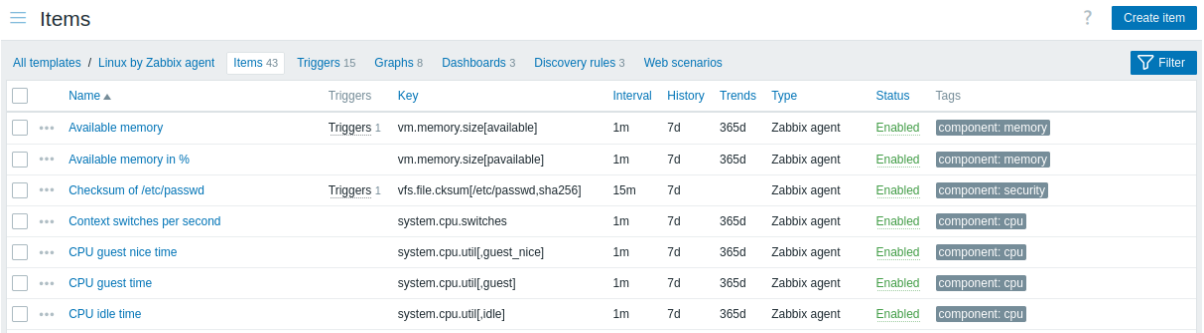
This menu features sections that are related to configuring data collection.

Items

Overview

The item list for a template can be accessed from *Data collection* > *Templates* by clicking *Items* for the respective template.

A list of existing items is displayed.



Displayed data:

Column	Description
Item menu	Click the three-dot icon to open the item menu.
Template	Template the item belongs to. Clicking the template name opens the template configuration form. This column is displayed only if multiple templates or no templates are selected in the filter.
Name	Name of the item displayed as a blue link to item details. Clicking the item name link opens the item configuration form. If the item is inherited from another template, the template name is displayed before the item name as a gray link. Clicking the template link opens the item list on that template level.
Triggers	Moving the mouse over Triggers will display an infobox displaying the triggers associated with the item. The number of the triggers is displayed in gray.
Key	Item key is displayed.
Interval	Frequency of the check is displayed.
History	How many days item data history will be kept is displayed.
Trends	How many days item trends history will be kept is displayed.
Type	Item type is displayed (Zabbix agent, SNMP agent, simple check, etc).
Status	Item status is displayed - *Enabled* or *Disabled*. By clicking the status you can change it - from Enabled to Disabled (and back).
Tags	Item tags are displayed. Up to three tags (name:value pairs) can be displayed. If there are more tags, a "..." link is displayed that allows all tags to be viewed on mouseover.

To configure a new item, click the *Create item* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

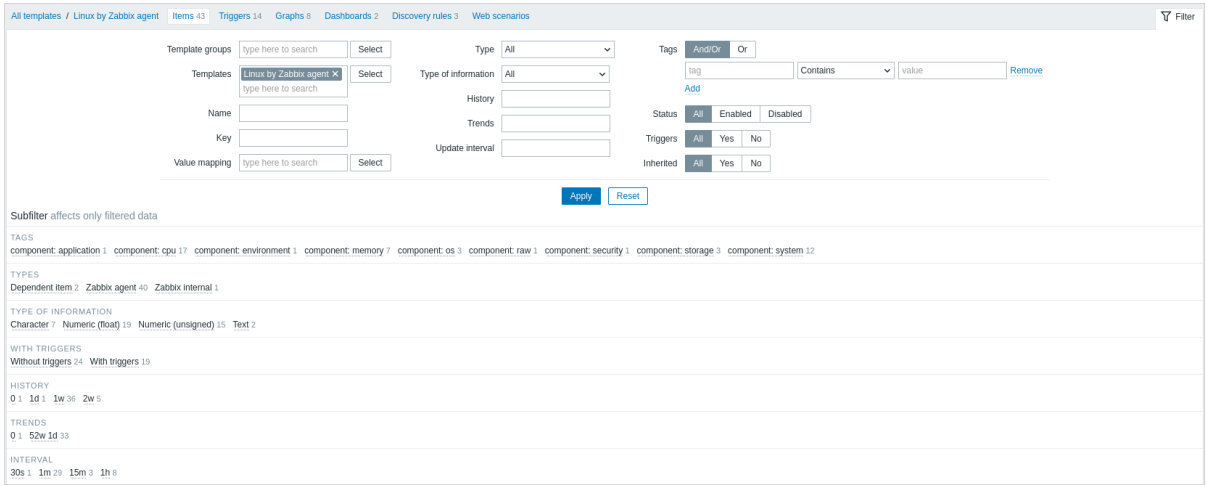
- *Enable* - change item status to *Enabled*.
- *Disable* - change item status to *Disabled*.
- *Copy* - copy the items to other hosts or templates.
- *Mass update* - update several properties for a number of items at once.
- *Delete* - delete the items.

To use these options, mark the checkboxes before the respective items, then click the required button.

Using filter

The item list may contain a lot of items. By using the filter, you can filter out some of them to quickly locate the items you're looking for. For better search performance, data is searched with macros unresolved.

The *Filter* icon is available at the upper-right corner. Clicking it will open a filter where you can specify the desired filtering criteria.



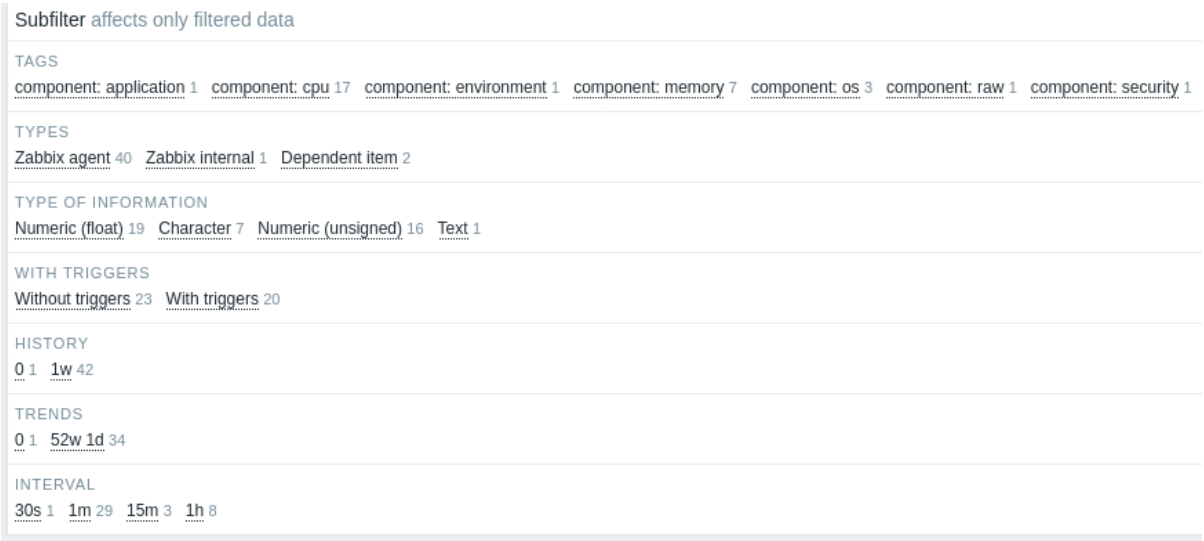
Parameter	Description
Template groups	Filter by one or more template groups. Specifying a parent template group implicitly selects all nested groups.
Templates	Filter by one or more templates.
Name	Filter by item name.

Parameter	Description
Key	Filter by item key.
Value mapping	Filter by the value map used. This parameter is not displayed if the *Templates* option is empty.
Type	Filter by item type (Zabbix agent, SNMP agent, etc.).
Type of information	Filter by type of information (Numeric unsigned, float, etc.).
History	Filter by how long item history is kept.
Trends	Filter by how long item trends are kept.
Update interval	Filter by item update interval.
Tags	Specify tags to limit the number of items displayed. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names **Equals** - include the specified tag names and values (case-sensitive) **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive) **Does not exist** - exclude the specified tag names **Does not equal** - exclude the specified tag names and values (case-sensitive) **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive) There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition **Or** - enough if one condition is met
Status	Filter by item status - *Enabled* or *Disabled*.
Triggers	Filter items with (or without) triggers.
Inherited	Filter items inherited (or not inherited) from linked templates.

Using subfilter

The subfilter allows further modification of the filtering applied by the main filter.

It contains clickable links for a quick access to related items. Items are related by common entity - tag, item type, item update interval, etc. When an entity is clicked, the entity is highlighted with a gray background, and items are immediately filtered (no need to click *Apply* in the main filter). Clicking another entity adds it to the filtered results. Clicking the entity again removes the filtering.



Subfilters are generated based on the filtered data, which is limited to 1000 records. If you want to see more records in the subfilter, you need to increase the value of *Limit for search and filter results* parameter (in *Administration* > *General* > *GUI*).

Unlike the main filter, the subfilter is updated with each table refresh request to always have up-to-date information of available filtering options and their counter numbers.

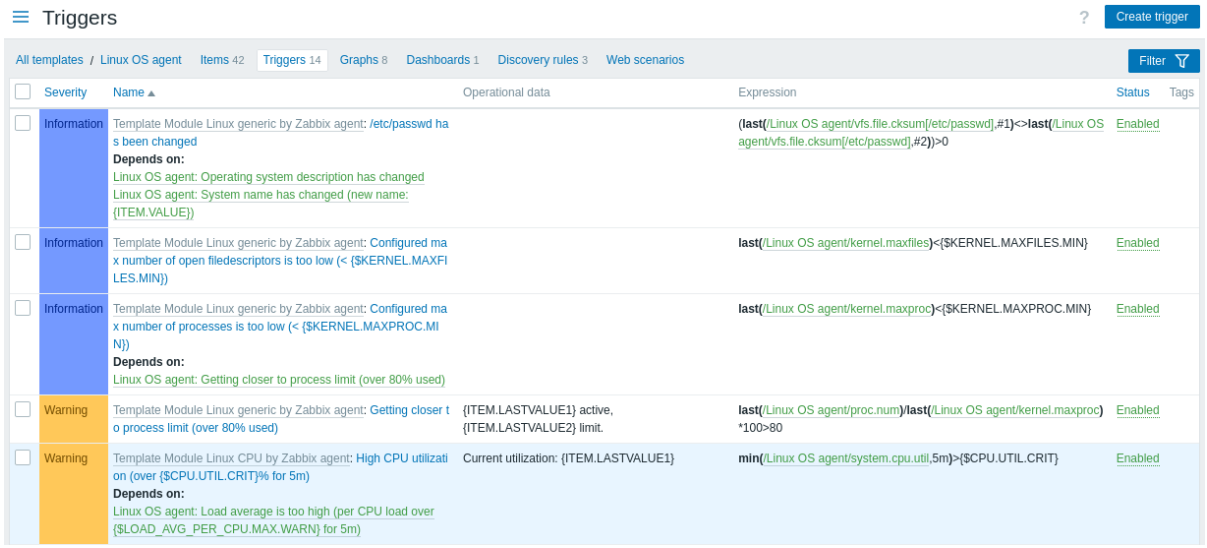
The number of entities displayed is limited to 100 horizontally. If there are more, a three-dot icon is displayed at the end; it is not clickable.

A number next to each clickable entity indicates the number of items grouped in it (based on the results of the main filter). When an entity is clicked, the numbers with other available entities are displayed with a plus sign indicating how many items may be added to the current selection. Entities without items are not displayed unless selected in the subfilter before.

Triggers

Overview

The trigger list for a template can be accessed from *Data collection* > *Templates* by clicking *Triggers* for the respective template.



Displayed data:

Column	Description
Severity	Severity of the trigger is displayed by both name and cell background color.
Template	Template the trigger belongs to. Clicking the template name opens the template configuration form. This column is displayed only if multiple templates or no templates are selected in the filter.
Name	Name of the trigger displayed as a blue link to trigger details. Clicking the trigger name link opens the trigger configuration form. If the trigger is inherited from another template, the template name is displayed before the trigger name as a gray link. Clicking the template link opens the trigger list on that template level.
Operational data	Definition of the trigger's operational data, containing arbitrary strings and macros that are resolved dynamically in *Monitoring* > *Problems* and in the *Problems* and *Problems by severity* dashboard widgets. For details, see the description of the *Operational data* field in trigger configuration.
Expression	Trigger expression is displayed. The template-item part of the expression is displayed as a link, leading to the item configuration form.
Status	Trigger status is displayed - *Enabled* or *Disabled*. By clicking the status you can change it - from Enabled to Disabled (and back).
Tags	If a trigger contains tags, tag name and value are displayed in this column.

To configure a new trigger, click the *Create trigger* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

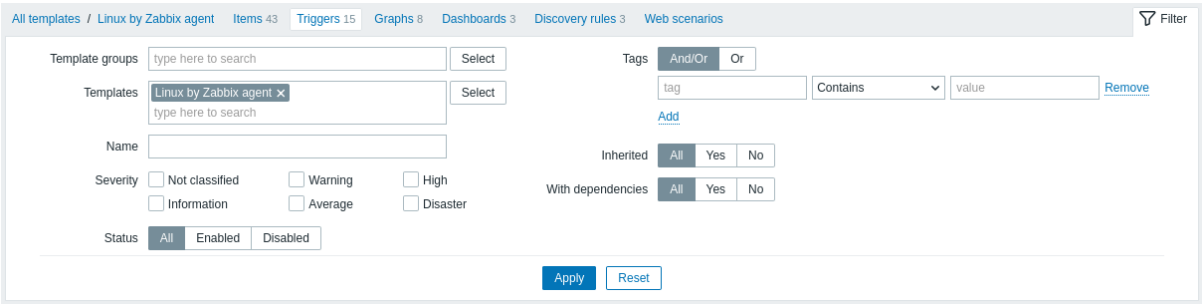
- *Enable* - change trigger status to *Enabled*
- *Disable* - change trigger status to *Disabled*
- *Copy* - copy the triggers to other hosts or templates
- *Mass update* - update several properties for a number of triggers at once
- *Delete* - delete the triggers

To use these options, mark the checkboxes before the respective triggers, then click the required button.

Using filter

You can use the filter to display only the triggers you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* icon is available at the upper-right corner. Clicking it will open a filter where you can specify the desired filtering criteria.



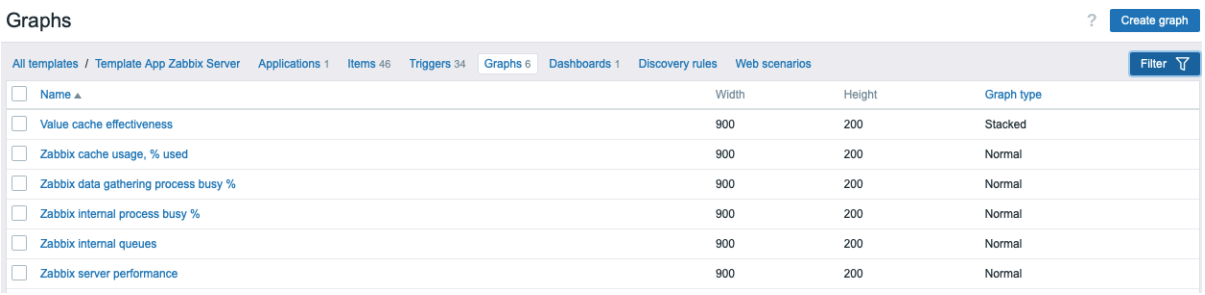
Parameter	Description
Template groups	Filter by one or more template groups. Specifying a parent template group implicitly selects all nested groups.
Templates	Filter by one or more templates. If template groups are already selected above, template selection is limited to those groups.
Name	Filter by trigger name.
Severity	Select to filter by one or several trigger severities.
Status	Filter by trigger activation status (Enabled/Disabled).
Tags	Filter by trigger tag name and value. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names **Equals** - include the specified tag names and values (case-sensitive) **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive) **Does not exist** - exclude the specified tag names **Does not equal** - exclude the specified tag names and values (case-sensitive) **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive) There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition **Or** - enough if one condition is met Macros and macro functions are supported in tag name and tag value fields.
Inherited	Filter triggers inherited (or not inherited) from linked templates.
With dependencies	Filter triggers with (or without) dependencies.

Graphs

Overview

The custom graph list for a template can be accessed from *Data collection* > *Templates* by clicking *Graphs* for the respective template.

A list of existing graphs is displayed.



Displayed data:

Column	Description
Template	Template the graph belongs to. Clicking the template name opens the template configuration form. This column is displayed only if multiple templates or no templates are selected in the filter.
Name	Name of the custom graph, displayed as a blue link to graph details. Clicking the graph name link opens the graph configuration form. If the graph is inherited from another template, the template name is displayed before the graph name as a gray link. Clicking the template link opens the graph list on that template level.
Width	Graph width is displayed.
Height	Graph height is displayed.
Graph type	Graph type is displayed - *Normal*, *Stacked*, *Pie* or *Exploded*.

To configure a new graph, click the *Create graph* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

- *Copy* - copy the graphs to other hosts or templates
- *Delete* - delete the graphs

To use these options, mark the checkboxes before the respective graphs, then click the required button.

Using filter

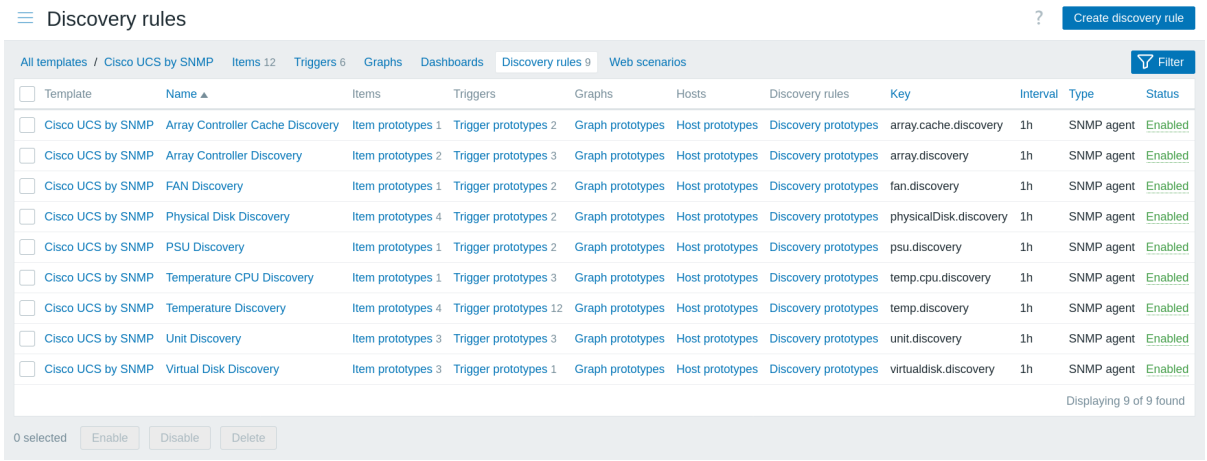
You can filter graphs by template group and template. For better search performance, data is searched with macros unresolved.

Discovery rules

Overview

The list of low-level discovery rules for a template can be accessed from *Data collection* > *Templates* by clicking *Discovery* for the respective template.

A list of existing low-level discovery rules is displayed. It is also possible to see all discovery rules independently of the template, or all discovery rules of a specific template group by changing the filter settings.



Displayed data:

Column	Description
Template	The template discovery rule belongs to. Clicking the template name opens the template configuration form.
Name	Name of the rule, displayed as a blue link. Clicking the rule name opens the low-level discovery rule configuration form. If the discovery rule is inherited from another template, the template name is displayed before the rule name as a gray link. Clicking the template link opens the discovery rule list on that template level.

Column	Description
Items	A link to the list of item prototypes is displayed. The number of existing item prototypes is displayed in gray.
Triggers	A link to the list of trigger prototypes is displayed. The number of existing trigger prototypes is displayed in gray.
Graphs	A link to the list of graph prototypes displayed. The number of existing graph prototypes is displayed in gray.
Hosts	A link to the list of host prototypes displayed. The number of existing host prototypes is displayed in gray.
Key	The item key used for discovery is displayed.
Interval	The frequency of performing discovery is displayed.
Type	The item type used for discovery is displayed (Zabbix agent, SNMP agent, etc).
Status	Discovery rule status is displayed - *Enabled* or *Disabled*. By clicking the status you can change it - from Enabled to Disabled (and back).

To configure a new low-level discovery rule, click the *Create discovery rule* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

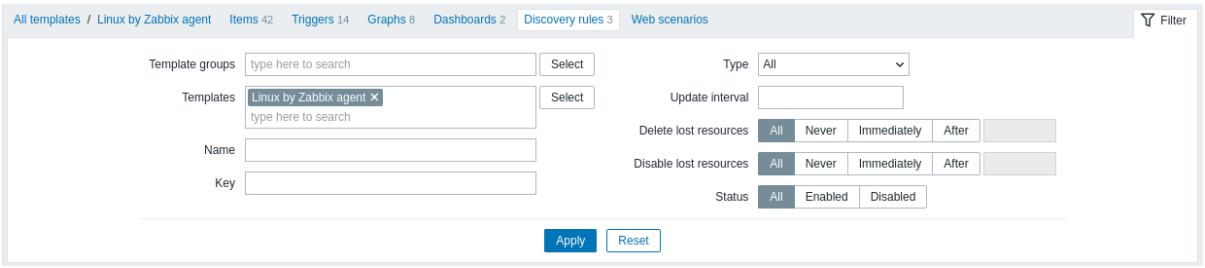
- *Enable* - change the low-level discovery rule status to *Enabled*
- *Disable* - change the low-level discovery rule status to *Disabled*
- *Delete* - delete the low-level discovery rules

To use these options, mark the checkboxes before the respective discovery rules, then click the required button.

Using filter

You can use the filter to display only the discovery rules you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* icon is available at the upper-right corner. Clicking it will open a filter where you can specify the desired filtering criteria such as template, discovery rule name, item key, item type, etc.



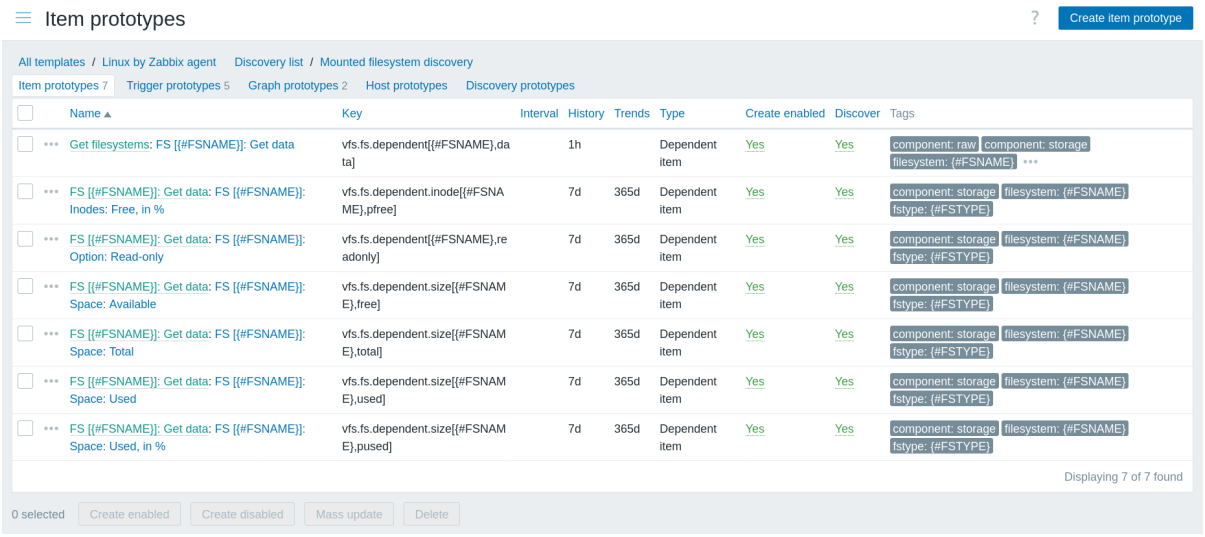
Parameter	Description
Template groups	Filter by one or more template groups. Specifying a parent template group implicitly selects all nested groups.
Templates	Filter by one or more templates.
Name	Filter by discovery rule name.
Key	Filter by discovery item key.
Type	Filter by discovery item type.
Update interval	Filter by update interval. Not available for Zabbix trapper and dependent items.
Delete lost resources	Filter by *Delete lost resources* period.
Disable lost resources	Filter by *Disable lost resources* period.
Status	Filter by discovery rule status (All/Enabled/Disabled).

Item prototypes

Overview

In this section the configured item prototypes of a low-level discovery rule on the template are displayed.

If the template is linked to the host, item prototypes will become the basis of creating real host items during low-level discovery.



Displayed data:

Column	Description
Name	Name of the item prototype, displayed as a blue link. Clicking the name opens the item prototype configuration form. If the item prototype belongs to a linked template, the template name is displayed before the item prototype name as a gray link. Clicking the template link opens the item prototype list on the linked template level. If the item prototype is of the ”Dependent item” type, the master item name is displayed before the item prototype name in green (as in the case of item prototypes on the host level).
Key	Key of the item prototype is displayed.
Interval	Frequency of the check is displayed.
History	How many days to keep item data history is displayed.
Trends	How many days to keep item trends history is displayed.
Type	Type of the item prototype is displayed (Zabbix agent, SNMP agent, simple check, etc).
Create enabled	Create the item based on this prototype as: **Yes** - enabled **No** - disabled. You can switch between ’Yes’ and ’No’ by clicking them.
Discover	Discover the item based on this prototype: **Yes** - discover **No** - do not discover. You can switch between ’Yes’ and ’No’ by clicking them.
Tags	Tags of the item prototype is displayed.

To configure a new item prototype, click the *Create item prototype* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

- *Create enabled* - create these items as *Enabled*
- *Create disabled* - create these items as *Disabled*
- *Mass update* - mass update these item prototypes
- *Delete* - delete these item prototypes

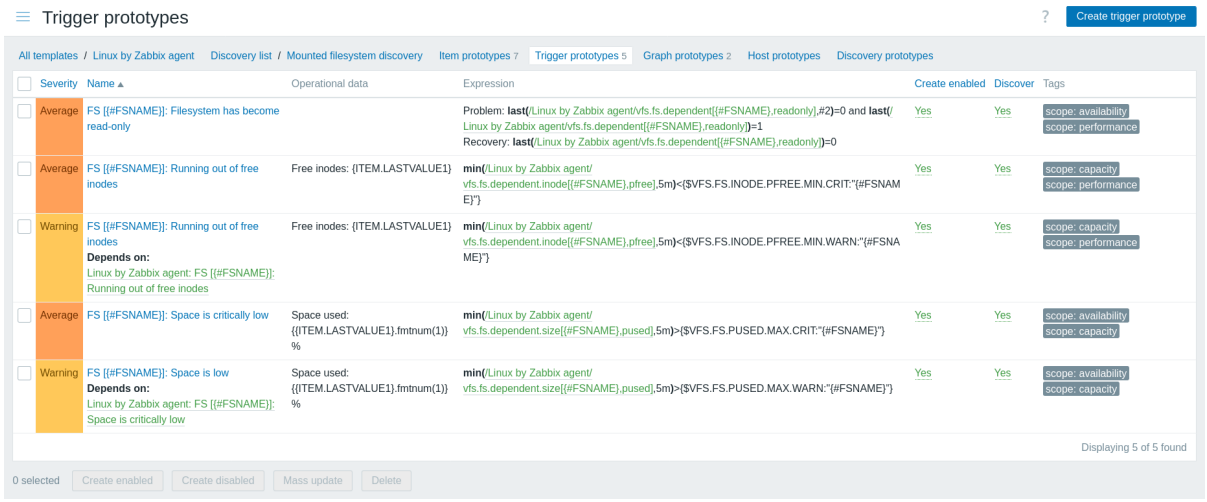
To use these options, mark the checkboxes before the respective item prototypes, then click the required button.

Trigger prototypes

Overview

In this section the configured trigger prototypes of a low-level discovery rule on the template are displayed.

If the template is linked to the host, trigger prototypes will become the basis of creating real host triggers during low-level discovery.



Displayed data:

Column	Description
Severity	Severity of the trigger is displayed by both name and cell background color.
Name	Name of the trigger prototype, displayed as a blue link. Clicking the name opens the trigger prototype configuration form. If the trigger prototype belongs to a linked template, the template name is displayed before the trigger name as a gray link. Clicking the template link opens the trigger prototype list on the linked template level.
Expression	Trigger expression is displayed. The template-item part of the expression is displayed as a link, leading to the item configuration form.
Operational data	Format of the trigger's operational data, containing arbitrary strings and macros that are resolved dynamically in *Monitoring > Problems* and in the *Problems* and *Problems by severity* dashboard widgets. For details, see the description of the *Operational data* field in trigger configuration.
Create enabled	Create the trigger based on this prototype as: **Yes** - enabled **No** - disabled. You can switch between 'Yes' and 'No' by clicking them.
Discover	Discover the trigger based on this prototype: **Yes** - discover **No** - do not discover. You can switch between 'Yes' and 'No' by clicking them.
Tags	Tags of the trigger prototype are displayed.

To configure a new trigger prototype, click the *Create trigger prototype* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

- *Create enabled* - create these triggers as *Enabled*
- *Create disabled* - create these triggers as *Disabled*
- *Mass update* - mass update these trigger prototypes
- *Delete* - delete these trigger prototypes

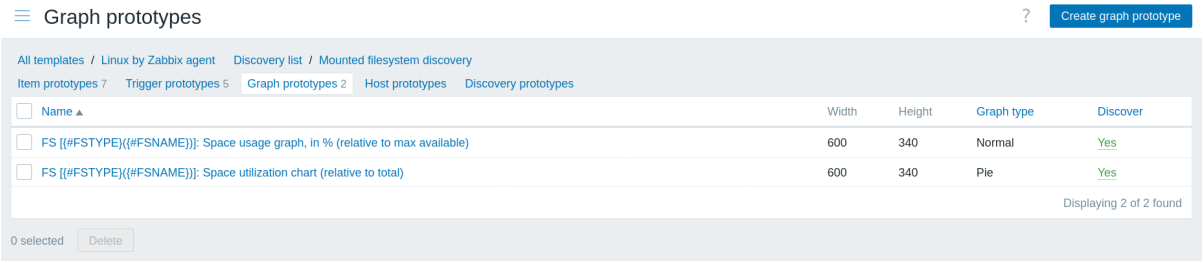
To use these options, mark the checkboxes before the respective trigger prototypes, then click the required button.

Graph prototypes

Overview

In this section the configured graph prototypes of a low-level discovery rule on the template are displayed.

If the template is linked to the host, graph prototypes will become the basis of creating real host graphs during low-level discovery.



Displayed data:

Column	Description
Name	Name of the graph prototype, displayed as a blue link. Clicking the name opens the graph prototype configuration form. If the graph prototype belongs to a linked template, the template name is displayed before the graph name as a gray link. Clicking the template link opens the graph prototype list on the linked template level.
Width	Width of the graph prototype is displayed.
Height	Height of the graph prototype is displayed.
Type	Type of the graph prototype is displayed - *Normal*, *Stacked*, *Pie* or *Exploded*.
Discover	Discover the graph based on this prototype: **Yes** - discover **No** - do not discover. You can switch between 'Yes' and 'No' by clicking them.

To configure a new graph prototype, click the *Create graph prototype* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

- *Delete* - delete these graph prototypes

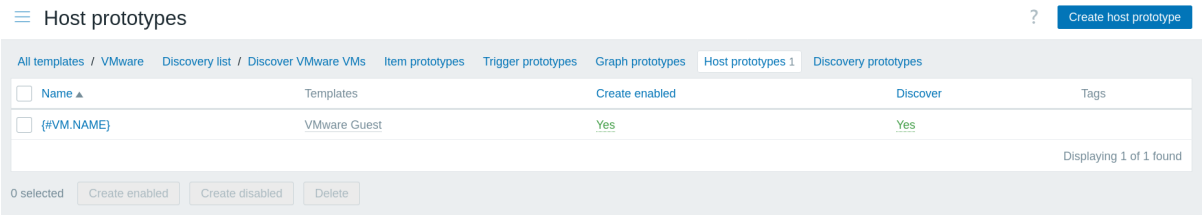
To use these options, mark the checkboxes before the respective graph prototypes, then click the required button.

Host prototypes

Overview

In this section the configured host prototypes of a low-level discovery rule on the template are displayed.

If the template is linked to the host, host prototypes will become the basis of creating real hosts during low-level discovery.



Displayed data:

Column	Description
Name	Name of the host prototype, displayed as a blue link. Clicking the name opens the host prototype configuration form. If the host prototype belongs to a linked template, the template name is displayed before the host name as a gray link. Clicking the template link opens the host prototype list on the linked template level.
Templates	Templates of the host prototype are displayed.
Create enabled	Create the host based on this prototype as: **Yes** - enabled **No** - disabled. You can switch between 'Yes' and 'No' by clicking them.

Column	Description
Discover	Discover the host based on this prototype: **Yes** - discover **No** - do not discover. You can switch between 'Yes' and 'No' by clicking them.
Tags	Tags of the host prototype are displayed.

To configure a new host prototype, click the *Create host prototype* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

- *Create enabled* - create these hosts as *Enabled*
- *Create disabled* - create these hosts as *Disabled*
- *Delete* - delete these host prototypes

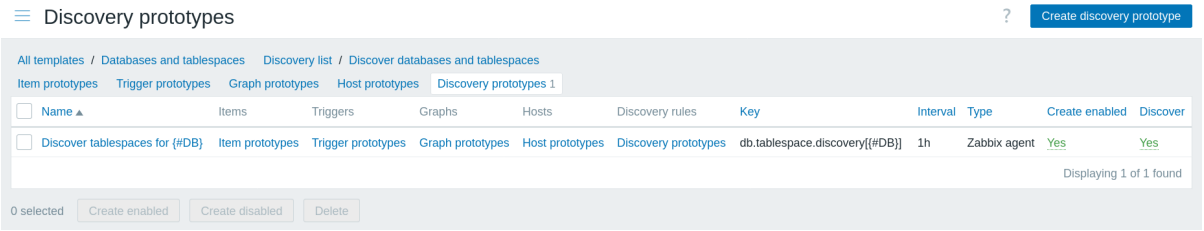
To use these options, mark the checkboxes before the respective host prototypes, then click the required button.

Discovery prototypes

Overview

In this section the configured discovery prototypes of a low-level discovery rule on the template are displayed.

If the template is linked to the host, discovery prototypes will become the basis of creating real discovery prototypes of the parent discovery rule.



Displayed data:

Column	Description
Name	Name of the discovery prototype, displayed as a blue link. Clicking the name opens the discovery prototype configuration form. If the discovery prototype belongs to a linked template, the template name is displayed before the discovery prototype name as a gray link. Clicking the template link opens the discovery prototype list on the linked template level.
Items	A link to the list of item prototypes is displayed. The number of existing item prototypes is displayed in gray.
Triggers	A link to the list of trigger prototypes is displayed. The number of existing trigger prototypes is displayed in gray.
Graphs	A link to the list of graph prototypes displayed. The number of existing graph prototypes is displayed in gray.
Hosts	A link to the list of host prototypes displayed. The number of existing host prototypes is displayed in gray.
Discovery rules	A link to the list of discovery prototypes displayed. The number of existing discovery prototypes is displayed in gray.
Key	The item key used for discovery is displayed.
Interval	The frequency of performing discovery is displayed.
Type	The item type used for discovery is displayed (Zabbix agent, SNMP agent, etc).
Create enabled	Discovery prototype creation status is displayed - create enabled (*Yes*) or create disabled (*No*). By clicking the status you can change it - from Yes to No (and back).
Discover	Discovery prototype discovery status is displayed - discover (*Yes*) or do not discover (*No*). By clicking the status you can change it - from Yes to No (and back).

To configure a new discovery prototype, click the *Create discovery prototype* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

- *Create enabled* - create these discovery prototypes as *Enabled*
- *Create disabled* - create these discovery prototypes as *Disabled*
- *Delete* - delete these discovery prototypes

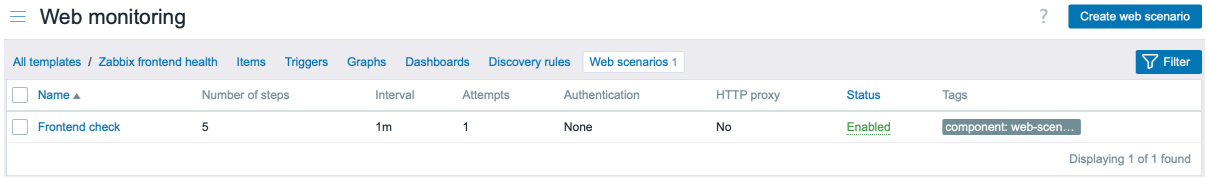
To use these options, mark the checkboxes before the respective discovery prototypes, then click the required button.

Web scenarios

Overview

The web scenario list for a template can be accessed from *Data collection* > *Templates* by clicking *Web* for the respective template.

A list of existing web scenarios is displayed.



Displayed data:

Column	Description
Name	Name of the web scenario. Clicking the web scenario name opens the web scenario configuration form. If the web scenario is inherited from another template, the template name is displayed before the web scenario name as a gray link. Clicking the template link opens the web scenario list on that template level.
Number of steps	The number of steps the scenario contains.
Update interval	How often the scenario is performed.
Attempts	How many attempts for executing web scenario steps are performed.
Authentication	Authentication method is displayed - Basic, NTLM, Kerberos, Digest, or None.
HTTP proxy	Displays HTTP proxy or 'No' if not used.
Status	Web scenario status is displayed - *Enabled* or *Disabled*. By clicking the status you can change it.
Tags	Web scenario tags are displayed. Up to three tags (name:value pairs) can be displayed. If there are more tags, a "..." link is displayed that allows all tags to be viewed on mouseover.

To configure a new web scenario, click the *Create web scenario* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

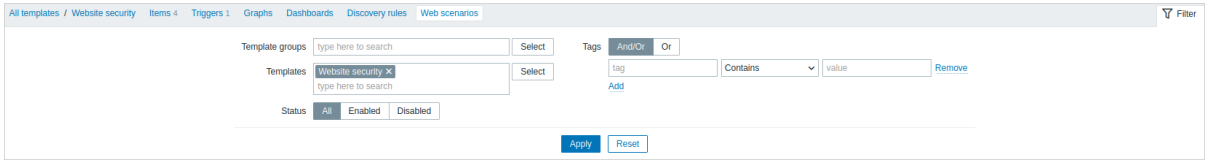
- *Enable* - change the scenario status to *Enabled*
- *Disable* - change the scenario status to *Disabled*
- *Delete* - delete the web scenarios

To use these options, mark the checkboxes before the respective web scenarios, then click the required button.

Using filter

You can use the filter to display only the scenarios you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* link is available above the list of web scenarios. If you click it, a filter becomes available where you can filter scenarios by template group, template, status and tags.

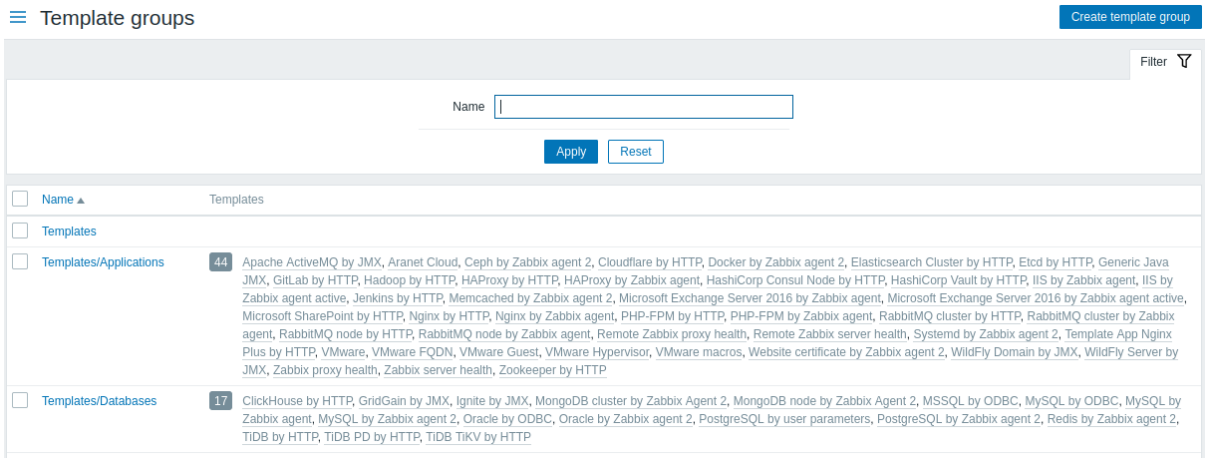


Template groups

Overview

In the *Data collection* > *Templates groups* section users can configure and maintain template groups.

A listing of existing template groups with their details is displayed. You can search and filter template groups by name.



Displayed data:

Column	Description
Name	Name of the template group. Clicking the group name opens the template group configuration form.
Templates	Number of templates in the group (displayed in gray) followed by the list of group members. Clicking a template name will open the template configuration form. Clicking the number opens the list of templates in this group.

Mass editing options

To delete several template groups at once, mark the checkboxes before the respective groups, then click the Delete button below the list.

Using filter

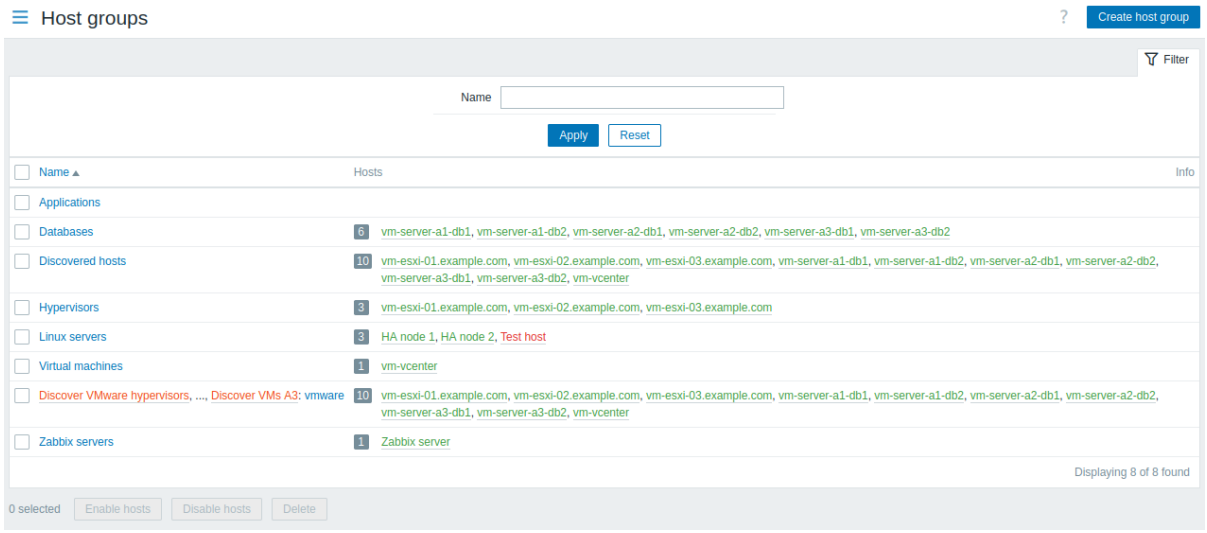
You can use the filter to display only the template groups you are interested in. For better search performance, data is searched with macros unresolved.

Host groups

Overview

In the *Data collection* > *Host groups* section users can configure and maintain host groups.

A listing of existing host groups with their details is displayed. You can search and filter host groups by name.



Displayed data:

Column	Description
Name	Name of the host group. Clicking the group name opens the host group configuration form. Discovered host groups are displayed with low-level discovery rule names as a prefix. Clicking the LLD rule name opens the host prototype configuration form. Note that discovered host groups are deleted when all LLD rules that discovered it are deleted.
Hosts	Number of hosts in the group (displayed in gray) followed by the list of group members. Clicking a host name will open the host configuration form. Clicking the number will, in the whole listing of hosts, filter out those that belong to the group.
Info	Error information (if any) regarding the host group is displayed.

Mass editing options

Buttons below the list offer some mass-editing options:

- *Enable hosts* - change the status of all hosts in the group to "Monitored"
- *Disable hosts* - change the status of all hosts in the group to "Not monitored"
- *Delete* - delete the host groups

To use these options, mark the checkboxes before the respective host groups, then click the required button.

Using filter

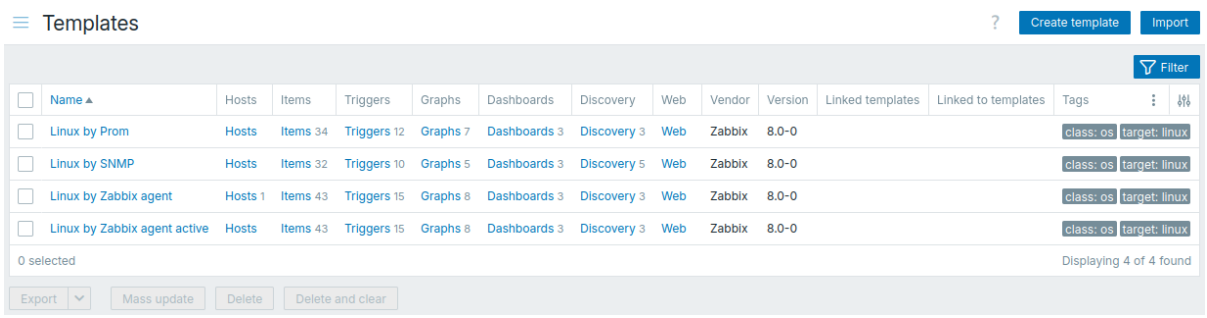
You can use the filter to display only the host groups you are interested in. For better search performance, data is searched with macros unresolved.

Templates

Overview

In the *Data collection* > *Templates* section users can configure and maintain templates.

A listing of existing templates with their details is displayed.



Template details are displayed in a customizable table. The *Customize table* button in the table header lets you show, hide, or reorder columns, and reset the layout to default. You can also resize each column by dragging its right border; double-click the border to reset. All customizations are applied immediately and always remembered for the current user.

The following columns are available:

Column	Description
Name	Name of the template. Clicking the template name opens the template configuration form.
Hosts	Number of editable hosts to which the template is linked; read-only hosts are not included. Clicking *Hosts* will open the host list with only those hosts filtered that are linked to the template.
Entities (Items, Triggers, Graphs, Dashboards, Discovery, Web)	Number of the respective entities in the template (displayed in gray). Clicking the entity name will, in the whole listing of that entity, filter out those that belong to the template.
Linked templates	Templates that are linked to the template.
Linked to templates	Templates that the template is linked to.
Vendor, *Version*	Template vendor and version; displayed if the template configuration contains such information, and only for out-of-the-box templates, imported templates, or templates modified through the Template API. For out-of-the-box templates, version is displayed as follows: major version of Zabbix, delimiter (”-”), revision number (increased with each new version of the template, and reset with each major version of Zabbix). For example, 8.0-0, 8.0-3, 8.2-0, 8.2-3.
Tags	Template tags are displayed, with macros unresolved. Hovering or clicking the three dots icon after the displayed tags opens a pop-up with all tags. Clicking the *Options* button in the column header lets you set the number of tags displayed (1–3), choose how tag names are displayed (full, shortened, or none), set tag display priority, or duplicate the column (duplicated columns can be renamed or deleted).
Tag value	Template tag values are displayed, with macros unresolved. Clicking the *Options* button in the column header lets you rename the column, choose which tag values to display, or duplicate the column (duplicated columns can be renamed or deleted).
Custom text	Custom text, macro values, or a combination of both is displayed. Clicking the *Options* button in the column header lets you rename the column, enter the custom text or macros, or duplicate the column (duplicated columns can be renamed or deleted). Supported macros: `{HOST.ID}`, `{HOST.HOST}`, `{HOST.NAME}`, `{HOST.DESCRIPTION}`, and user macros.

To configure a new template, click the *Create template* button in the upper-right corner.

To import a template from a YAML, XML, or JSON file, click the *Import* button in the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

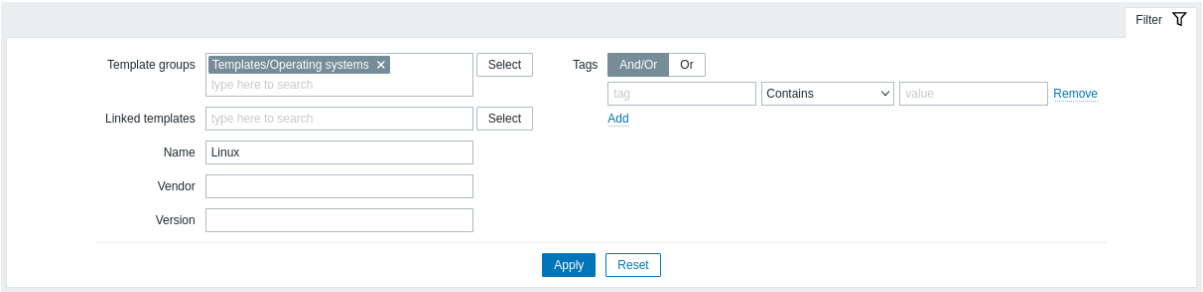
- *Export* - export the template to a YAML, XML or JSON file;
- *Mass update* - update several properties for a number of templates at once;
- *Delete* - delete the template while leaving its linked entities (items, triggers etc.) with the hosts;
- *Delete and clear* - delete the template and its linked entities from the hosts.

To use these options, mark the checkboxes before the respective templates, then click the required button.

Using filter

You can use the filter to display only the templates you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* link is available below *Create template* and *Import* buttons. If you click it, a filter becomes available where you can filter templates by template group, linked templates, name and tags.



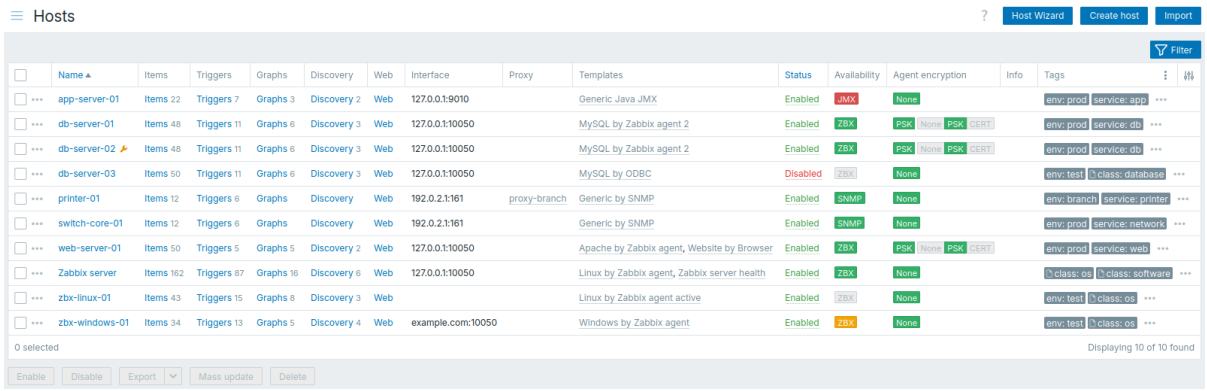
Parameter	Description
Template groups	Filter by one or more template groups. Specifying a parent template group implicitly selects all nested groups.
Linked templates	Filter by directly linked templates.
Name	Filter by template name.
Vendor	Filter by template vendor.
Version	Filter by template version.
Tags	Filter by template tag name and value. Filtering is possible only by template-level tags (not inherited ones). It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met.

Hosts

Overview

In the *Data collection* > *Hosts* section users can configure and maintain hosts.

A listing of existing hosts with their details is displayed.



Host details are displayed in a customizable table. The *Customize table* button in the table header lets you show, hide, or reorder columns, and reset the layout to default. You can also resize each column by dragging its right border; double-click the border to reset. All customizations are applied immediately and always remembered for the current user.

The following columns are available:

Column	Description
Name	Name of the host. Clicking the host name opens the host configuration form. An orange wrench icon after the host name indicates that this host is in maintenance. Maintenance details are displayed when the mouse pointer is positioned on the icon.
Entities (Items, Triggers, Graphs, Discovery, Web)	Clicking the entity name will display items, triggers, etc., of the host. The number of the respective entities is displayed in gray.
Interface	The main interface of the host is displayed.

Column	Description
Proxy	Assigned proxies are displayed in this column: <Proxy name> - host is monitored by standalone proxy (even if the proxy is part of proxy group); <Proxy group name: proxy name> - host is monitored by proxy group, and Zabbix server has assigned a proxy to monitor the host; <Proxy group name> - host is monitored by proxy group without any proxies, or if Zabbix server has not assigned a proxy to monitor the host; Nothing - host is not monitored by either proxy or proxy group. This column is only displayed if the *Monitored by* filter option is set to 'Any', 'Proxy', or 'Proxy group'.
Templates	The templates linked to the host are displayed. If other templates are contained in the linked template, those are displayed in parentheses, separated by a comma. Clicking a template name will open its configuration form.
Status	Host status is displayed - *Enabled* or *Disabled*. By clicking the status you can change it manually. Discovered hosts that have been lost are marked with an info icon. The tooltip text provides details on their status.
Availability	Host availability per configured interface is displayed. Availability icons represent host interface current status on Zabbix server. Therefore, if you disable a host in the frontend, its availability will update after Zabbix server has synchronized the configuration changes. Similarly, if you enable a host, its availability will update after Zabbix server has synchronized the configuration changes and polled the host. Availability icons represent only those interface types (Agent, SNMP, IPMI, JMX) that are configured. Hovering over the icon displays a pop-up with a list of all interfaces of the same type with details, status, and errors. For Agent interface, the pop-up displays interfaces (passive) and active checks. If a host has active checks only, the Agent interface icon is displayed even if the host does not have an Agent interface configured. The column is empty for hosts with no interfaces. The status of a single host interface is determined by the connection between an enabled item using the interface and the host. The status can be: **Available** - the connection to the host was successful; **Not available** - the connection to the host was unsuccessful (timeout, firewall issues, etc.); **Unknown** - the connection to the host has not been attempted or the result is unknown. For additional details on how Zabbix server determines interface status, see Unknown interface status and Unreachable/unavailable host interface settings. The status of all host interfaces of a single type (Agent, SNMP, IPMI, JMX) is determined by those interfaces that are used by at least one enabled item. The status is indicated by the icon color: **Green** - all interfaces are available; **Yellow** - at least one interface is not available, and at least one is available or unknown; **Red** - all interfaces are not available; **Gray** - at least one interface is unknown, but none are not available. **Active check availability.** If at least one active check is enabled on the host, active checks also affect the total Agent interface availability as described above. To determine active check availability, heartbeat messages are sent in the agent active check thread. The frequency of heartbeat messages is controlled by the `HeartbeatFrequency` parameter in Zabbix agent or agent 2 configuration (default 60 seconds, range 0-3600). Active checks are considered not available when the active check heartbeat is older than 2 x HeartbeatFrequency seconds. **Note:** Zabbix agents older than version 6.2.x do not send active check heartbeats, so the availability of their hosts remains unknown.
Agent encryption	Encryption status for connections to the host and from the host is displayed: **None** - no encryption; **PSK** - using pre-shared key; **Cert** - using certificate.
Info	Error information (if any) regarding the host is displayed.

Column	Description
Tags	Host tags are displayed, with macros unresolved. Hovering or clicking the three dots icon after the displayed tags opens a pop-up with all tags. Clicking the *Options* button in the column header lets you set the number of tags displayed (1–3), choose how tag names are displayed (full, shortened, or none), set tag display priority, or duplicate the column (duplicated columns can be renamed or deleted).
Tag value	Host tag values are displayed, with macros unresolved. Clicking the *Options* button in the column header lets you rename the column, choose which tag values to display, or duplicate the column (duplicated columns can be renamed or deleted).
Custom text	Custom text, macro values, or a combination of both is displayed. Clicking the *Options* button in the column header lets you rename the column, enter the custom text or macros, or duplicate the column (duplicated columns can be renamed or deleted). Supported macros: `{HOST.*}` (except `{HOST.METADATA}`), `{INVENTORY.*}`, and user macros.

To open the host menu, click the *Options* button before the host name.

To configure a new host, click the *Host Wizard* or *Create host* button in the upper-right corner.

To import a host from a YAML, XML, or JSON file, click the *Import* button in the upper-right corner.

Mass editing options

Buttons below the list offer some mass editing options:

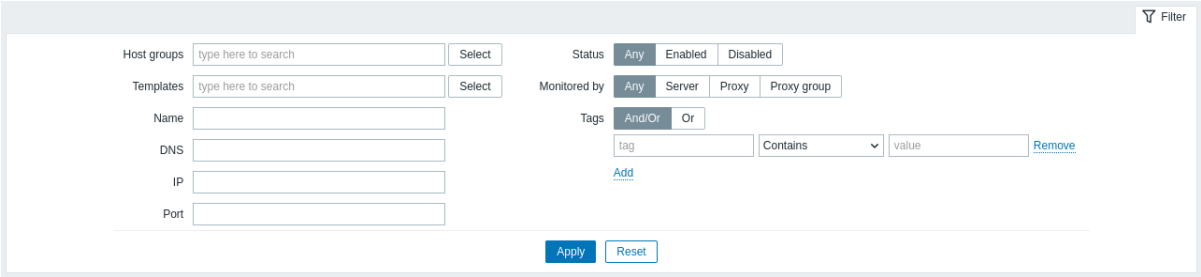
- *Enable* - change host status to *Monitored*;
- *Disable* - change host status to *Not monitored*;
- *Export* - export the hosts to a YAML, XML or JSON file;
- *Mass update* - update several properties for a number of hosts at once;
- *Delete* - delete the hosts.

To use these options, mark the checkboxes before the respective hosts, then click the required button.

Using filter

You can use the filter to display only the hosts you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* icon is available at the upper-right corner. Clicking it will open a filter where you can specify the desired filtering criteria.



Parameter	Description
Host groups	Filter by one or more host groups. Specifying a parent host group implicitly selects all nested host groups.
Templates	Filter by linked templates.
Name	Filter by visible host name.
DNS	Filter by DNS name.
IP	Filter by IP address.
Port	Filter by port number.
Status	Filter by host status.
Monitored by	Filter hosts that are monitored by Zabbix server, proxy, or proxy group.
Proxies	Filter hosts that are monitored by the proxies specified here. This field is only available if "Proxy" is selected in the *Monitored by* field.
Proxy groups	Filter hosts that are monitored by the proxy groups specified here. This field is only available if "Proxy group" is selected in the *Monitored by* field.

Parameter	Description
Tags	Filter by host tag name and value. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met.

Unknown interface status

Zabbix server displays the ”Unknown” status for a host interface (Agent, SNMP, IPMI, JMX) in the following cases:

- Host is disabled.
- Host is set to be monitored by proxy, different proxy, or server if it was previously monitored by proxy.
- Host is monitored by a proxy that appears to be offline (no updates received from the proxy during the maximum heartbeat interval - 1 hour).
- All host items with that interface type are disabled.
- No pollers for that interface type have been configured (for example, the `StartAgentPollers` or `StartSNMPPollers` server configuration parameter is set to 0).

Interface availability is set to ”Unknown” after Zabbix server configuration cache synchronization.

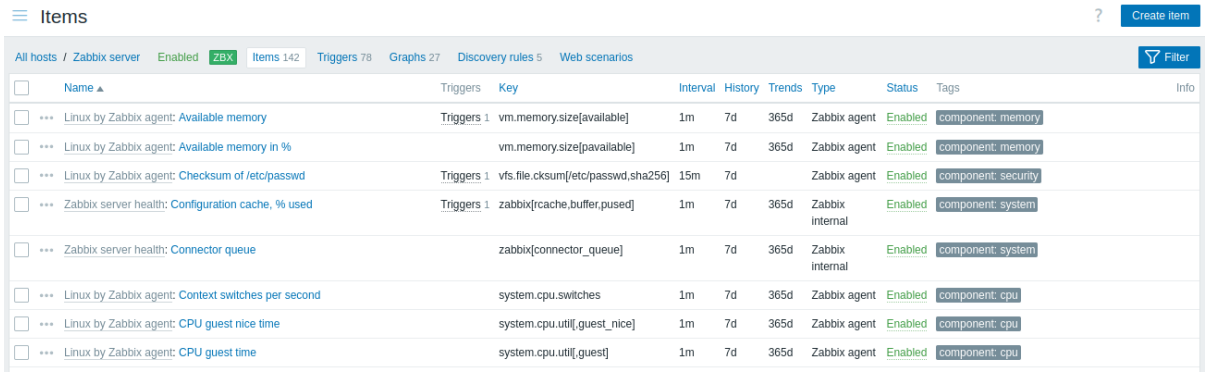
Interface availability (Available/Not available) on hosts monitored by proxies is restored after proxy configuration cache synchronization.

Items

Overview

The item list for a host can be accessed from *Data collection* > *Hosts* by clicking *Items* for the respective host.

A list of existing items is displayed.



Displayed data:

Column	Description
Item menu	Click the three-dot icon to open the item menu.

Column	Description
Host	Host of the item. Clicking the host name opens the host configuration form. This column is displayed only if multiple hosts or no hosts are selected in the filter.
Name	Name of the item displayed as a blue link to item details. Clicking the item name link opens the item configuration form. If the host item belongs to a template, the template name is displayed before the item name as a gray link. Clicking the template link will open the item list on the template level. If the item has been created from an item prototype, its name is preceded by the low-level discovery rule name, in orange. Clicking the discovery rule name will open the item prototype list.
Triggers	Moving the mouse over Triggers will display an infobox displaying the triggers associated with the item. The number of the triggers is displayed in gray.
Key	Item key is displayed.
Interval	Frequency of the check is displayed. *Note* that passive items can also be checked immediately by pushing the *Execute now* button.
History	How many days item data history will be kept is displayed.
Trends	How many days item trends history will be kept is displayed.
Type	Item type is displayed (Zabbix agent, SNMP agent, simple check, etc).
Status	Item status is displayed - *Enabled*, *Disabled* or *Not supported*. You can change the status manually by clicking it - from Enabled to Disabled (and back); from Not supported to Disabled (and back). Discovered items that have been lost are marked with an info icon. The tooltip text provides details on their status.
Tags	Item tags are displayed. Up to three tags (name:value pairs) can be displayed. If there are more tags, a "..." link is displayed that allows all tags to be viewed on mouseover.
Info	If the item is working correctly, no icon is displayed in this column. In case of errors, a square icon with the letter "i" is displayed. Hover over the icon to see a tooltip with the error description.

To configure a new item, click the *Create item* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

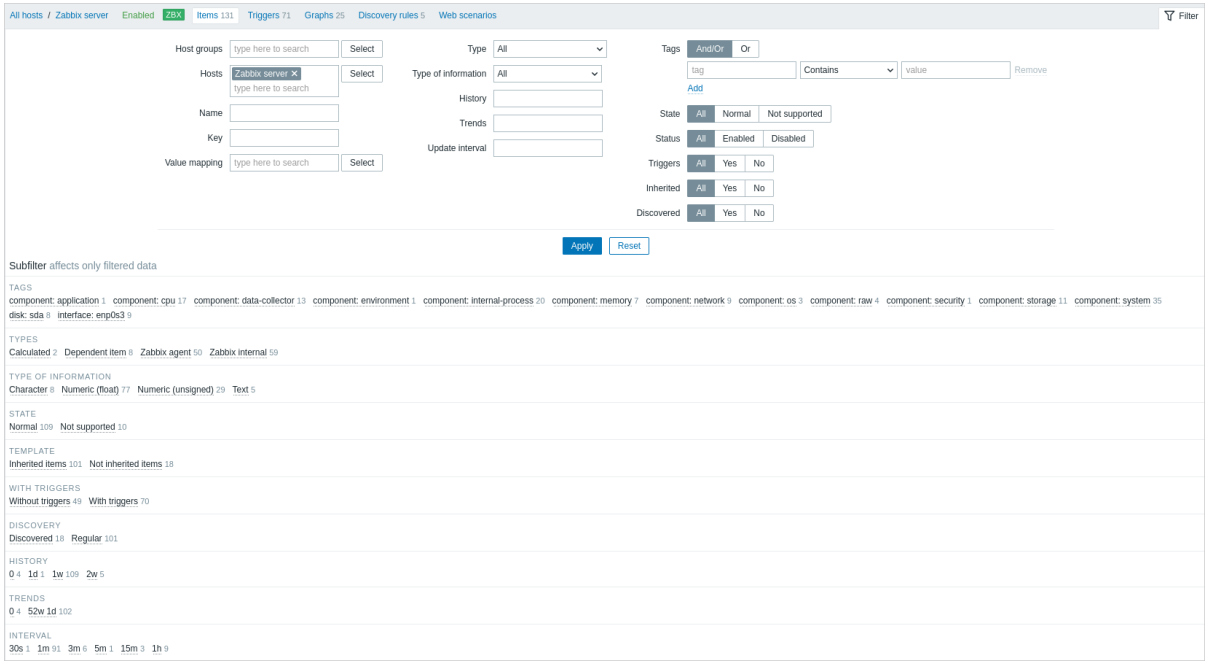
- *Enable* - change item status to *Enabled*
- *Disable* - change item status to *Disabled*
- *Execute now* - execute a check for new item values immediately. Supported for **passive** checks only (see more details). Note that when checking for values immediately, configuration cache is not updated, thus the values will not reflect very recent changes to item configuration.
- *Clear history and trends* - delete history and trend data for items.
- *Copy* - copy the items to other hosts or templates.
- *Mass update* - update several properties for a number of items at once.
- *Delete* - delete the items.

To use these options, mark the checkboxes before the respective items, then click the required button.

Using filter

You can use the filter to display only the items you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* icon is available at the upper-right corner. Clicking it will open a filter where you can specify the desired filtering criteria.



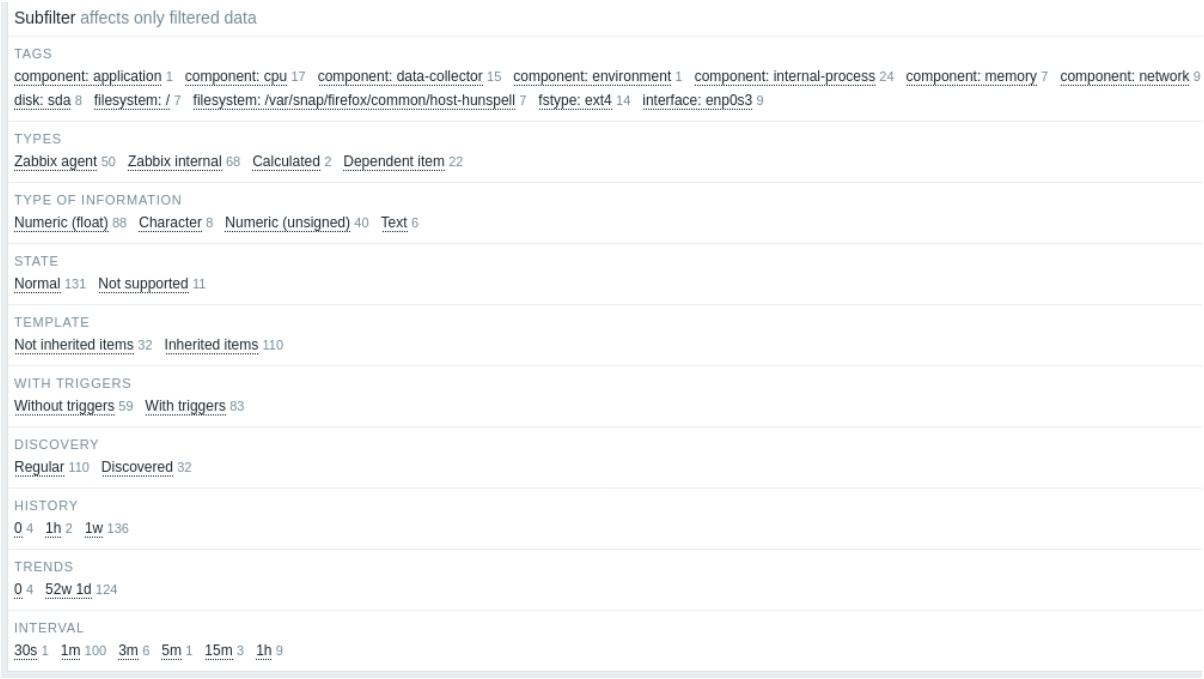
Parameter	Description
Host groups	Filter by one or more host groups. Specifying a parent host group implicitly selects all nested host groups. Host groups containing templates only cannot be selected.
Hosts	Filter by one or more hosts.
Name	Filter by item name.
Key	Filter by item key.
Value mapping	Filter by the value map used. This parameter is not displayed if the *Hosts* option is empty.
Type	Filter by item type (Zabbix agent, SNMP agent, etc.).
Type of information	Filter by type of information (Numeric unsigned, float, etc.).
History	Filter by how long item history is kept.
Trends	Filter by how long item trends are kept.
Update interval	Filter by item update interval.
Tags	Specify tags to limit the number of items displayed. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names **Equals** - include the specified tag names and values (case-sensitive) **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive) **Does not exist** - exclude the specified tag names **Does not equal** - exclude the specified tag names and values (case-sensitive) **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive) There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition **Or** - enough if one condition is met
State	Filter by item state - *Normal* or *Not supported*.
Status	Filter by item status - *Enabled* or *Disabled*.
Triggers	Filter items with (or without) triggers.
Inherited	Filter items inherited (or not inherited) from a template.
Discovery	Filter items discovered (or not discovered) by low-level discovery.

Using subfilter

The subfilter further modification of the filtering applied by the main filter.

It contains clickable links for a quick access to related items. Items are related by common entity - tag, item type, item state, item

update interval, etc. When an entity is clicked, the entity is highlighted with a gray background, and items are immediately filtered (no need to click *Apply* in the main filter). Clicking another entity adds it to the filtered results. Clicking the entity again removes the filtering.



Subfilters are generated based on the filtered data, which is limited to 1000 records. If you want to see more records in the subfilter, you need to increase the value of *Limit for search and filter results* parameter (in *Administration > General > GUI*).

Unlike the main filter, the subfilter is updated with each table refresh request to always have up-to-date information of available filtering options and their counter numbers.

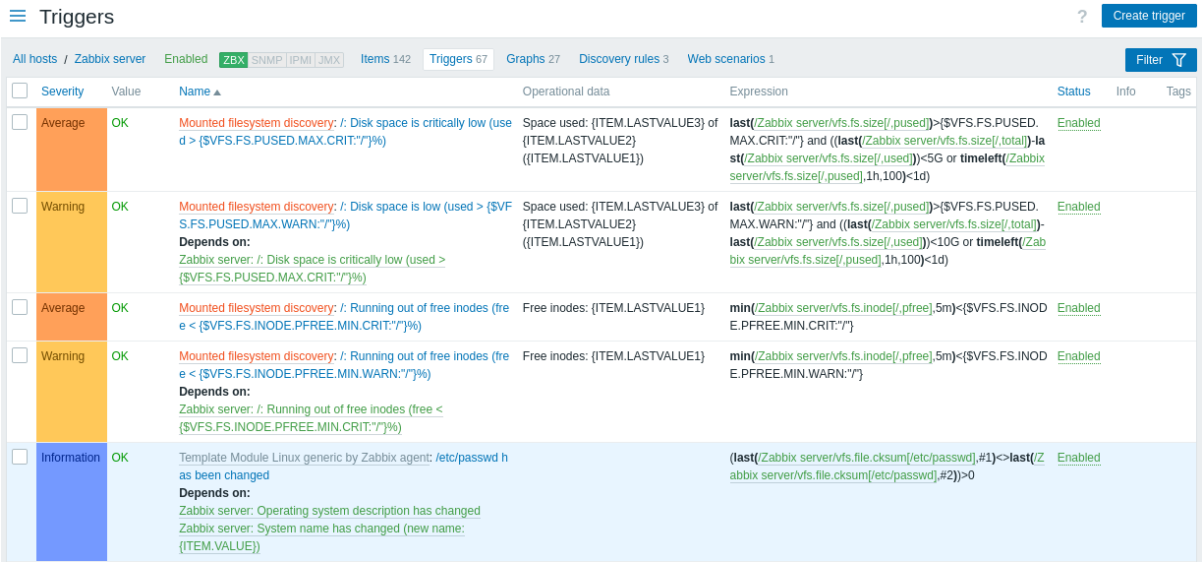
The number of entities displayed is limited to 100 horizontally. If there are more, a three-dot icon is displayed at the end; it is not clickable.

A number next to each clickable entity indicates the number of items grouped in it (based on the results of the main filter). When an entity is clicked, the numbers with other available entities are displayed with a plus sign indicating how many items may be added to the current selection. Entities without items are not displayed unless selected in the subfilter before.

Triggers

Overview

The trigger list for a host can be accessed from *Data collection > Hosts* by clicking *Triggers* for the respective host.



Displayed data:

Column	Description
Severity	Severity of the trigger is displayed by both name and cell background color.
Value	Trigger state is displayed: **OK** - OK state **PROBLEM** - Problem state
Host	Host of the trigger. Clicking the host name opens the host configuration form. This column is displayed only if multiple hosts or no hosts are selected in the filter.
Name	Name of the trigger, displayed as a blue link to trigger details. Clicking the trigger name link opens the trigger configuration form. If the host trigger belongs to a template, the template name is displayed before the trigger name as a gray link. Clicking the template link will open the trigger list on the template level. If the trigger has been created from a trigger prototype, its name is preceded by the low-level discovery rule name, in orange. Clicking the discovery rule name opens the trigger prototype list.
Operational data	Definition of the trigger's operational data, containing arbitrary strings and macros that are resolved dynamically in *Monitoring* > *Problems* and in the *Problems* and *Problems by severity* dashboard widgets. For details, see the description of the *Operational data* field in trigger configuration.
Expression	Trigger expression is displayed. The host-item part of the expression is displayed as a link, leading to the item configuration form.
Status	Trigger status is displayed - *Enabled*, *Disabled*, or *Unknown*. By clicking the status you can manually change it - from Enabled to Disabled (and back); from Unknown to Disabled (and back). Problems of a disabled trigger are no longer displayed in the frontend, but are not deleted. Discovered triggers that have been lost are marked with an info icon. The tooltip text provides details on their status.
Info	If everything is working correctly, no icon is displayed in this column. In case of errors, a square icon with the letter "i" is displayed. Hover over the icon to see a tooltip with the error description.
Tags	If a trigger contains tags, tag name and value are displayed in this column.

To configure a new trigger, click the *Create trigger* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

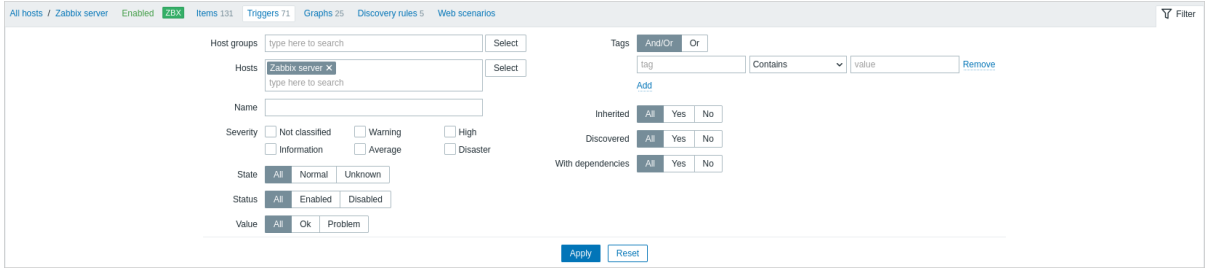
- *Enable* - change trigger status to *Enabled*.
- *Disable* - change trigger status to *Disabled*.
- *Copy* - copy the triggers to other hosts or templates.
- *Mass update* - update several properties for a number of triggers at once.
- *Delete* - delete the triggers.

To use these options, mark the checkboxes before the respective triggers, then click the required button.

Using filter

You can use the filter to display only the triggers you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* icon is available at the upper-right corner. Clicking it will open a filter where you can specify the desired filtering criteria.



Parameter	Description
Host groups	Filter by one or more host groups. Specifying a parent host group implicitly selects all nested host groups. Host groups containing templates only cannot be selected.
Hosts	Filter by one or more hosts. If host groups are already selected above, host selection is limited to those groups.

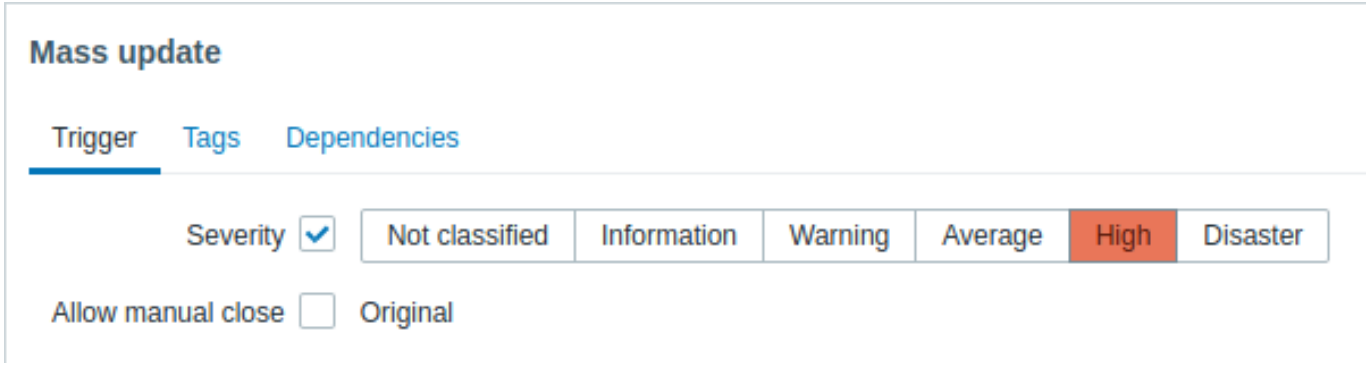
Parameter	Description
Name	Filter by trigger name.
Severity	Select to filter by one or several trigger severities.
State	Filter by trigger expression state (Normal/Unknown).
Status	Filter by trigger activation status (Enabled/Disabled).
Value	Filter by trigger value.
Tags	Filter by trigger tag name and value. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names **Equals** - include the specified tag names and values (case-sensitive) **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive) **Does not exist** - exclude the specified tag names **Does not equal** - exclude the specified tag names and values (case-sensitive) **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive) There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition **Or** - enough if one condition is met Macros and macro functions are supported both in tag name and tag value fields.
Inherited	Filter triggers inherited (or not inherited) from a template.
Discovered	Filter triggers discovered (or not discovered) by low-level discovery.
With dependencies	Filter triggers with (or without) dependencies.

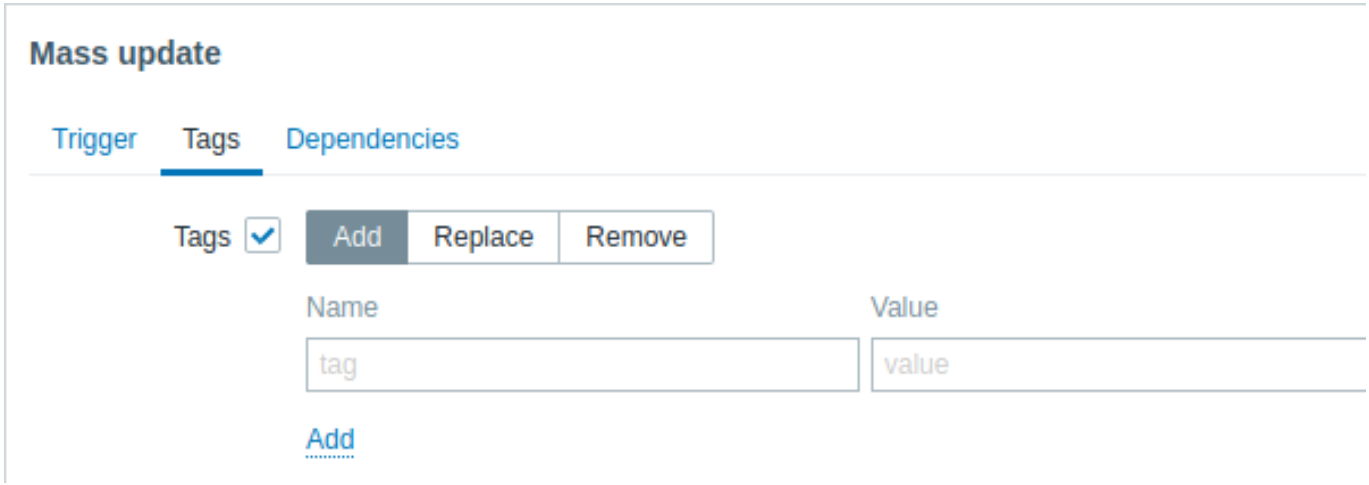
Using mass update

With mass update you may change some attribute for a number of triggers at once, saving you the need to open each individual trigger for editing.

To mass-update some triggers, do the following:

- Mark the checkboxes of the triggers you want to update in the list
- Click *Mass update* below the list
- Navigate to the tab with required attributes (*Trigger*, *Tags* or *Dependencies*)
- Mark the checkboxes of any attribute to update

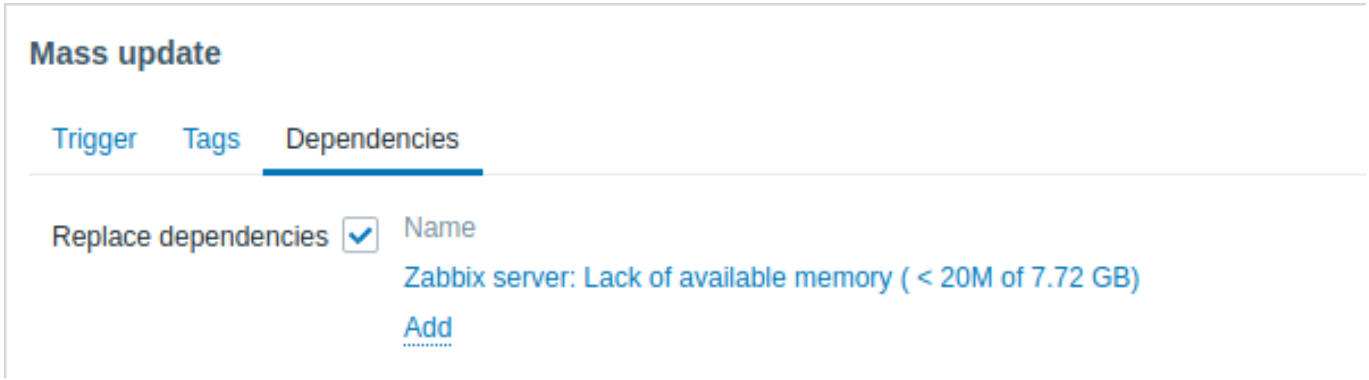




The following options are available when selecting the respective button for tag update:

- *Add* - add new tags for the triggers;
- *Replace* - remove any existing tags from the trigger and replace them with the one(s) specified below;
- *Remove* - remove specified tags from triggers.

Note that tags with the same name but different values are not considered 'duplicates' and can be added to the same trigger.



Replace dependencies - will remove any existing dependencies from the trigger and replace them with the one(s) specified.

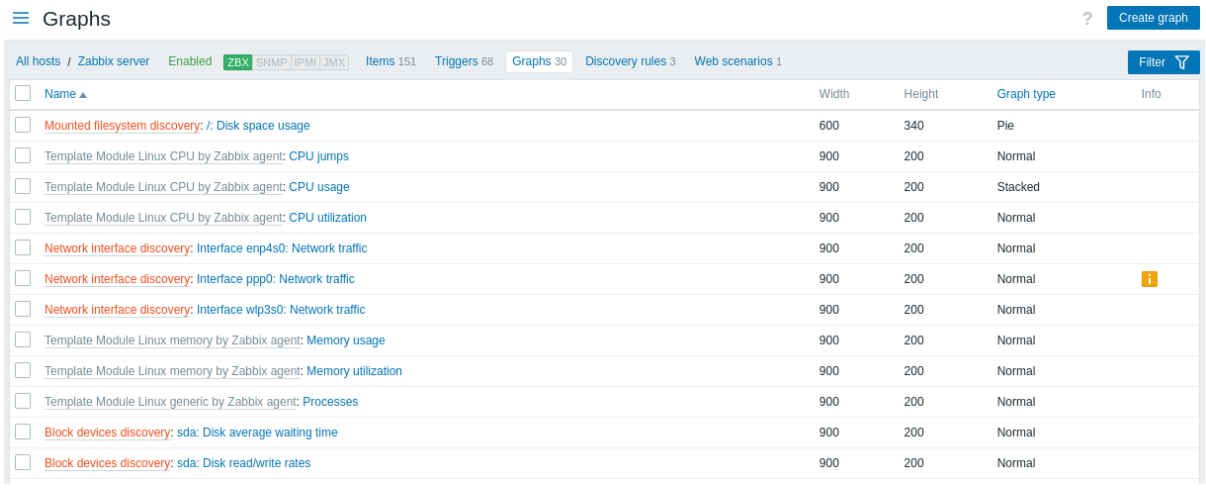
Click *Update* to apply the changes.

Graphs

Overview

The custom graph list for a host can be accessed from *Data collection* > *Hosts* by clicking *Graphs* for the respective host.

A list of existing graphs is displayed.



Displayed data:

Column	Description
Host	Host of the graph. Clicking the host name opens the host configuration form. This column is displayed only if multiple hosts or no hosts are selected in the filter.
Name	Name of the custom graph, displayed as a blue link to graph details. Clicking the graph name link opens the graph configuration form. If the host graph belongs to a template, the template name is displayed before the graph name as a gray link. Clicking the template link will open the graph list on the template level. If the graph has been created from a graph prototype, its name is preceded by the low-level discovery rule name, in orange. Clicking the discovery rule name opens the graph prototype list.
Width	Graph width is displayed.
Height	Graph height is displayed.
Graph type	Graph type is displayed - *Normal*, *Stacked*, *Pie* or *Exploded*.
Info	If the graph is working correctly, no icon is displayed in this column. In case of errors, a square icon with the letter "i" is displayed. Hover over the icon to see a tooltip with the error description.

To configure a new graph, click the *Create graph* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

- *Copy* - copy the graphs to other hosts or templates
- *Delete* - delete the graphs

To use these options, mark the checkboxes before the respective graphs, then click the required button.

Using filter

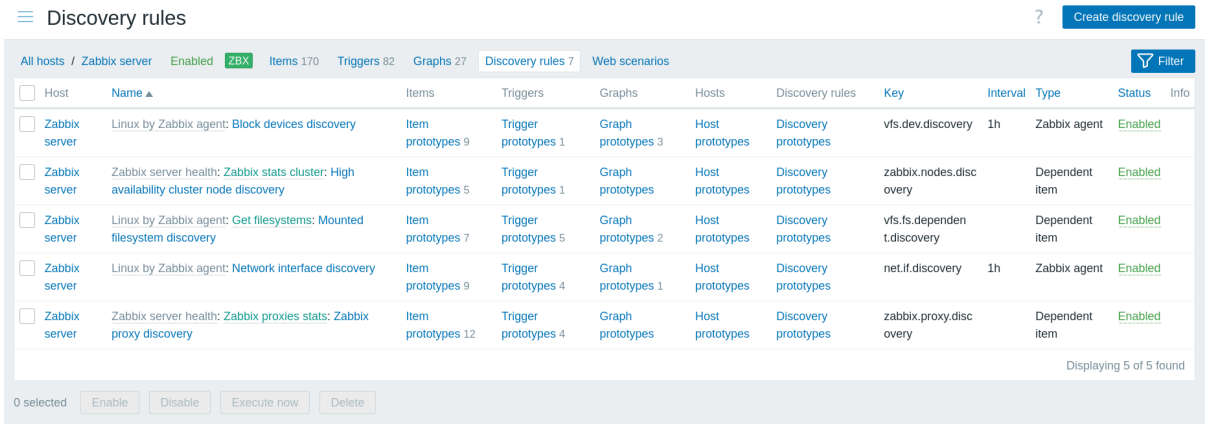
You can filter graphs by host group and host. For better search performance, data is searched with macros unresolved.

Discovery rules

Overview

The list of low-level discovery rules for a host can be accessed from *Data collection* > *Hosts* by clicking *Discovery* for the respective host.

A list of existing low-level discovery rules is displayed. It is also possible to see all discovery rules independently of the host, or all discovery rules of a specific host group by changing the filter settings.



Displayed data:

Column	Description
Host	The visible host name is displayed. Clicking the host name opens the host configuration form. In the absence of a visible host name, the technical host name is displayed.

Column	Description
Name	Name of the rule, displayed as a blue link. Clicking the rule name opens the low-level discovery rule configuration form. If the discovery rule belongs to a template, the template name is displayed before the rule name as a gray link. Clicking the template link opens the rule list on the template level.
Items	A link to the list of item prototypes is displayed. The number of existing item prototypes is displayed in gray.
Triggers	A link to the list of trigger prototypes is displayed. The number of existing trigger prototypes is displayed in gray.
Graphs	A link to the list of graph prototypes is displayed. The number of existing graph prototypes is displayed in gray.
Hosts	A link to the list of host prototypes is displayed. The number of existing host prototypes is displayed in gray.
Key	The item key used for discovery is displayed.
Interval	The frequency of performing discovery is displayed. *Note* that discovery can also be performed immediately by pushing the *Execute now* button below the list.
Type	The item type used for discovery is displayed (Zabbix agent, SNMP agent, etc).
Status	Discovery rule status is displayed - *Enabled*, *Disabled* or *Not supported*. By clicking the status you can change it - from Enabled to Disabled (and back); from Not supported to Disabled (and back).
Info	If everything is fine, no icon is displayed in this column. In case of errors, a square icon with the letter "i" is displayed. Hover over the icon to see a tooltip with the error description.

To configure a new low-level discovery rule, click the *Create discovery rule* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

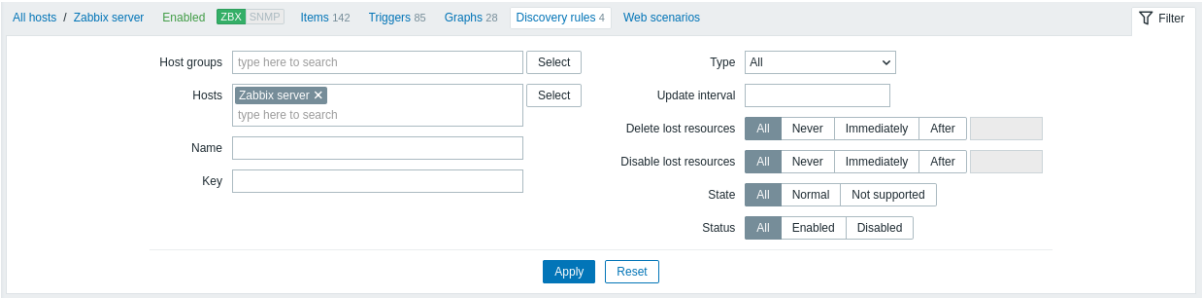
- *Enable* - change the low-level discovery rule status to *Enabled*.
- *Disable* - change the low-level discovery rule status to *Disabled*.
- *Execute now* - perform discovery based on the discovery rules immediately. See more details. Note that when performing discovery immediately, the configuration cache is not updated, thus the result will not reflect very recent changes to discovery rule configuration.
- *Delete* - delete the low-level discovery rules.

To use these options, mark the checkboxes before the respective discovery rules, then click the required button.

Using filter

You can use the filter to display only the discovery rules you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* link is available above the list of discovery rules. If you click it, a filter becomes available where you can filter discovery rules by host group, host, name, item key, item type, and other parameters.



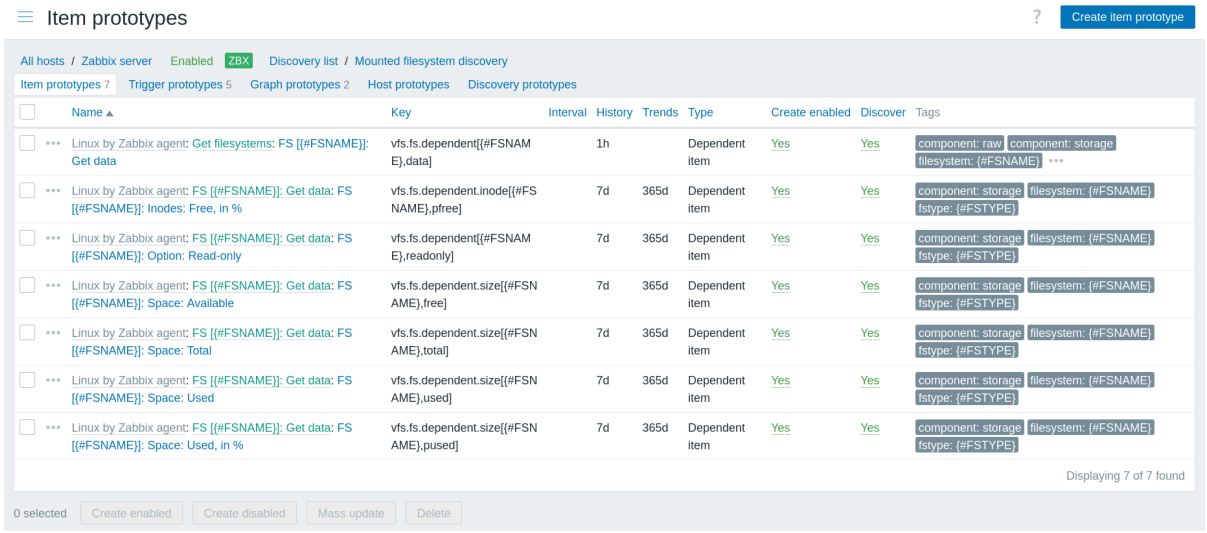
Parameter	Description
Host groups	Filter by one or more host groups. Specifying a parent host group implicitly selects all nested host groups.
Hosts	Filter by one or more hosts.
Name	Filter by discovery rule name.
Key	Filter by discovery item key.
Type	Filter by discovery item type.

Parameter	Description
Update interval	Filter by update interval. Not available for Zabbix trapper and dependent items.
Delete lost resources	Filter by *Delete lost resources* period.
Disable lost resources	Filter by *Disable lost resources* period.
SNMP OID	Filter by SNMP OID. Only available if *SNMP agent* is selected as type.
State	Filter by discovery rule state (All/Normal/Not supported).
Status	Filter by discovery rule status (All/Enabled/Disabled).

Item prototypes

Overview

In this section the item prototypes of a low-level discovery rule on the host are displayed. Item prototypes are the basis of real host items that are created during low-level discovery.



Displayed data:

Column	Description
Name	Name of the item prototype, displayed as a blue link. Clicking the name opens the item prototype configuration form. If the item prototype belongs to a template, the template name is displayed before the rule name as a gray link. Clicking the template link opens the item prototype list on the template level. If the item prototype is of the ”Dependent item” type, the master item name is displayed before the item prototype name in green.
Key	Key of the item prototype is displayed.
Interval	Frequency of the check is displayed.
History	How many days to keep item data history is displayed.
Trends	How many days to keep item trends history is displayed.
Type	Type of the item prototype is displayed (Zabbix agent, SNMP agent, simple check, etc).
Create enabled	Create the item based on this prototype as: **Yes** - enabled **No** - disabled. You can switch between ’Yes’ and ’No’ by clicking them.
Discover	Discover the item based on this prototype: **Yes** - discover **No** - do not discover. You can switch between ’Yes’ and ’No’ by clicking them.
Tags	Tags of the item prototype are displayed.

To configure a new item prototype, click the *Create item prototype* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

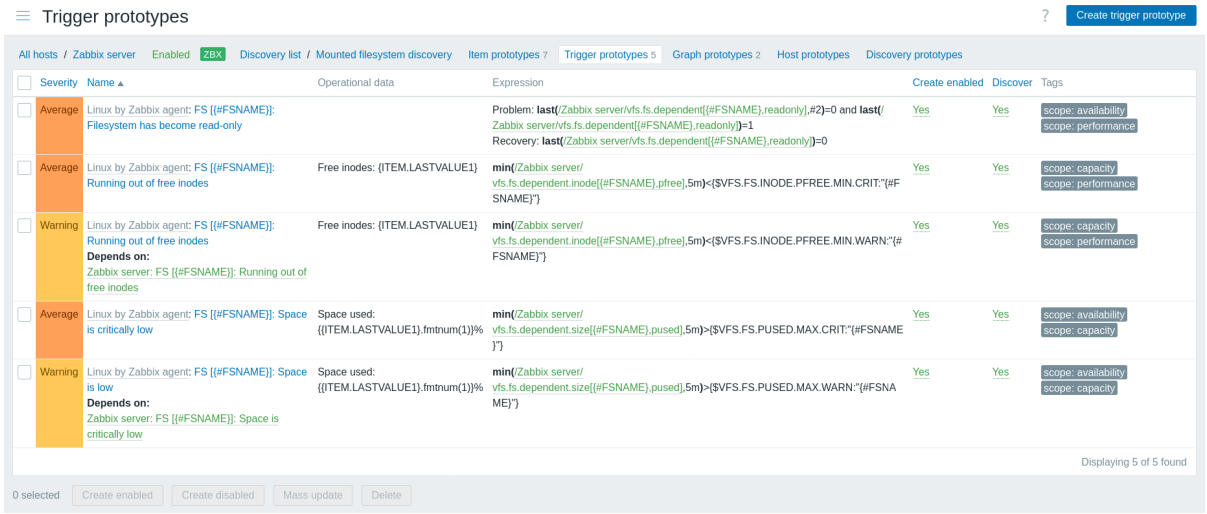
- *Create enabled* - create these items as *Enabled*
- *Create disabled* - create these items as *Disabled*
- *Mass update* - mass update these item prototypes
- *Delete* - delete these item prototypes

To use these options, mark the checkboxes before the respective item prototypes, then click the required button.

Trigger prototypes

Overview

In this section the trigger prototypes of a low-level discovery rule on the host are displayed. Trigger prototypes are the basis of real host triggers that are created during low-level discovery.



Displayed data:

Column	Description
Name	Name of the trigger prototype, displayed as a blue link. Clicking the name opens the trigger prototype configuration form. If the trigger prototype belongs to a linked template, the template name is displayed before the trigger name as a gray link. Clicking the template link opens the trigger prototype list on the linked template level.
Operational data	Format of the trigger's operational data, containing arbitrary strings and macros that are resolved dynamically in *Monitoring > Problems* and in the *Problems* and *Problems by severity* dashboard widgets. For details, see the description of the *Operational data* field in trigger configuration.
Create enabled	Create the trigger based on this prototype as: **Yes** - enabled **No** - disabled. You can switch between 'Yes' and 'No' by clicking them.
Discover	Discover the trigger based on this prototype: **Yes** - discover **No** - do not discover. You can switch between 'Yes' and 'No' by clicking them.
Tags	Tags of the trigger prototype are displayed.

To configure a new trigger prototype, click the *Create trigger prototype* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

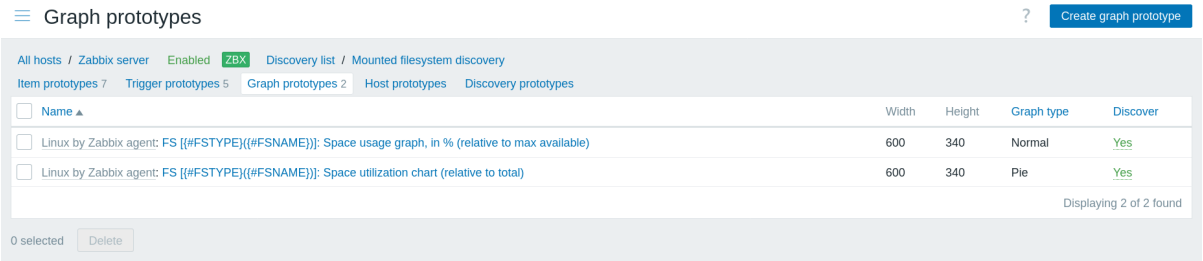
- *Create enabled* - create these triggers as *Enabled*
- *Create disabled* - create these triggers as *Disabled*
- *Mass update* - mass update these trigger prototypes
- *Delete* - delete these trigger prototypes

To use these options, mark the checkboxes before the respective trigger prototypes, then click the required button.

Graph prototypes

Overview

In this section the graph prototypes of a low-level discovery rule on the host are displayed. Graph prototypes are the basis of real host graphs that are created during low-level discovery.



Displayed data:

Column	Description
Name	Name of the graph prototype, displayed as a blue link. Clicking the name opens the graph prototype configuration form. If the graph prototype belongs to a linked template, the template name is displayed before the graph name as a gray link. Clicking the template link opens the graph prototype list on the linked template level.
Width	Width of the graph prototype is displayed.
Height	Height of the graph prototype is displayed.
Type	Type of the graph prototype is displayed - *Normal*, *Stacked*, *Pie* or *Exploded*.
Discover	Discover the graph based on this prototype: **Yes** - discover **No** - do not discover. You can switch between 'Yes' and 'No' by clicking them.

To configure a new graph prototype, click the *Create graph prototype* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

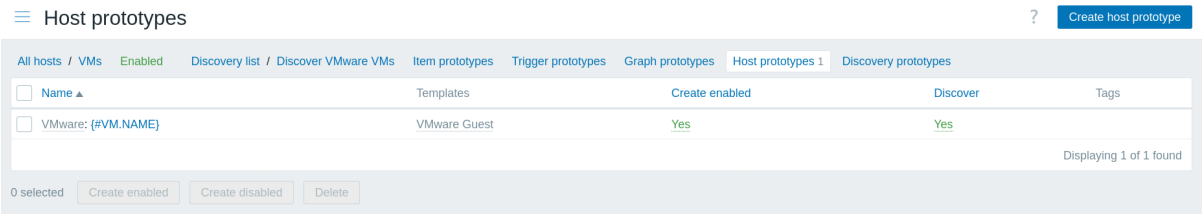
- *Delete* - delete these graph prototypes

To use these options, mark the checkboxes before the respective graph prototypes, then click the required button.

Host prototypes

Overview

In this section the host prototypes of a low-level discovery rule on the host are displayed. Host prototypes are the basis of real hosts that are created during low-level discovery.



Displayed data:

Column	Description
Name	Name of the host prototype, displayed as a blue link. Clicking the name opens the host prototype configuration form. If the host prototype belongs to a linked template, the template name is displayed before the host name as a gray link. Clicking the template link opens the host prototype list on the linked template level.
Templates	Templates of the host prototype are displayed.
Create enabled	Create the host based on this prototype as: **Yes** - enabled **No** - disabled. You can switch between 'Yes' and 'No' by clicking them.
Discover	Discover the host based on this prototype: **Yes** - discover **No** - do not discover. You can switch between 'Yes' and 'No' by clicking them.
Tags	Tags of the host prototype are displayed.

To configure a new host prototype, click the *Create host prototype* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

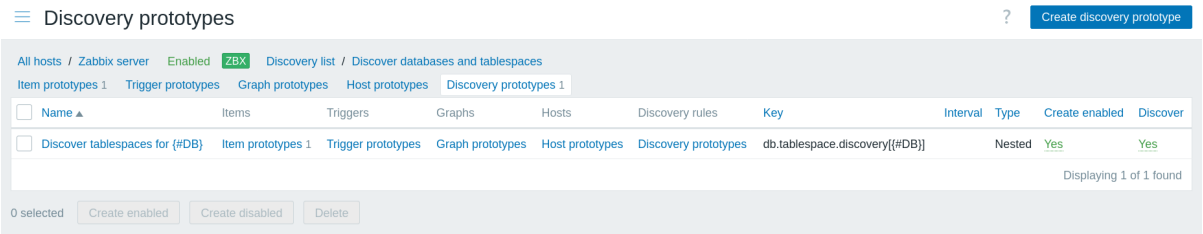
- *Create enabled* - create these hosts as *Enabled*
- *Create disabled* - create these hosts as *Disabled*
- *Delete* - delete these host prototypes

To use these options, mark the checkboxes before the respective host prototypes, then click the required button.

Discovery prototypes

Overview

In this section the discovery prototypes of a low-level discovery rule on the host are displayed. Discovery prototypes are nested discovery rules of the parent discovery rule.



Displayed data:

Column	Description
Name	Name of the discovery prototype, displayed as a blue link. Clicking the name opens the discovery prototype configuration form. If the discovery prototype belongs to a linked template, the template name is displayed before the host name as a gray link. Clicking the template link opens the discovery prototype list on the linked template level.
Items	A link to the list of item prototypes is displayed. The number of existing item prototypes is displayed in gray.
Triggers	A link to the list of trigger prototypes is displayed. The number of existing trigger prototypes is displayed in gray.
Graphs	A link to the list of graph prototypes is displayed. The number of existing graph prototypes is displayed in gray.
Hosts	A link to the list of host prototypes is displayed. The number of existing host prototypes is displayed in gray.
Discovery rules	A link to the list of discovery prototypes is displayed. The number of existing discovery prototypes is displayed in gray.
Key	The item key used for discovery is displayed.

Column	Description
Interval	The frequency of performing discovery is displayed. *Note* that discovery can also be performed immediately by pushing the *Execute now* button below the list.
Type	The item type used for discovery is displayed (Zabbix agent, SNMP agent, etc).
Create enabled	Discovery prototype creation status is displayed - create enabled (*Yes*) or create disabled (*No*). By clicking the status you can change it - from Yes to No (and back).
Discover	Discovery prototype discovery status is displayed - discover (*Yes*) or do not discover (*No*). By clicking the status you can change it - from Yes to No (and back).

To configure a new discovery prototype, click the *Create discovery prototype* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

- *Create enabled* - create these discovery prototypes as *Enabled*
- *Create disabled* - create these discovery prototypes as *Disabled*
- *Delete* - delete these discovery prototypes

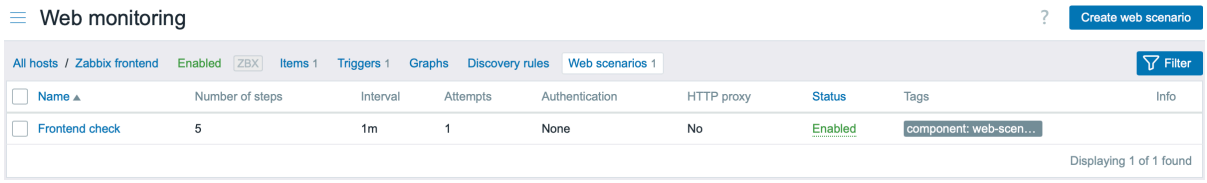
To use these options, mark the checkboxes before the respective discovery prototypes, then click the required button.

Web scenarios

Overview

The web scenario list for a host can be accessed from *Data collection* > *Hosts* by clicking *Web* for the respective host.

A list of existing web scenarios is displayed.



Displayed data:

Column	Description
Name	Name of the web scenario. Clicking the web scenario name opens the web scenario configuration form. If the host web scenario belongs to a template, the template name is displayed before the web scenario name as a gray link. Clicking the template link will open the web scenario list on the template level.
Number of steps	The number of steps the scenario contains.
Update interval	How often the scenario is performed.
Attempts	How many attempts for executing web scenario steps are performed.
Authentication	Authentication method is displayed - Basic, NTLM, Kerberos, Digest, or None.
HTTP proxy	Displays HTTP proxy or 'No' if not used.
Status	Web scenario status is displayed - *Enabled* or *Disabled*. By clicking the status you can change it.
Tags	Web scenario tags are displayed. Up to three tags (name:value pairs) can be displayed. If there are more tags, a "..." link is displayed that allows all tags to be viewed on mouseover.
Info	If everything is working correctly, no icon is displayed in this column. In case of errors, a square icon with the letter "i" is displayed. Hover over the icon to see a tooltip with the error description.

To configure a new web scenario, click the *Create web scenario* button at the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

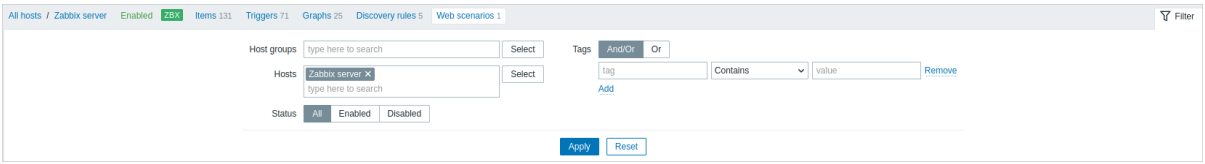
- *Enable* - change the scenario status to *Enabled*
- *Disable* - change the scenario status to *Disabled*
- *Clear history and trends* - clear history and trend data for the scenarios
- *Delete* - delete the web scenarios

To use these options, mark the checkboxes before the respective web scenarios, then click the required button.

Using filter

You can use the filter to display only the scenarios you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* link is available above the list of web scenarios. Clicking it will open a filter where you can specify the desired filtering criteria.



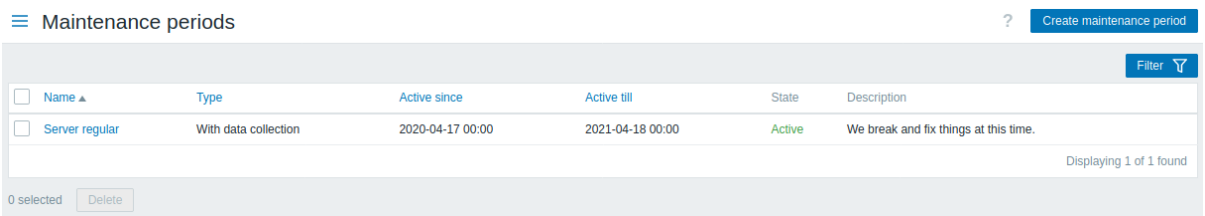
Parameter	Description
Host groups	Filter by one or more host groups. Specifying a parent host group implicitly selects all nested host groups.
Hosts	Filter by one or more hosts. If host groups are already selected above, host selection is limited to those groups.
Status	Filter by web scenario status.
Tags	Filter by web scenario tag name and value. It is possible to include as well as exclude specific tags and tag values. Several conditions can be set. Tag name matching is always case-sensitive. There are several operators available for each condition: **Exists** - include the specified tag names; **Equals** - include the specified tag names and values (case-sensitive); **Contains** - include the specified tag names where the tag values contain the entered string (substring match, case-insensitive); **Does not exist** - exclude the specified tag names; **Does not equal** - exclude the specified tag names and values (case-sensitive); **Does not contain** - exclude the specified tag names where the tag values contain the entered string (substring match, case-insensitive). There are two calculation types for conditions: **And/Or** - all conditions must be met, conditions having the same tag name will be grouped by the Or condition; **Or** - enough if one condition is met.

Maintenance

Overview

In the *Data collection* > *Maintenance* section users can configure and maintain maintenance periods for hosts.

A listing of existing maintenance periods with their details is displayed.



Displayed data:

Column	Description
Name	Name of the maintenance period. Clicking the maintenance period name opens the maintenance period configuration form.
Type	The type of maintenance is displayed: *With data collection* or *No data collection*
Active since	The date and time when executing maintenance periods becomes active. Note: This time does not activate a maintenance period; maintenance periods need to be set separately.
Active till	The date and time when executing maintenance periods stops being active.
State	The state of the maintenance period: **Upcoming** - will become active soon **Active** - is active **Expired** - is not active any more
Description	Description of the maintenance period is displayed.

To configure a new maintenance period, click the *Create maintenance period* button in the upper-right corner.

Mass editing options

A button below the list offers one mass-editing option:

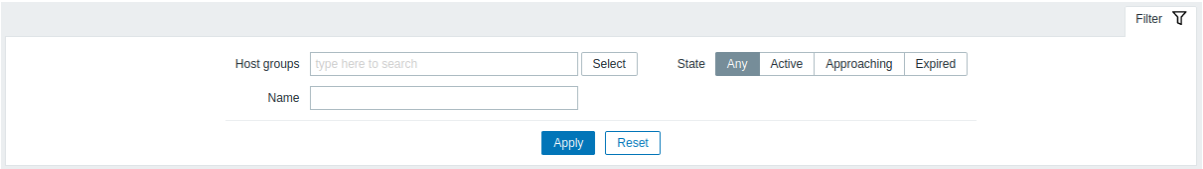
- *Delete* - delete the maintenance periods

To use this option, mark the checkboxes before the respective maintenance periods and click *Delete*.

Using filter

You can use the filter to display only the maintenance periods you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* link is available above the list of maintenance periods. If you click it, a filter becomes available where you can filter maintenance periods by host group, name and state.



Calculation of queues during maintenance

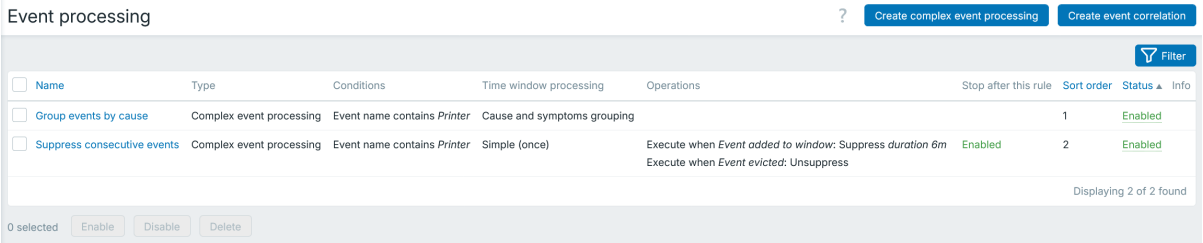
Attention:
The Zabbix proxy is not aware of maintenance periods; see Calculation of queues during maintenance for details.

Event processing

Overview

In the *Data collection* > *Event processing* section users can configure and maintain:

- Complex event processing rules (CEP rules)
- Global event correlation rules (deprecated)



Displayed data:

Column	Description
Name	Name of the rule. Clicking the name opens the rule configuration form - for CEP rule or global correlation rule respectively.
Type	Type of the rule - *Complex event processing* or *Correlation*.
Conditions	Rule conditions are displayed.
Time window processing	Time window type is displayed.
Operations	Rule operations are displayed.
Stop after this rule	Stop execution after this rule - *Enabled* or *Disabled*.
Sort order	Sort order of the CEP rule.
Status	Rule status is displayed - *Enabled* or *Disabled*. By clicking the status you can change it.
Info	If everything is working correctly, no icon is displayed in this column. In case of errors, a square icon with the letter "i" is displayed. Hover over the icon to see a tooltip with the error description.

To configure a new CEP rule, click the *Create complex event processing* button in the upper-right corner.

To configure a new correlation rule, click the *Create event correlation* button in the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

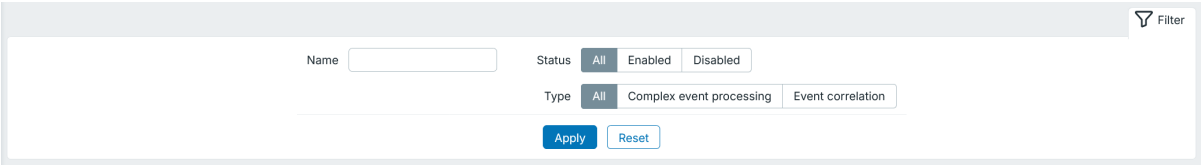
- *Enable* - change the rule status to *Enabled*
- *Disable* - change the rule status to *Disabled*
- *Delete* - delete the rules

To use these options, mark the checkboxes before the respective rules, then click the required button.

Using filter

You can use the filter to display only the rules you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* link is available above the list of CEP/correlation rules. If you click it, a filter becomes available where you can filter rules by name and status.

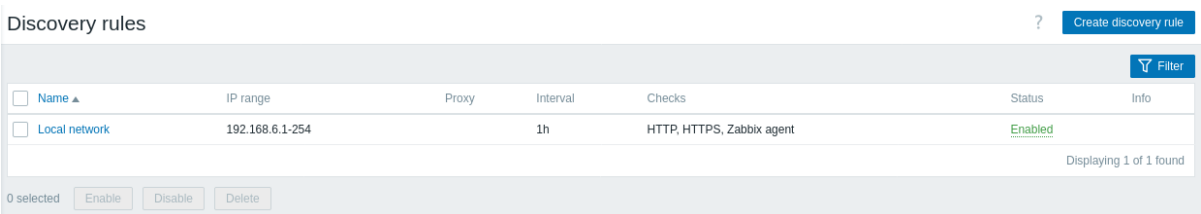


Discovery

Overview

In the *Data collection* > *Discovery* section users can configure and maintain discovery rules.

A listing of existing discovery rules with their details is displayed.



Displayed data:

Column	Description
Name	Name of the discovery rule. Clicking the discovery rule name opens the discovery rule configuration form.
IP range	The range of IP addresses to use for network scanning is displayed.
Proxy	The proxy name is displayed, if discovery is performed by the proxy.
Interval	The frequency of performing discovery displayed.

Column	Description
Checks	The types of checks used for discovery are displayed.
Status	The discovery rule status is displayed - *Enabled* or *Disabled*. By clicking the status you can change it.
Info	If everything is working correctly, nothing is displayed in this column. In case of errors, a red info icon with the letter "i" is displayed. Hover over the icon to see a tooltip with the error description.

To configure a new discovery rule, click the *Create discovery rule* button in the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

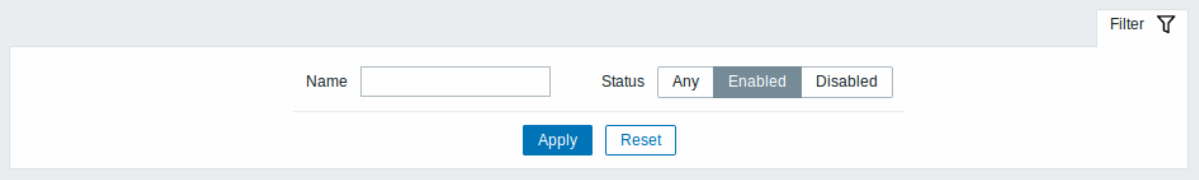
- *Enable* - change the discovery rule status to *Enabled*
- *Disable* - change the discovery rule status to *Disabled*
- *Delete* - delete the discovery rules

To use these options, mark the checkboxes before the respective discovery rules, then click the required button.

Using filter

You can use the filter to display only the discovery rules you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* link is available above the list of discovery rules. If you click it, a filter becomes available where you can filter discovery rules by name and status.



Alerts

Overview

This menu features sections that are related to configuring alerts in Zabbix.

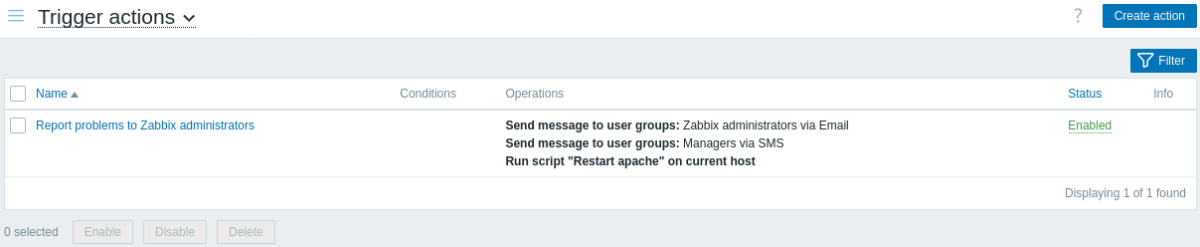
Actions

Overview

In the *Alerts* > *Actions* section users can configure and maintain actions.

The actions displayed are actions assigned to the selected event source (trigger, services, discovery, autoregistration, internal actions).

To switch to a different event source, click *Actions* in the *Alerts* menu section. It is also possible to switch between sources by using the title dropdown in the upper-left corner.



Displayed data:

Column	Description
Name	Name of the action. Clicking the action name opens the action configuration form.
Conditions	Action conditions are displayed.

Column	Description
Operations	Action operations are displayed. The operation list also displays the media type (email, SMS or script) used for notification as well as the name and surname (in parentheses after the username) of a notification recipient. Action operation can both be a notification or a remote command depending on the selected type of operation.
Status	Action status is displayed - *Enabled* or *Disabled*. By clicking the status you can change it. See the Escalations section for more details as to what happens if an action is disabled during an escalation in progress.
Info	If everything is working correctly, no icon is displayed in this column. In case of errors, such as missing action operations or conditions after upgrade/low-level discovery, a warning icon is displayed. Hover over the icon to see a tooltip with the error description.

To configure a new action, click the *Create action* button in the upper-right corner.

For users without Super admin rights actions are displayed according to the permission settings. That means in some cases a user without Super admin rights isn't able to view the complete action list because of certain permission restrictions. An action is displayed to the user without Super admin rights if the following conditions are fulfilled:

- The user has read-write access to host groups, hosts, templates, and triggers in action conditions
- The user has read-write access to host groups, hosts, and templates in action operations, recovery operations, and update operations
- The user has read access to user groups and users in action operations, recovery operations, and update operations

Mass editing options

Buttons below the list offer some mass-editing options:

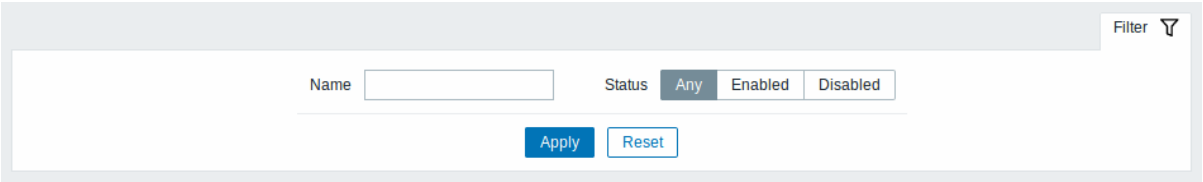
- *Enable* - change the action status to *Enabled*
- *Disable* - change the action status to *Disabled*
- *Delete* - delete the actions

To use these options, mark the checkboxes before the respective actions, then click the required button.

Using filter

You can use the filter to display only the actions you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* link is available above the list of actions. If you click on it, a filter becomes available where you can filter actions by name and status.



Media types

Overview

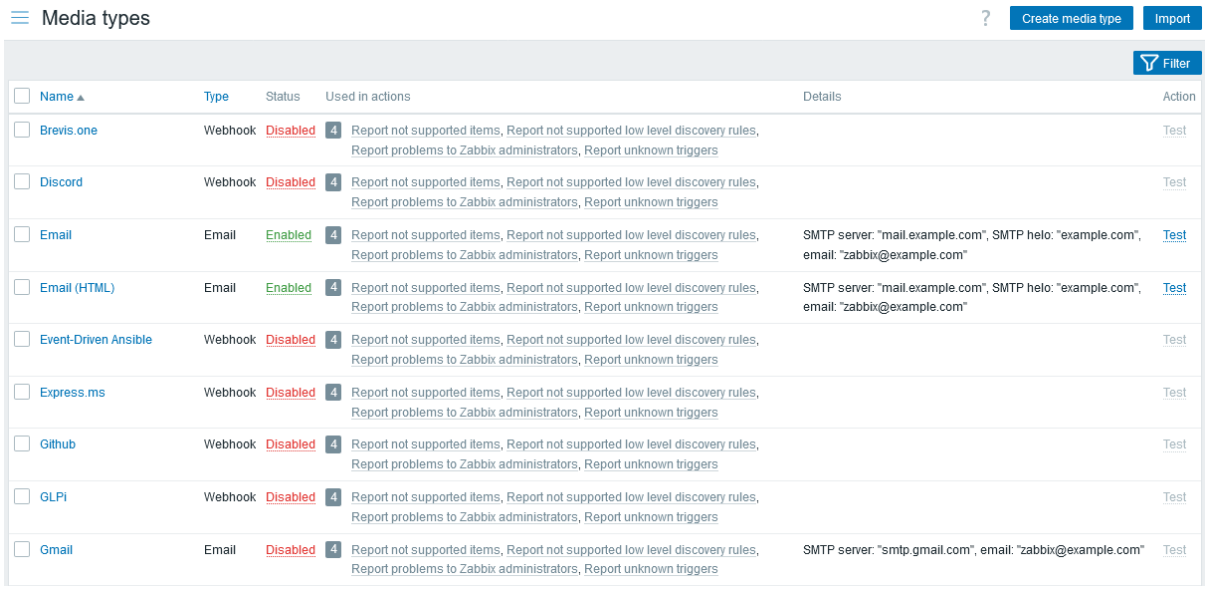
In the *Alerts* > *Media types* section users can configure and maintain media type information.

This menu is available to SuperAdmin user type users only.

Media type information contains general instructions for using a medium as delivery channel for notifications. Specific details, such as the individual email addresses to send a notification to are kept with individual users.

A listing of existing media types with their details is displayed.

Note:
If delivery methods (*Email*, *SMS*, *Webhook*, *Script*, or *Push*) are specified in the `$ZBX_FEATURE_FLAGS['media_type_denylist']` feature flag in the frontend configuration file, the corresponding pre-configured media types are disabled in this section.



Displayed data:

Column	Description
Name	Name of the media type. Clicking the name opens the media type configuration form.
Type	Type of the media (email, SMS, etc) is displayed.
Status	Media type status is displayed - *Enabled* or *Disabled*. By clicking the status you can change it.
Used in actions	Actions where the media type is used are displayed, preceded by the total number of these actions. Clicking the action name opens the action configuration form. If the user has no permissions to the action, the name is not clickable.
Details	Detailed information of the media type is displayed.
Actions	The following action is available: **Test** - click to open a testing form where you can enter media type parameters (e.g., a recipient address with test subject and body) and send a test message to verify that the configured media type works. See also: Media type testing.

To configure a new media type, click the *Create media type* button in the upper-right corner.

To import a media type, click the *Import* button in the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

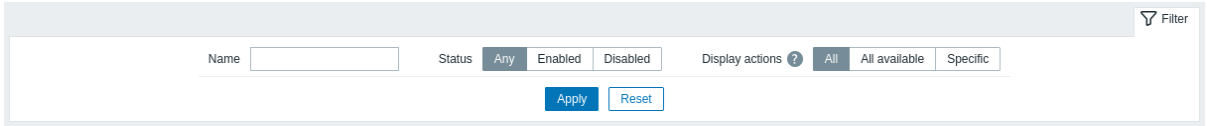
- *Enable* - change the media type status to *Enabled*
- *Disable* - change the media type status to *Disabled*
- *Export* - export the media types to a YAML, XML or JSON file
- *Delete* - delete the media types

To use these options, mark the checkboxes before the respective media types, then click the required button.

Using filter

You can use the filter to display only the media types you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* link is available above the list of media types. If you click it, a filter becomes available where you can filter media types by name and status. Additionally, you can use the filter to display actions in the *Used in actions* column based on the scope of their media type usage (defined by the *Send to media type* parameter in action operation details).



Scripts

Overview

In the *Alerts* > *Scripts* section user-defined global scripts can be configured and maintained.

This menu is available to SuperAdmin user type users only.

Global scripts, depending on the configured scope and also user permissions, are available for execution:

- From the host menu in various frontend locations (*Dashboard*, *Problems*, *Latest data*, *Maps*, etc.).
- From the event menu.
- Can be run as an action operation.

The scripts are executed by Zabbix agent, Zabbix server (proxy) or Zabbix server only. See also Command execution.

Remote commands are disabled by default on both the Zabbix agent and Zabbix proxy. They can be enabled as follows:

- For remote commands executed by Zabbix agent, add an `AllowKey=system.run[<command>,*]` or `AllowKeyRegexp=^system\.r` parameter for each allowed command in the agent configuration, where `*` matches either execution mode (`wait` or `nowait`).
- For remote commands executed by Zabbix proxy, set the `EnableRemoteCommands` parameter to `1` in the proxy configuration. Note that remote commands do not need to be enabled on Zabbix proxy when they are executed by Zabbix agent on a host monitored by the proxy.

Attention:
Do not use `AllowKey=system.run[*]` or `AllowKeyRegexp=^system\.run\[.*\]$` unless necessary, as it allows the execution of any command on the host. For security reasons, it is recommended to allow only exact commands, while `system.run[*]` is denied by default (`DenyKey=system.run[*]`).

Global script execution on Zabbix server can be disabled by setting `EnableGlobalScripts=0` in server configuration. For new installations starting with Zabbix 7.0, global script execution on Zabbix server is disabled by default.

A listing of existing scripts with their details is displayed.



Displayed data:

Column	Description
Name	Name of the script. Clicking the script name opens the script configuration form.
Scope	Scope of the script - action operation, manual host action or manual event action. This setting determines where the script is available - in remote commands of action operations, from the host menu, or from the event menu respectively.
Used in actions	All actions where the script is used are displayed, preceded by the total number of these actions. Clicking the action name opens the action configuration form. If the user has no permissions to the action, the name is not clickable.
Type	Script type is displayed - *URL*, *Webhook*, *Script*, *SSH*, *Telnet* or *IPMI* command.
Execute on	It is displayed whether the script will be executed on Zabbix agent, Zabbix proxy or server, or Zabbix server only.
Commands	All commands to be executed within the script are displayed. Nothing is displayed here for webhooks.
User group	The user group that the script is available to is displayed (or *All* for all user groups).
Host group	The host group that the script is available for is displayed (or *All* for all host groups).

Column	Description
Host access	The permission level for the host group is displayed - *Read* or *Write*. Only users with the required permission level will have access to executing the script.

To configure a new script, click the *Create script* button in the upper-right corner.

Mass editing options

A button below the list offers one mass-editing option:

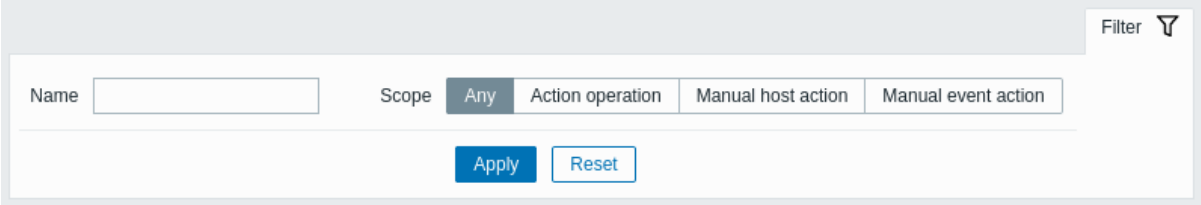
- *Delete* - delete the scripts

To use this option, mark the checkboxes before the respective scripts and click *Delete*.

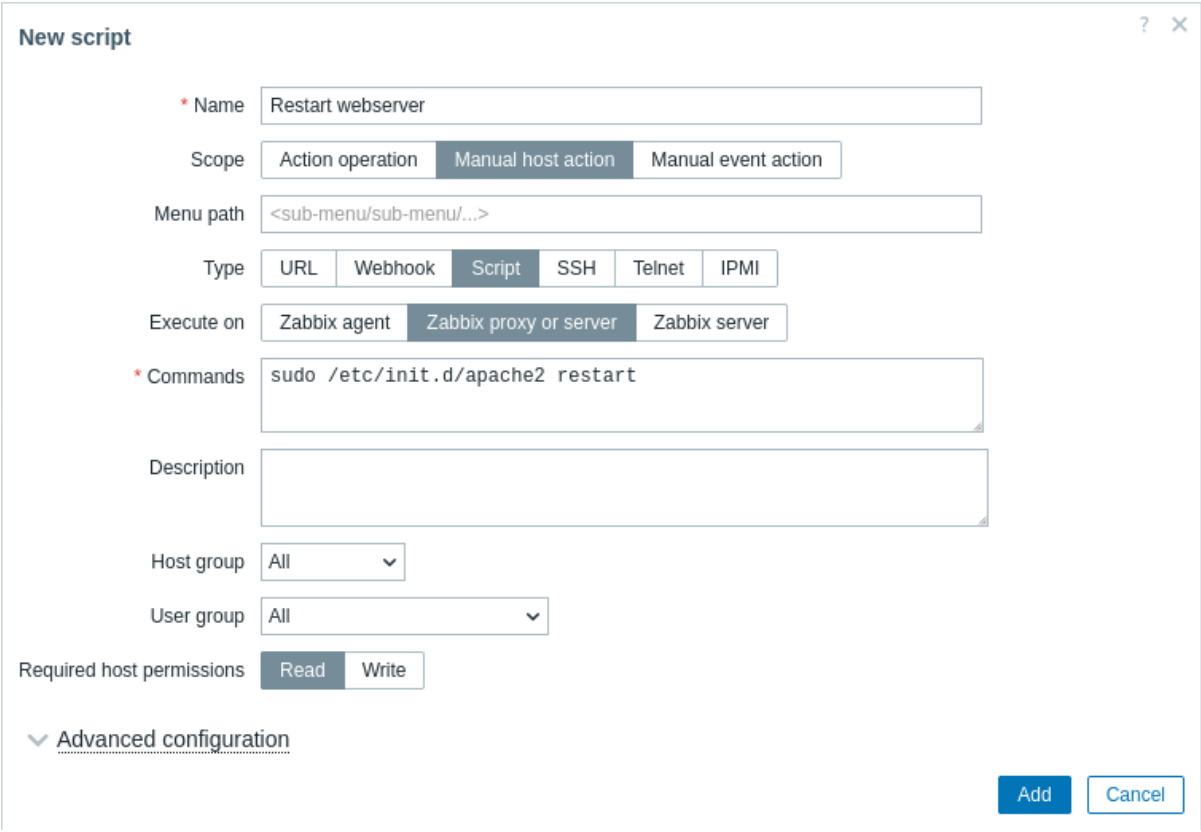
Using filter

You can use the filter to display only the scripts you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* link is available above the list of scripts. If you click it, a filter becomes available where you can filter scripts by name and scope.



Configuring a global script



Note:
It is recommended to use macro functions instead of plain macros in global scripts to enhance security, as macros are not automatically escaped.

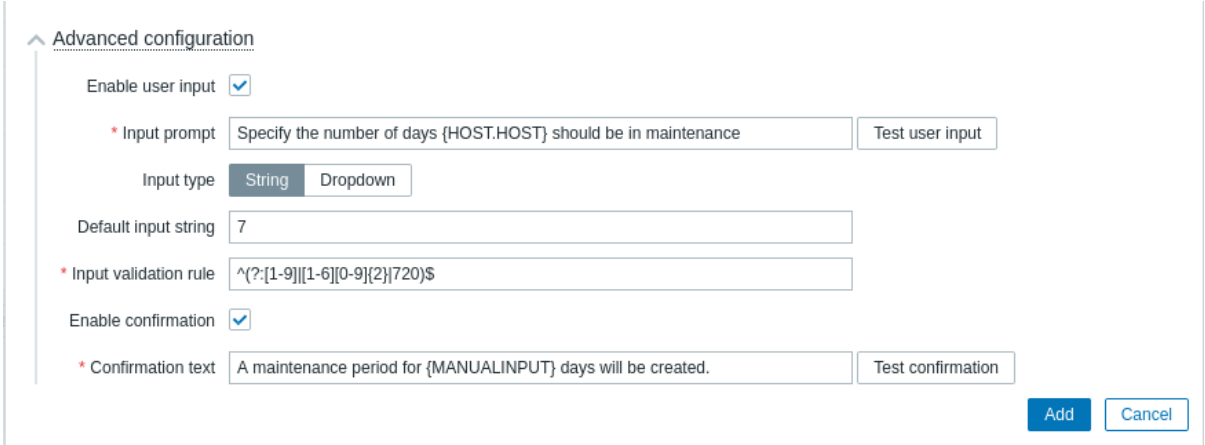
Script attributes:

Parameter	Description
Name	Unique name of the script. E.g., `Clear /tmp filesystem`
Scope	Scope of the script - action operation, manual host action or manual event action. This setting determines where the script can be used - in remote commands of action operations, from the host menu or from the event menu respectively. Setting the scope to 'Action operation' makes the script available for all users with access to *Alerts* > *Actions*. If a script is actually used in an action, its scope cannot be changed away from 'action operation'. **Macro support** The scope affects the range of available macros. For example, user-related macros (`{USER.*}`) are supported in scripts to allow passing information about the user that launched the script. However, they are not supported if the script scope is action operation, as action operations are executed automatically. A `{MANUALINPUT}` macro allows you to specify manual input at script execution time. It is supported for manual host action and manual event action scripts. To find out which other macros are supported, do a search for 'Trigger-based notifications and commands/Trigger-based commands', 'Manual host action scripts' and 'Manual event action scripts' in the supported macro table. Note that if a macro may resolve to a value with spaces (for example, host name), don't forget to quote as needed.
Menu path	The desired menu path to the script. For example, `Default` or `Default/`, will display the script in the respective directory. Menus can be nested, e.g., `Main menu/Sub menu1/Sub menu2`. When accessing scripts through the host/event menu in monitoring sections, they will be organized according to the given directories. This field is displayed only if 'Manual host action' or 'Manual event action' is selected as *Scope*.
Type	Click the respective button to select script type: **URL**, **Webhook**, **Script**, **SSH**, **Telnet** or **IPMI** command. The type **URL** is available only when 'Manual host action' or 'Manual event action' is selected as *Scope*.
Script type: **URL**	
URL	Specify the URL for quick access from the host menu or event menu. Macros and custom user macros are supported. Macro support depends on the scope of the script (see *Scope* above). Use the {MANUALINPUT} macro in this field to be able to specify manual input at script execution time, for example: `http://{MANUALINPUT}/zabbix/zabbix.php?action=dashboard.view` Macro values must not be URL-encoded.
Open in new window	Determines whether the URL should be opened in a new or the same browser tab.
Script type: **Webhook**	
Parameters	Specify the webhook variables as attribute-value pairs. See also: Webhook media configuration. Macros and custom user macros are supported in parameter values. Macro support depends on the scope of the script (see *Scope* above).
Script	Enter the JavaScript code in the modal editor that opens when clicking in the parameter field or on the pencil icon next to it. Macro support depends on the scope of the script (see *Scope* above). See also: Webhook media configuration, Additional Javascript objects.
Timeout	JavaScript execution timeout (1-60s, default 30s). Time suffixes are supported (e.g., `30s`, `1m`).
Script type: **Script**	
Execute on	Click the respective button to execute the shell script on: **Zabbix agent** - the script will be executed by Zabbix agent (if the system.run item is allowed) on the host **Zabbix proxy or server** - the script will be executed by Zabbix proxy or server - depending on whether the host is monitored by proxy or server. It will be executed on the proxy if enabled by EnableRemoteCommands. It will be executed on the server if global scripts are enabled by the EnableGlobalScripts server parameter. **Zabbix server** - the script will be executed by Zabbix server only. This option will not be available if global scripts are disabled by the EnableGlobalScripts server parameter.

Parameter	Description
Commands	Enter full path to the commands to be executed within the script. Macro support depends on the scope of the script (see *Scope* above). Custom user macros are supported.
Script type: **SSH**	
Authentication method	Select authentication method - password or public key.
Username	Enter the username.
Password	Enter the password. This field is available if 'Password' is selected as the authentication method.
Public key file	Enter the path to the public key file. This field is available if 'Public key' is selected as the authentication method.
Private key file	Enter the path to the private key file. This field is available if 'Public key' is selected as the authentication method.
Passphrase	Enter the passphrase. This field is available if 'Public key' is selected as the authentication method.
Port	Enter the remote SSH service port on the target host that Zabbix will connect to.
Commands	Enter the commands. Macro support depends on the scope of the script (see *Scope* above). Custom user macros are supported.
Script type: **Telnet**	
Username	Enter the username.
Password	Enter the password.
Port	Enter the remote Telnet service port on the target host that Zabbix will connect to.
Commands	Enter the commands. Macro support depends on the scope of the script (see *Scope* above). Custom user macros are supported.
Script type: **IPMI**	
Command	Enter the IPMI command. Macro support depends on the scope of the script (see *Scope* above). Custom user macros are supported.
Description	Enter a description for the script.
Host group	Select the host group that the script will be available for (or *All* for all host groups).
User group	Select the user group that the script will be available to (or *All* for all user groups). This field is displayed only if 'Manual host action' or 'Manual event action' is selected as *Scope*.
Required host permissions	Select the permission level for the host group - *Read* or *Write*. Only users with the required permission level will have access to executing the script. This field is displayed only if 'Manual host action' or 'Manual event action' is selected as *Scope*.
Advanced configuration	Click the *Advanced configuration* header to display advanced configuration options. This field is displayed only if 'Manual host action' or 'Manual event action' is selected as *Scope*.

Advanced configuration

Advanced configuration options are available in a collapsible *Advanced configuration* section:



Parameter	Description
Enable user input	Mark the checkbox to enable manual user input before executing the script. Manual user input will replace the {MANUALINPUT} macro value in the script. See also: Manual user input.
Input prompt	Enter custom text prompting for custom user input. This text will be displayed above the input field in the *Manual input* popup. To see a preview of the *Manual input* popup, click *Test user input*. The preview also allows you to test if the input string complies with the input validation rule (see parameters below). Macro and user macro support depends on the scope of the script (see *Scope* in general script configuration parameters).
Input type	Select the manual input type: **String** - single string; **Dropdown** - value is selected from multiple dropdown options.
Dropdown options	Enter unique values for the user input dropdown in a comma-delimited list. To include an empty option in the dropdown, add an extra comma at the beginning, middle, or end of the list. This field is displayed only if 'Dropdown' is selected as *Input type*.
Default input string	Enter the default string for user input (or none). This field will be validated against the regular expression provided in the *Input validation rule* field. The value entered here will be displayed by default in the *Manual input* popup. This field is displayed only if 'String' is selected as *Input type*.
Input validation rule	Enter a regular expression to validate the user input string. Global regular expressions are supported. This field is displayed only if 'String' is selected as *Input type*.
Enable confirmation	Mark the checkbox to display a confirmation message before executing the script. This feature might be especially useful with potentially dangerous operations (like a reboot script) or ones that might take a long time.
Confirmation text	Enter custom confirmation text for the confirmation popup enabled with the checkbox above (for example, *Remote system will be rebooted. Are you sure?*). To see how the text will look like, click *Test confirmation* next to the field. Macros and custom user macros are supported. *Note:* the macros will not be expanded when testing the confirmation message.

If both manual user input and a confirmation message are configured, they will be displayed in consecutive popup windows.

Manual user input

Manual user input allows a custom parameter to be supplied on each execution of the script. This saves the necessity to create multiple similar user scripts with only a single parameter difference.

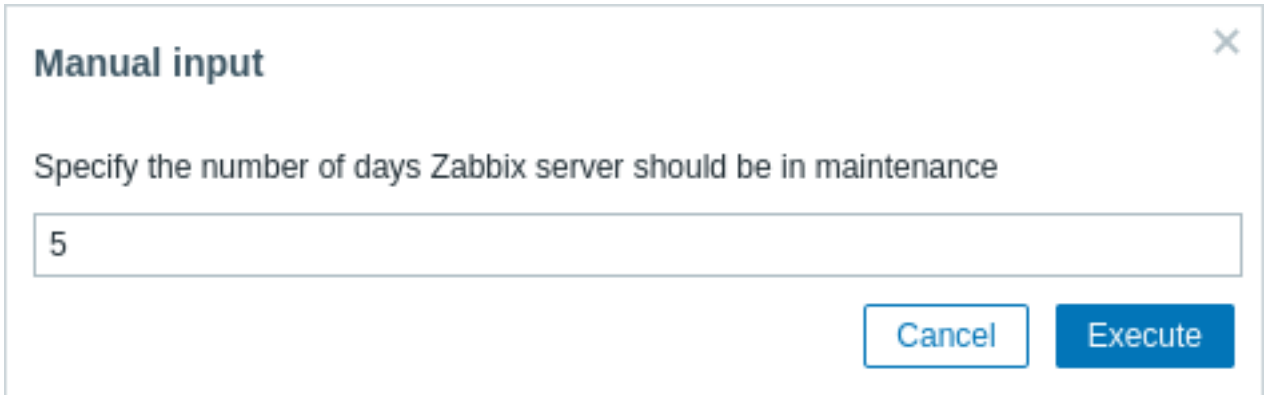
For example, you may want to supply a different integer or a different URL address to the script during execution.

To enable manual user input:

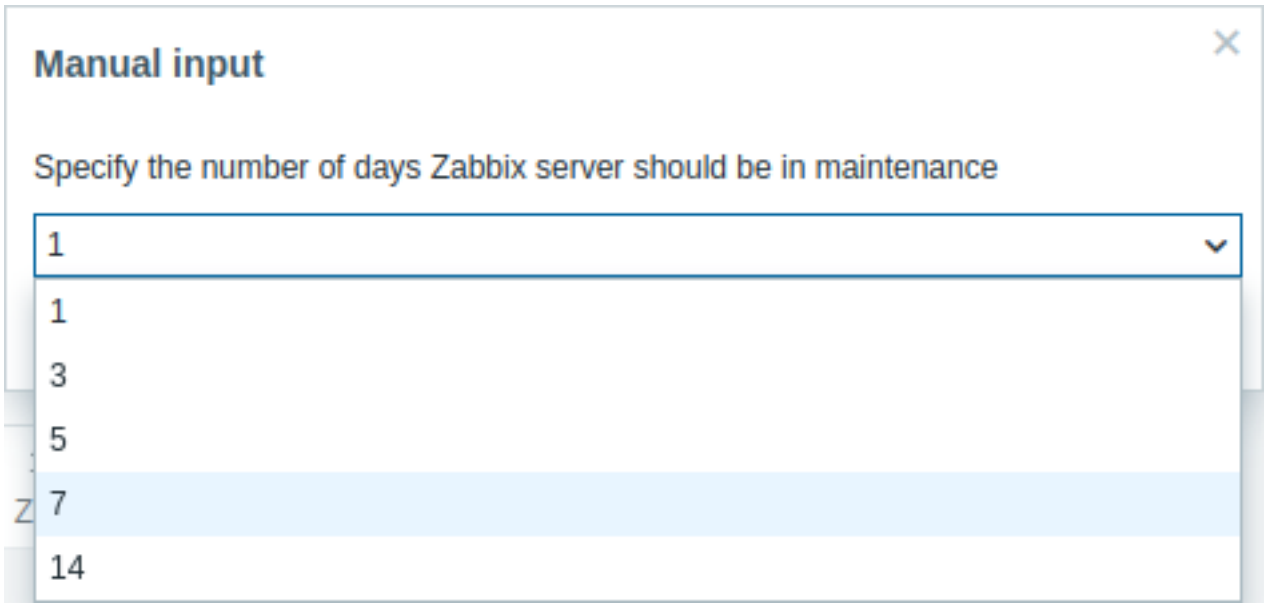
- use the {MANUALINPUT} macro in the script (commands, script, script parameter) where required; or in the URL field of URL scripts;
- in advanced script configuration, enable manual user input and configure input options.

With user input enabled, before script execution, a *Manual input* popup will appear to the user asking to supply a custom value. The supplied value will replace {MANUALINPUT} in the script.

Depending on the configuration, the user will be asked to enter a string value:



Or select the value from a dropdown of pre-determined options:



Manual user input is available only for scripts where the scope is 'Manual host action' or 'Manual event action'.

Script execution and result

Scripts run by Zabbix server are executed in the order described in the command execution page.

The script result is displayed in a pop-up window that appears after the script is run. The return value of the script is a standard output:

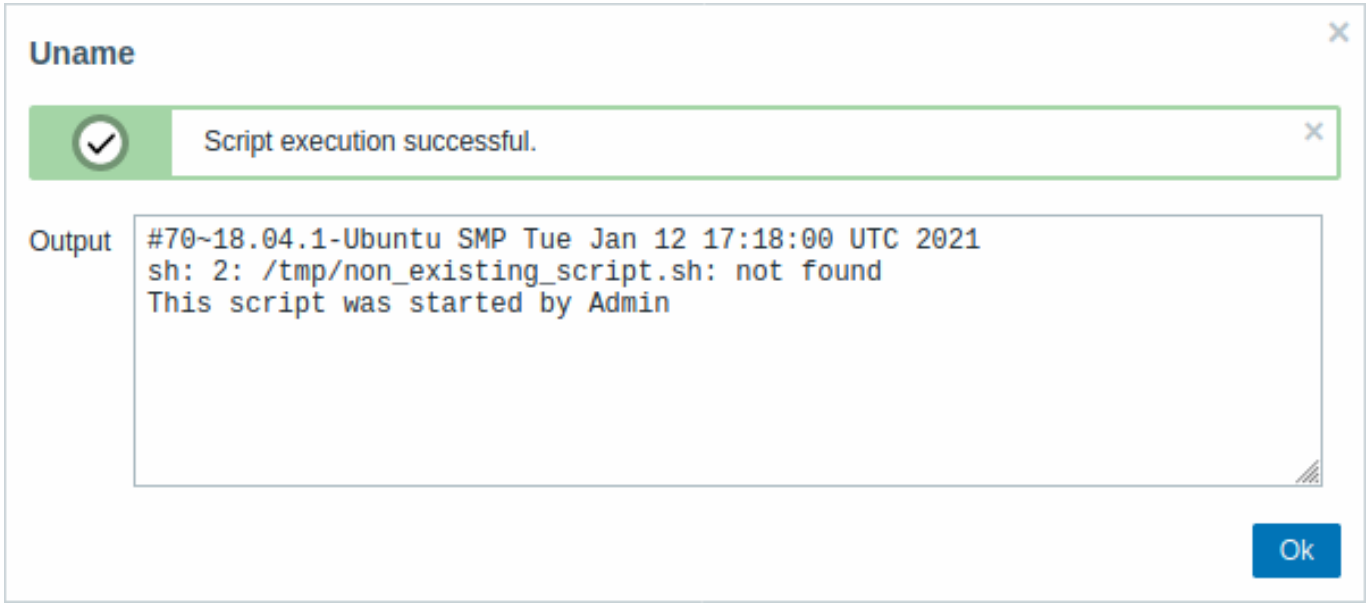
- If the script finishes successfully (exit code 0), the return value is limited to 16MB (including trailing whitespace that is truncated).
- If the script exits with an error (non-zero exit code), the return value is a standard error limited to 2KB.

Zabbix does not store extended script output by default. To preserve full output details, you can implement logging within the script itself (e.g., redirect output to a local log file).

Note that for scripts executed either on Zabbix server or Zabbix proxy, database limits also apply.

Below is an example of a script and the result window:

```
uname -v
/tmp/non_existing_script.sh
echo "This script was started by {USER.USERNAME}"
```



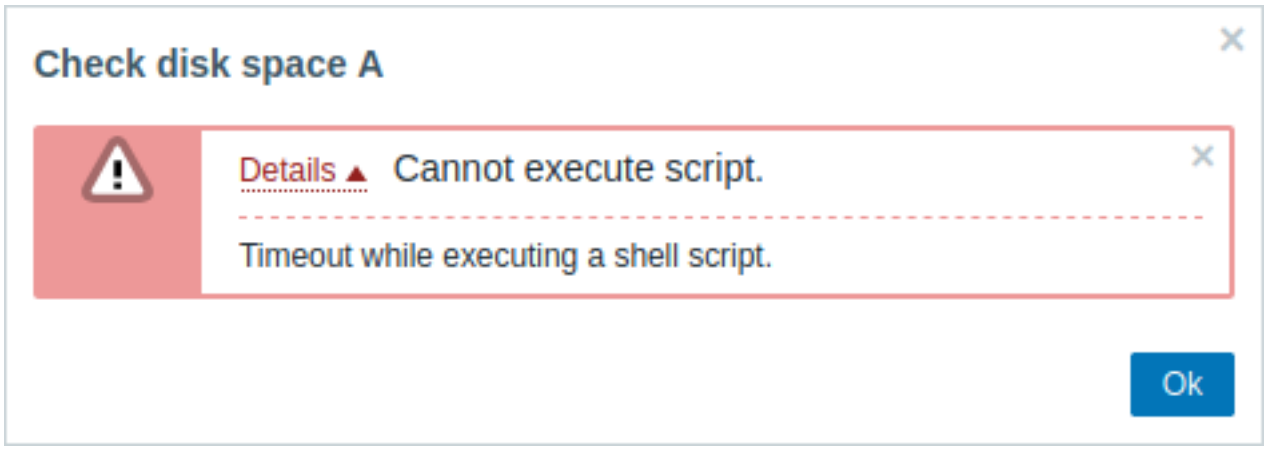
The script result does not display the script itself.

Script timeout on Zabbix agent

You may encounter a situation when a timeout occurs while executing a script.

See an example of a script running on Zabbix agent and the result window below:

```
sleep 5
df -h
```



The error message, in this case, is the following:

```
Timeout while executing a shell script.
```

To avoid such situations, it is advised to optimize the script itself (in the example above, "5") instead of adjusting the `Timeout` parameter in Zabbix agent configuration and Zabbix server configuration.

In case the `Timeout` parameter has been changed in Zabbix agent configuration, the following error message will appear:

```
Get value from agent failed: ZBX_TCP_READ() timed out.
```

It means that the modification has been made in Zabbix agent configuration, but it is required to modify the `Timeout` parameter in Zabbix server configuration as well.

Script timeout on active agents

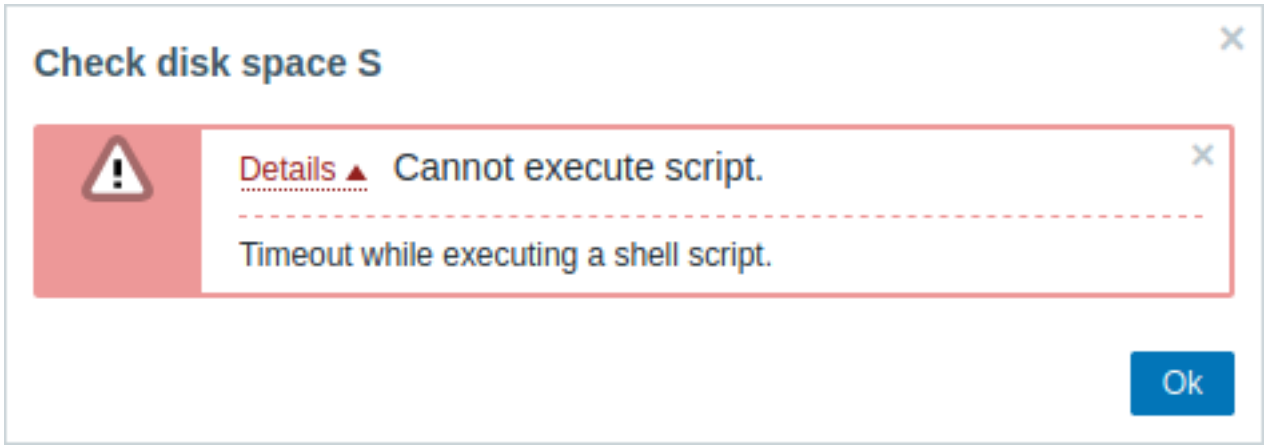
It is possible to execute scripts on agents operating in active mode only. Once the script execution is triggered by action operation or manual script execution, the command is included in active check configuration and executed once the active agent receives it.

Manual scripts are sent to the active agent along with the server/proxy timeout for script execution. Please increase the default server/proxy timeout for script execution. The timeout must be higher than the active check refresh frequency; otherwise, the timeout is exceeded before the active agent receives the script and is able to return the result.

Script timeout on Zabbix server/proxy

See an example of a script running on Zabbix server and the result window below:

```
sleep 11
df -h
```



It is also advised to optimize the script itself (instead of adjusting `TrapperTimeout` parameter to a corresponding value (in our case, > 11) by modifying the Zabbix server configuration).

Users

Overview

This menu features sections that are related to configuring users in Zabbix. This menu is available to SuperAdmin user type users only.

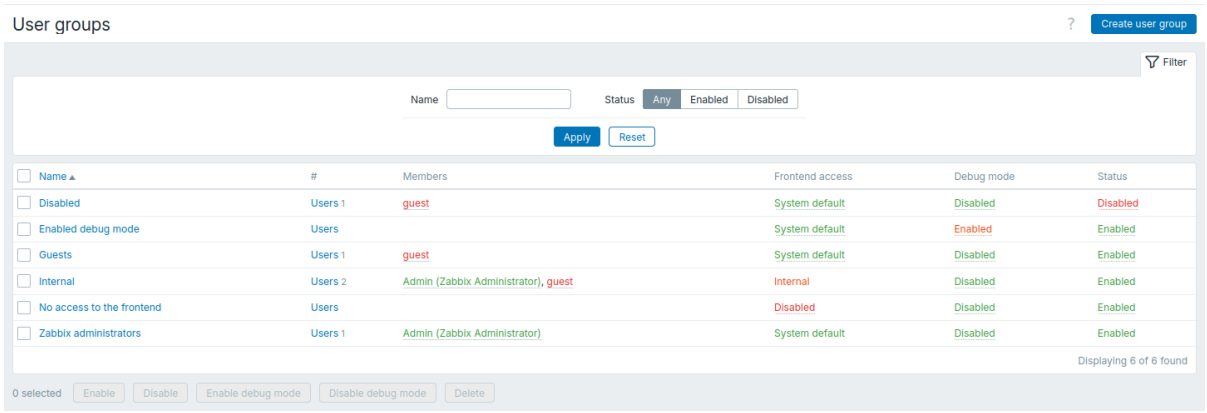
User groups

Overview

In the *Users* > *User groups* section user groups of the system are maintained.

User groups

A listing of existing user groups with their details is displayed.



Displayed data:

Column	Description
Name	Name of the user group. Clicking the user group name opens the user group configuration form.
#	The number of users in the group. Clicking *Users* will display the respective users filtered out in the user list.
Members	Usernames of individual users in the user group (with name and surname in parentheses). Clicking the username will open the user configuration form. Users from disabled groups are displayed in red.

Column	Description
Frontend access	Frontend access level is displayed: **System default** - users are authenticated by Zabbix, LDAP or HTTP (depending on the authentication method set globally); **Internal** - users are authenticated by Zabbix; ignored if HTTP authentication is the global default; **LDAP** - users are authenticated by LDAP; ignored if HTTP authentication is the global default; **Disabled** - access to Zabbix frontend is forbidden for this group. By clicking the current level, you can change it.
Debug mode	Debug mode status is displayed - *Enabled* or *Disabled*. By clicking the status you can change it.
Status	User group status is displayed - *Enabled* or *Disabled*. By clicking the status you can change it.

To configure a new user group, click the *Create user group* button in the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

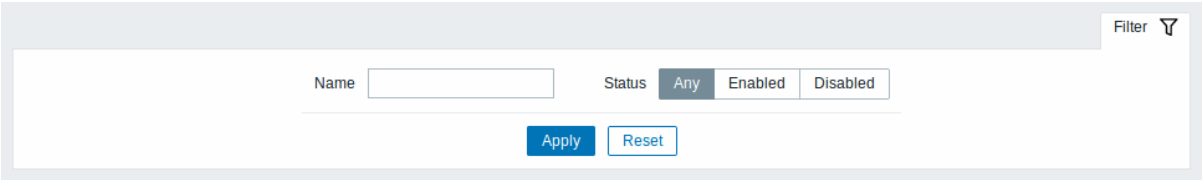
- *Enable* - change the user group status to *Enabled*
- *Disable* - change the user group status to *Disabled*
- *Enable debug mode* - enable debug mode for the user groups
- *Disable debug mode* - disable debug mode for the user groups
- *Delete* - delete the user groups

To use these options, mark the checkboxes before the respective user groups, then click the required button.

Using filter

You can use the filter to display only the user groups you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* link is available above the list of user groups. If you click it, a filter becomes available where you can filter user groups by name and status.



User roles

Overview

In the *Users* > *User roles* section you may create user roles.

User roles allow you to create fine-grained permissions based on the initially selected user type (*User*, *Admin*, *Super admin*).

Upon selecting a user type, all available permissions for this user type are granted (checked by default).

Permissions can only be revoked from the subset that is available for the user type; they cannot be extended beyond what is available for the user type.

Checkboxes for unavailable permissions are grayed out; users will not be able to access the element even by entering a direct URL to this element into the browser.

Note:
Restricting access to some UI elements only prevents opening that page—it does not remove the ability to access underlying data on other parts of the interface.

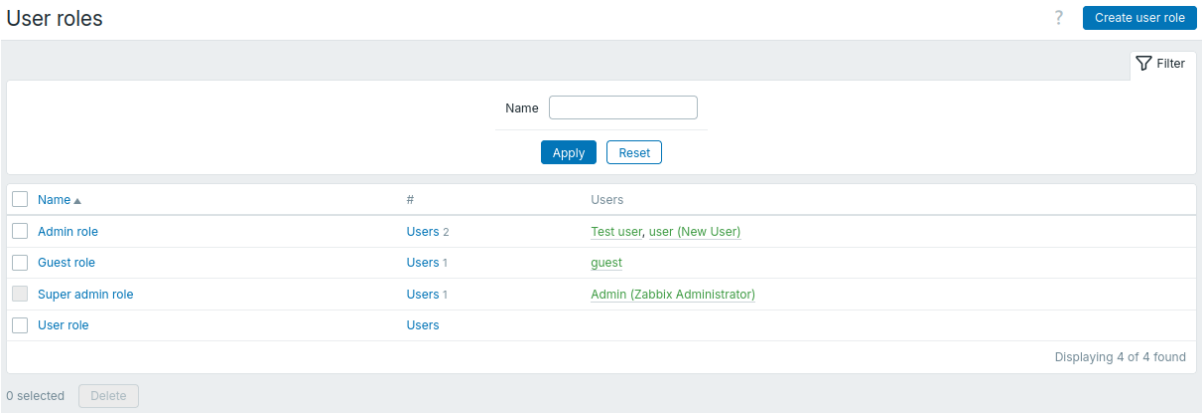
User roles can be assigned to system users. Each user may have only one role assigned.

Default user roles

By default, Zabbix is configured with four user roles, which have a pre-defined set of permissions:

- Guest role
- User role
- Admin role

- Super admin role



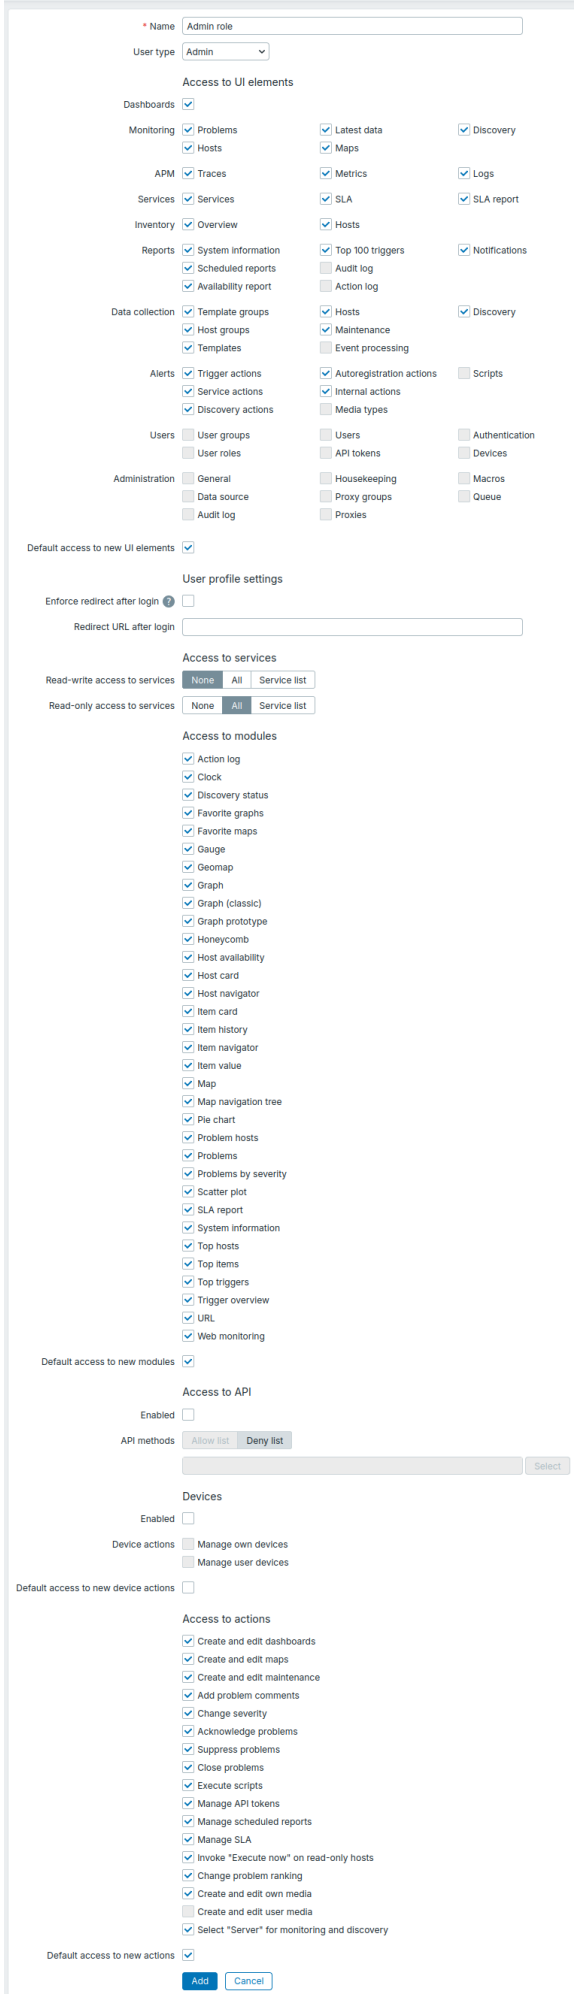
These are based on the main user types in Zabbix. The list of all users assigned the respective role is displayed. The users included in disabled groups are stated in red. The *Guest role* is a user-type role with the only permissions to view some frontend sections.

Note:
The default *Super admin role* cannot be modified or deleted, because at least one Super admin user with unlimited privileges must exist in Zabbix. Users of type *Super admin* can modify settings of their own role, but not the user type.

Configuration

To create a new role, click the *Create user role* button at the upper-right corner. To update an existing role, click the role name to open the configuration form.

User roles


* Name
Admin role
User type
Admin
Access to UI elements
Dashboards
Monitoring
Problems
Latest data
Discovery
Hosts
Maps
APM
Traces
Metrics
Logs
Services
Services
SLA
SLA report
Inventory
Overview
Hosts
Reports
System information
Top 100 triggers
Notifications
Scheduled reports
Audit log
Availability report
Action log
Data collection
Template groups
Hosts
Discovery
Host groups
Maintenance
Templates
Event processing
Alerts
Trigger actions
Autoregistration actions
Scripts
Service actions
Internal actions
Discovery actions
Media types
Users
User groups
Users
Authentication
User roles
API tokens
Devices
Administration
General
Housekeeping
Macros
Data source
Proxy groups
Queue
Audit log
Proxies
Default access to new UI elements
User profile settings
Enforce redirect after login
Redirect URL after login
Access to services
Read-write access to services
None
All
Service list
Read-only access to services
None
All
Service list
Access to modules
Action log
Clock
Discovery status
Favorite graphs
Favorite maps
Gauge
Geomap
Graph
Graph (classic)
Graph prototype
Honeycomb
Host availability
Host card
Host navigator
Item card
Item history
Item navigator
Item value
Map
Map navigation tree
Pie chart
Problem hosts
Problems
Problems by severity
Scatter plot
SLA report
System information
Top hosts
Top items
Top triggers
Trigger overview
URL
Web monitoring
Default access to new modules
Access to API
Enabled
API methods
Allow list
Deny list
Select
Devices
Enabled
Device actions
Manage own devices
Manage user devices
Default access to new device actions
Access to actions
Create and edit dashboards
Create and edit maps
Create and edit maintenance
Add problem comments
Change severity
Acknowledge problems
Suppress problems
Close problems
Execute scripts
Manage API tokens
Manage scheduled reports
Manage SLA
Invoke "Execute now" on read-only hosts
Change problem ranking
Create and edit own media
Create and edit user media
Select "Server" for monitoring and discovery
Default access to new actions
Add
Cancel

Available permissions are displayed. To revoke a certain permission, unmark its checkbox.

Available permissions along with the defaults for each pre-configured user role in Zabbix are described below.

Default permissions

Access to UI elements

The default access to menu sections depends on the user type. See the Permissions page for details.

Access to other options

Some of the sections described in the table below may be hidden:

- *Access to modules* is hidden if `$ZBX_FEATURE_FLAGS['modules_config_enabled']` is set to `false` in the frontend configuration file.
- *Devices* is hidden if `EnableMobileDevices` is disabled.

Parameter	Description	Default user roles			
		Super admin role	*Admin role*	*User role*	*Guest role*
Default access to new UI elements	This option specifies how new menu sections will be accessible after a Zabbix upgrade. Existing menu sections of modules remain unaffected.	Yes	Yes	Yes	Yes
User profile settings					
Enforce redirect after login	Mark the checkbox to restrict users to redirect URL after login. If unmarked, users can set their personal URL after login in their profile; redirect URL after login is used only as fallback.	Yes	No	No	No
Redirect URL after login	Specify the redirect URL after login.				
Access to services					
Read-write access to services	Select read-write access to services: **None** - no access at all **All** - access to all services is read-write **Service list** - select services for read-write access The read-write access, if granted, takes precedence over the read-only access settings and is dynamically inherited by the child services.	All	All	None	None
Read-write access to services with tag	Specify tag name and, optionally, value to additionally grant read-write access to services matching the tag. This option is available if 'Service list' is selected in the *Read-write access to services* parameter. The read-write access, if granted, takes precedence over the read-only access settings and is dynamically inherited by the child services.				

Read-only access to services	Select read-only access to services: **None** - no access at all **All** - access to all services is read-only **Service list** - select services for read-only access The read-only access does not take precedence over the read-write access and is dynamically inherited by the child services.			All	All
Read-only access to services with tag	Specify tag name and, optionally, value to additionally grant read-only access to services matching the tag. This option is available if 'Service list' is selected in the *Read-only access to services* parameter. The read-only access does not take precedence over the read-write access and is dynamically inherited by the child services.				
Access to modules					
<Module name>	Allow/deny access to a specific module. Only enabled modules are shown in this section. It is not possible to grant or restrict access to a module that is currently disabled.	Yes	Yes	Yes	Yes
Default access to new modules	This option specifies how new modules and widgets will be accessible after a Zabbix upgrade. It also applies to modules and widgets added in the *Administration > General > Modules section*.				
Access to API					
Enabled	Enable/disable access to API.	Yes	No	No	No
API methods	Select either *Allow list* to allow, or *Deny list* to deny the API methods specified in the search field. Note that it is not possible to allow some API methods and deny others. In the search field, start typing the method name, then select the method from the auto-complete list. You can also click the *Select* button and select methods from the full list available for this user type. Note that if certain action from the Access to actions block is unchecked, users will not be able to use API methods related to this action. Wildcards are supported. Examples: `dashboard.*` (all methods of 'dashboard.' API service) `*` (any method), `*.export` (methods with '.export' name from all API services). If no methods are specified, *Allow list* denies access to all API methods, while *Deny list* allows access to all API methods available for the user type. To allow all methods with *Allow list*, specify `*`. To block API access completely, clear the *Enabled* checkbox instead of leaving the *Allow list* empty. The API methods list does not apply to methods that do not require authentication, such as `apiinfo.version` and `user.login`.				

Access to actions					
Create and edit dashboards	Clearing this checkbox will also revoke the rights to use `.create`, `.update` and `.delete` API methods for the corresponding elements.	Yes	Yes	Yes	No
Create and edit maps					
Create and edit maintenance				No	
Add problem comments	Clearing this checkbox will also revoke the rights to perform corresponding action via `event.acknowledge` API method.			Yes	
Change severity					
Acknowledge problems					
Suppress problems					
Close problems					
Execute scripts	Clearing this checkbox will also revoke the rights to use the `script.execute` API method.				
Manage API tokens	Clearing this checkbox will also revoke the rights to use all `token.` API methods.				
Manage scheduled reports	Clearing this checkbox will also revoke the rights to use all `report.` API methods.			No	
Manage SLA	Enable/disable the rights to manage SLA.				
Invoke ”Execute now” on read-only hosts	Allow users to use the ”Execute now” option in Latest data for items on read-only hosts.			Yes	
Change problem ranking	Allow users to change the problem ranking from cause to symptom and vice versa.				
Create and edit own media	Allow users to create and edit their own media.				
Create and edit user media	Allow users to create and edit media for other users. This option is only available for Super admin users.		No	No	
Select ”Server” for monitoring and discovery	Clearing this checkbox restricts the creation of hosts monitored by the server and network discovery rules executed by the server, as well as the testing of items with the server. The *Server* option will be disabled in the host configuration for the *Monitored by* parameter, in the network discovery rule configuration for the *Discovery by* parameter, and in item testing for the *Test with* parameter.		Yes	No	
Default access to new actions	This option specifies how new actions will be accessible after a Zabbix upgrade.		Yes	Yes	
Devices					
Enabled	Enable/disable access to *User settings* > *Devices* (accessible to all users) and *Users* > *Devices* (accessible only to Super admins).	Yes	No	No	No
Manage own devices	Allow users to manage their own devices in *User settings* > *Devices*.				
Manage user devices	Allow users to manage devices belonging to other users in *Users* > *Devices*.				
Default access to new actions	This option specifies how new actions will be accessible after a Zabbix upgrade.				

See also:

- Configuring a user

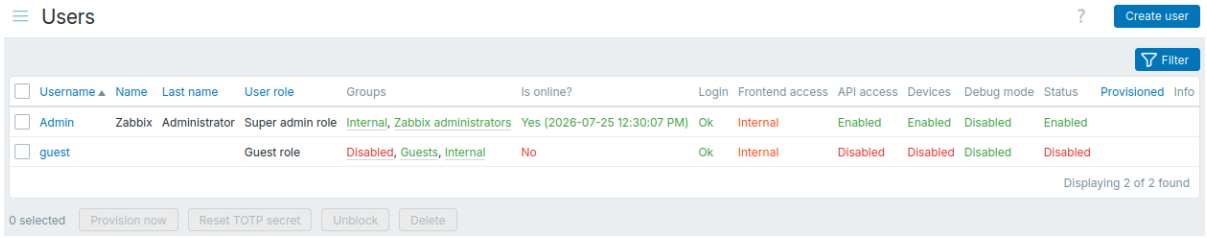
Users

Overview

In the *Users* > *Users* section users of the system are maintained.

Users

A listing of existing users with their details is displayed.



Displayed data:

Column	Description
Username	Username for logging into Zabbix. Clicking the username opens the user configuration form.
Name	First name of the user.
Last name	Second name of the user.
User role	User role is displayed.
Groups	Groups that the user is a member of are listed. Clicking the user group name opens the user group configuration form. Disabled groups are displayed in red.
Is online?	The on-line status of the user is displayed - *Yes* or *No*. The time of last user activity is displayed in parentheses.
Login	The login status of the user is displayed - *Ok* or *Blocked*. A user can become temporarily blocked upon exceeding the number of unsuccessful login attempts set in the *Administration > General > Other* section (five by default). By clicking *Blocked* you can unblock the user.
Frontend access	Frontend access level is displayed - *System default*, *Internal*, *LDAP*, or *Disabled*, depending on the one set for the whole user group.
API access	API access status is displayed - *Enabled* or *Disabled*, depending on the one set for the user role. If API access is enabled, mouseover on *Enabled* displays the API methods of the user role under *Allowed methods* or *Denied methods*, depending on the API methods mode. *None* is displayed if no methods are specified.
Devices	Device access status is displayed - *Enabled* or *Disabled*, depending on the one set for the user role. This column is displayed if `EnableMobileDevices` is enabled.
Debug mode	Debug mode status is displayed - *Enabled* or *Disabled*, depending on the one set for the whole user group.
Status	User status is displayed - *Enabled* or *Disabled*, depending on the one set for the whole user group.
Provisioned	The date when the user was last provisioned is displayed. Used for users created by JIT provisioning from LDAP/SAML.
Info	Information about errors is displayed. A yellow warning is displayed for users without user groups. A red warning is displayed for users without roles, and for users without roles and user groups.

To configure a new user, click the *Create user* button in the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

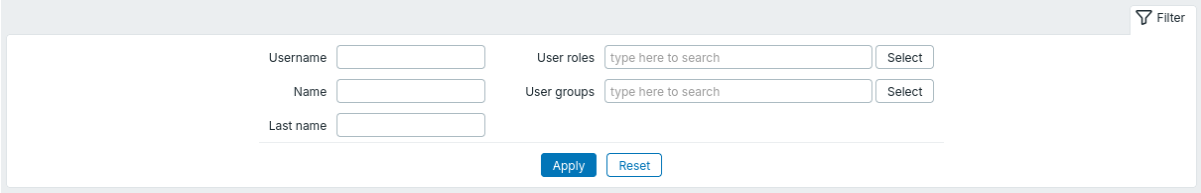
- *Provision now* - update user information from LDAP (this option is only enabled if an LDAP user is selected)
- *Reset TOTP secret* - reset user TOTP secrets for all TOTP methods and delete the user session (this option is only enabled if MFA is enabled; for users without TOTP secrets, their session will not be deleted)
- *Unblock* - re-enable system access to blocked users
- *Delete* - delete the users

To use these options, mark the check-boxes before the respective users, then click the required button.

Using filter

You can use the filter to display only the users you are interested in. For better search performance, data is searched with macros unresolved.

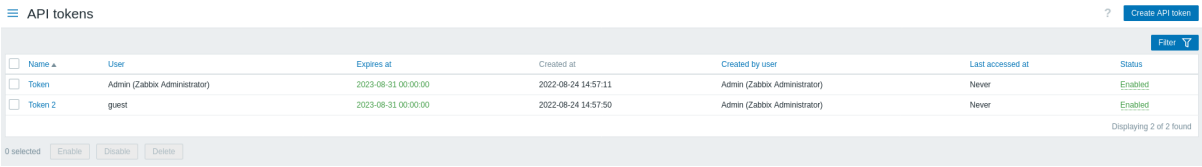
The *Filter* link is available above the list of users. If you click it, a filter becomes available where you can filter users by username, name, last name, user role and user group.



API tokens

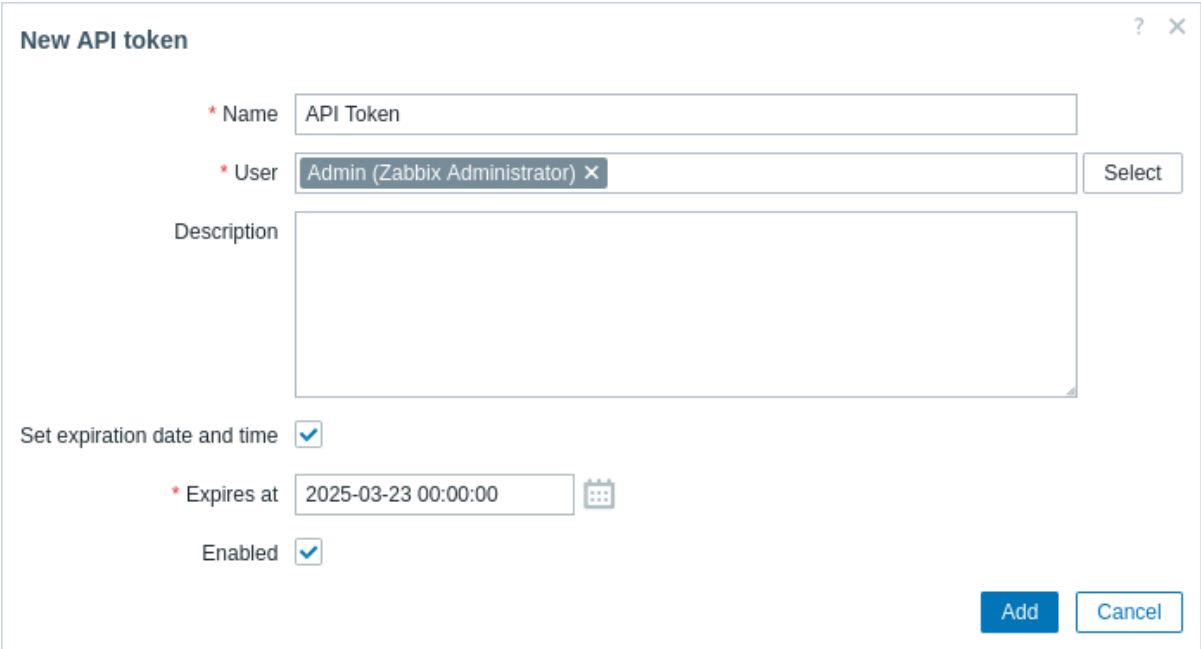
Overview

This section allows you to create and manage API tokens.
 API tokens are used for authentication in the Zabbix API.
 Each API token is assigned to a single user and inherits that user's permissions. Requests authenticated with an API token are authorized according to the associated user's role and permissions.



You may filter API tokens by name, users to whom the tokens are assigned, expiry date, users that created tokens, or status (enabled/disabled). Click the token status in the list to quickly enable/disable a token. You may also mass enable/disable tokens by selecting them in the list and then clicking the *Enable/Disable* buttons below the list.

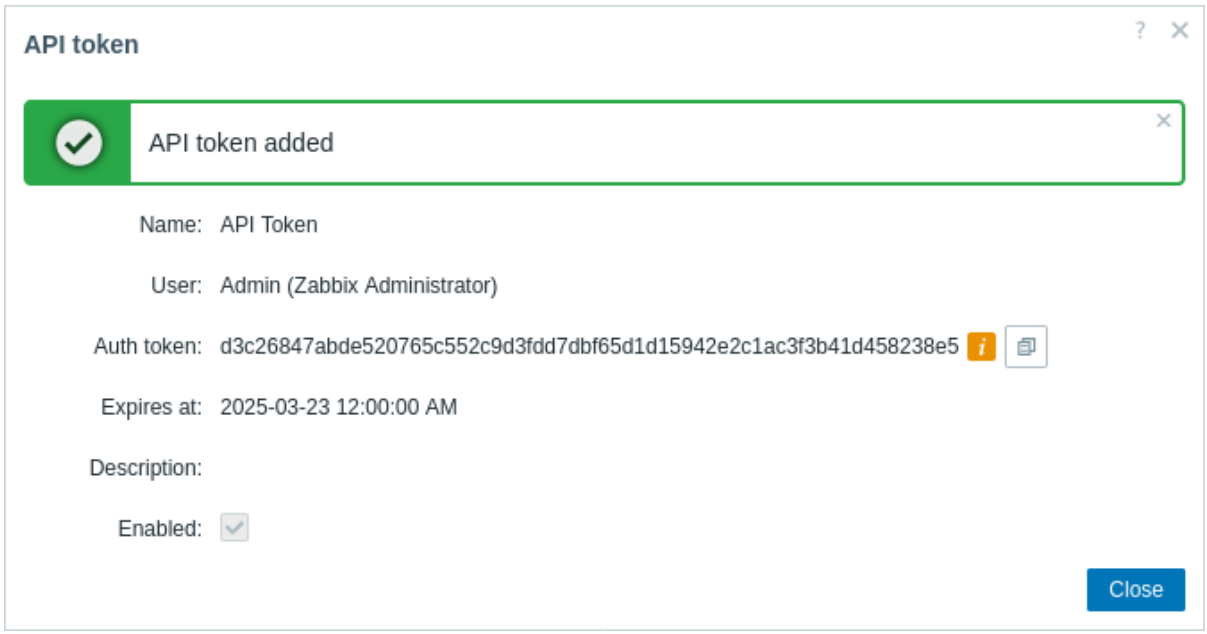
To create a new token, click the *Create API token* button at the upper-right corner, then fill out the required fields in the token configuration screen:



Parameter	Description
Name	Token's visible name.
User	User the token should be assigned to. To quickly select a user, start typing the username, first or last name, then select the required user from the auto-complete list. Alternatively, you can click the *Select* button and select a user from the full user list. A token can be assigned only to one user.
Description	Optional token description.
Set expiration date and time	Unmark this checkbox if a token should not have an expiry date.

Parameter	Description
Expiry date	Click the calendar icon to select token expiry date or enter the date manually in a format YYYY-MM-DD hh:mm:ss
Enabled	Unmark this checkbox if you need to create a token in a disabled state.

Click *Add* to create a token.



Copy the *Auth token* value and save in a safe place **before closing the page**, then click *Close*. The token will appear in the list.

Warning:
Auth token value cannot be viewed again later. It is only available immediately after creating a token. If you lose a saved token you will have to regenerate it and doing so will create a new authorization string.

Click the token name to edit the name, description, expiry date settings, or token status. Note that it is not possible to change to which user the token is assigned. Click the *Update* button to save changes. If a token has been lost or exposed, you may click the *Regenerate* button to generate new token value. A confirmation dialog box will appear, asking you to confirm this operation since after proceeding the previously generated token will become invalid.

Users without access to the *Administration* menu section can see and modify details of tokens assigned to them in the *User profile* > *API tokens* section only if *Manage API tokens* is allowed in their user role permissions.

Authentication

Overview

The *Users* > *Authentication* section allows you to specify the user authentication method for Zabbix and internal password requirements.

The available authentication methods are internal, HTTP, LDAP, SAML, and MFA authentication.

Default authentication

By default, Zabbix uses **internal** Zabbix authentication for all users.

It is possible to change the default authentication method to LDAP system-wide. To do so, navigate to the *LDAP* tab and configure LDAP parameters, then return to the *Authentication* tab and switch the *Default authentication* selector to LDAP.

Note that the authentication method can be fine-tuned on the user group level. Even if LDAP authentication is set globally, some user groups can still be authenticated by Zabbix. These groups must have frontend access set to Internal.

It is also possible to enable LDAP authentication only for specific user groups, if internal authentication is used globally. In this case LDAP authentication details can be specified and used for specific user groups whose frontend access must then be set to LDAP. If a user is included into at least one user group with LDAP authentication, this user will not be able to use the internal authentication method.

HTTP, SAML 2.0, and MFA authentication methods can be used in addition to the default authentication method.

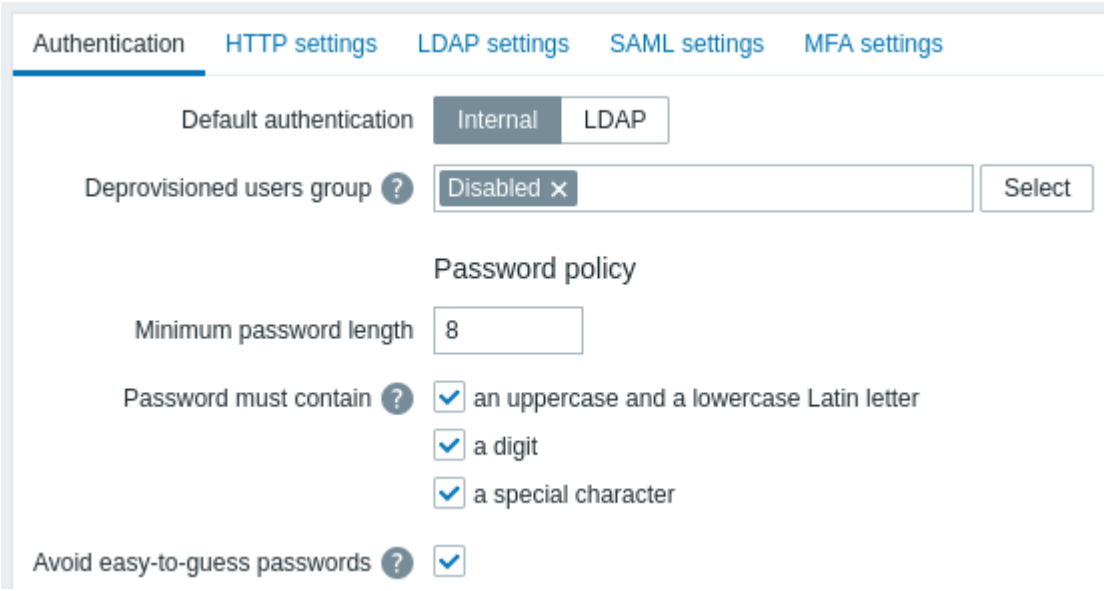
Zabbix supports just-in-time (JIT) provisioning that allows you to create user accounts in Zabbix the first time an external user authenticates and provision these user accounts. JIT provisioning is supported for LDAP and SAML.

See also:

- HTTP authentication
- LDAP authentication
- SAML authentication
- MFA authentication

Configuration

The *Authentication* tab allows you to set the default authentication method, specify a group for deprovisioned users and set password complexity requirements for Zabbix users.



Configuration parameters:

Parameter	Description
Default authentication	Select the default authentication method for Zabbix - *Internal* or *LDAP*.
Deprovisioned users group	Specify a user group for deprovisioned users. This setting is required only for JIT provisioning, regarding users that were created in Zabbix from LDAP or SAML systems, but no longer need to be provisioned. A disabled user group must be specified.
Minimum password length	By default, the minimum password length is set to 8. Supported range: 1-70. Note that passwords longer than 72 characters will be truncated.
Password must contain	Mark one or several checkboxes to require usage of specified characters in a password: - an uppercase and a lowercase Latin letter - a digit - a special character Hover over the question mark to see a hint with the list of characters for each option.
Avoid easy-to-guess passwords	If marked, a password will be checked against the following requirements: - must not contain user's name, surname, or username - must not be one of the common or context-specific passwords. The list of common and context-specific passwords is generated automatically from the list of NCSC "Top 100k passwords", the list of SecLists "Top 1M passwords" and the list of Zabbix context-specific passwords. Internal users will not be allowed to set passwords included in this list as such passwords are considered weak due to their common use.

Changes in password complexity requirements will not affect existing user passwords, but if an existing user chooses to change a password, the new password will have to meet current requirements. A hint with the list of requirements will be displayed next to the *Password* field in the user profile and in the user configuration form accessible from the *Users* > *Users* menu.

HTTP

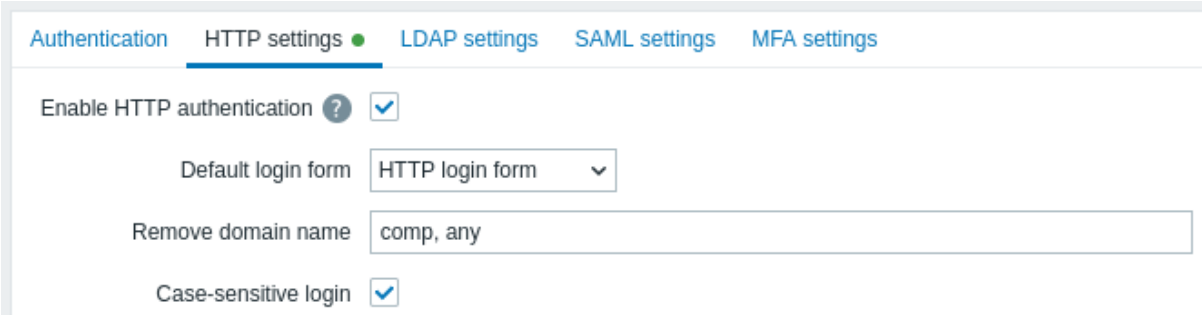
Overview

HTTP or web server-based authentication (for example: BasicAuthentication, NTLM/Kerberos) can be used to check user names and passwords. Note that a user must exist in Zabbix as well, however its Zabbix password will not be used.

Attention:
Be careful! Make sure that web server authentication is configured and works properly before switching it on.

Note:
HTTP authentication configuration can be disabled by setting `$ZBX_FEATURE_FLAGS['http_auth_enabled'] = false` in the frontend configuration file (`zabbix.conf.php`). In this case, the *HTTP settings* tab with HTTP authentication options will be hidden.

Configuration



Configuration parameters:

Parameter	Description
Enable HTTP authentication	Mark the checkbox to enable HTTP authentication. Hovering the mouse over will bring up a hint box warning that in the case of web server authentication, all users (even with frontend access set to LDAP/Internal) will be authenticated by the web server, not by Zabbix.
Default login form	Specify whether to direct non-authenticated users to: **Zabbix login form** - standard Zabbix login page. **HTTP login form** - HTTP login page. It is recommended to enable web-server based authentication for the `index_http.php` page only. If *Default login form* is set to 'HTTP login page' the user will be logged in automatically if web server authentication module will set valid user login in the `$_SERVER` variable. Supported `$_SERVER` keys are `PHP_AUTH_USER`, `REMOTE_USER`, `AUTH_USER`.
Remove domain name	A comma-delimited list of domain names that should be removed from the username. E.g., `comp,any` - if username is 'Admin@any', 'comp\Admin', user will be logged in as 'Admin'; if username is 'notacompany\Admin', login will be denied.
Case-sensitive login	Unmark the checkbox to disable case-sensitive login for usernames (enabled by default). Disabling case-sensitive login allows, for example, logging in as "admin" even if the Zabbix user is "Admin" or "ADMIN". Please note that if case-sensitive login is disabled and there are multiple Zabbix users with similar usernames (e.g., Admin and admin), the login for those users will always be denied with the following error message: "Authentication failed: supplied credentials are not unique."

Note:
For internal users who are unable to log in using HTTP credentials (with HTTP login form set as default) leading to the 401 error, you may want to add a `ErrorDocument 401 /index.php?form=default` line to basic authentication directives, which will redirect to the regular Zabbix login form.

LDAP

Overview

External LDAP authentication can be used to check user names and passwords.

Zabbix LDAP authentication works at least with Microsoft Active Directory and OpenLDAP.

If only LDAP sign-in is configured, then the user must also exist in Zabbix, however, its Zabbix password will not be used. If authentication is successful, then Zabbix will match a local username with the username attribute returned by LDAP.

User provisioning

It is possible to configure JIT (just-in-time) **user provisioning** for LDAP users. In this case, it is not required that a user already exists in Zabbix. The user account can be created when the user logs into Zabbix for the first time.

When an LDAP user enters their LDAP login and password, Zabbix checks the *default* LDAP server if this user exists. If the user exists and does not have an account in Zabbix yet, a new user is created in Zabbix and the user is able to log in.

A user created via JIT provisioning is associated with the LDAP server (directory) that is set as the default at the time of creation. Changing the default LDAP server later does not change or update the LDAP server linked to users that were already provisioned.

Attention:
If JIT provisioning is enabled, a user group for deprovisioned users must be specified in the *Authentication* tab.

JIT provisioning also allows provisioned user accounts to be updated based on changes in LDAP. For example, if a user is moved from one LDAP group to another, the user will also be moved from one group to another in Zabbix; if a user is removed from an LDAP group, the user will also be removed from the group in Zabbix and, if not belonging to any other group, added to the user group for deprovisioned users. Provisioned user accounts are updated based on the configured provisioning period or when the user logs into Zabbix.

Note that background provisioning is done by Zabbix frontend when the user is interacting with it or at least has an open frontend page in the browser. There are no dedicated background processes to provision users.

LDAP supports three ways to bind to the directory for authentication and searches:

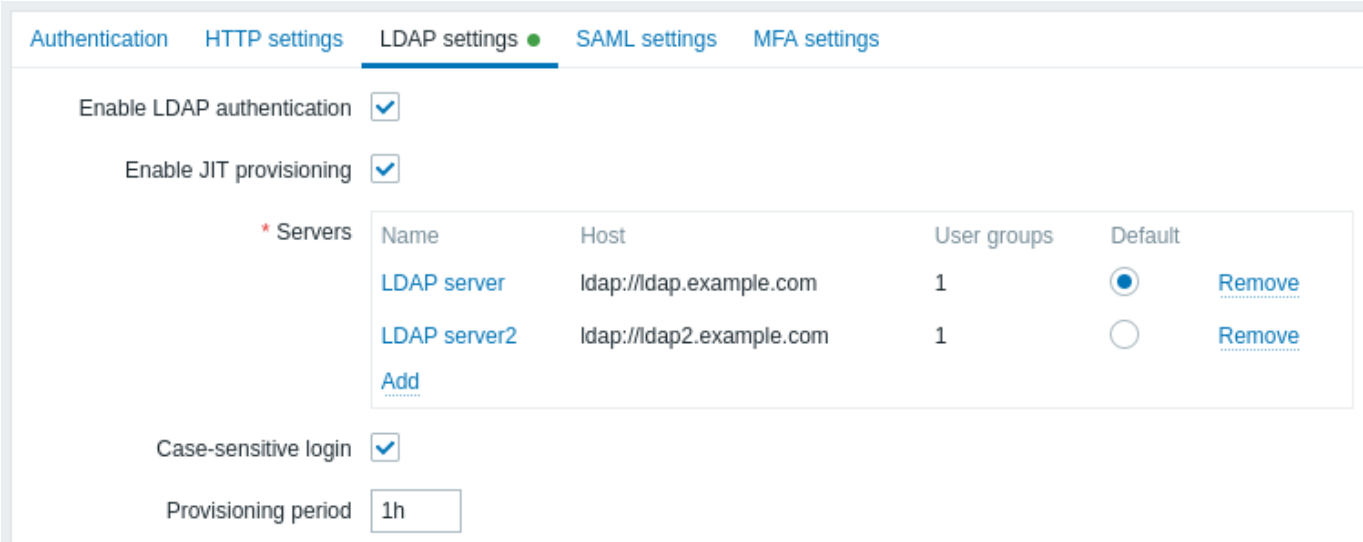
- Anonymous binding - no *Bind DN* / *Bind password* is provided and the LDAP server allows anonymous queries.
- Dedicated bind user (service account) - a specific LDAP account is set in *Bind DN* / *Bind password* and is used by Zabbix for searching and provisioning. This is the recommended and most flexible option because Zabbix can perform searches and background provisioning without end-user credentials.
- Direct user binding - Zabbix performs the bind using the credentials the user enters at login (no *Bind DN* / *Bind password* configured); this is configured by including a placeholder such as `uid=%{user}` in the Base DN. In this mode Zabbix only has access to the user's password during the interactive login. As a result, provisioning actions that require authenticating to LDAP outside of the user's login session (for example, using the *Provision now* button or running background provisioning when the user is not actively logged in) cannot authenticate and therefore will not work. Provisioning and updates for direct user binding happen only at the moment the user logs in.

Multiple servers

Several LDAP servers can be defined, if necessary. For example, a different server can be used to authenticate a different user group. Once LDAP servers are configured, in user group configuration it becomes possible to select the required LDAP server for the respective user group.

If a user is in multiple user groups and multiple LDAP servers, the first server in the list of LDAP servers sorted by name in ascending order will be used for authentication.

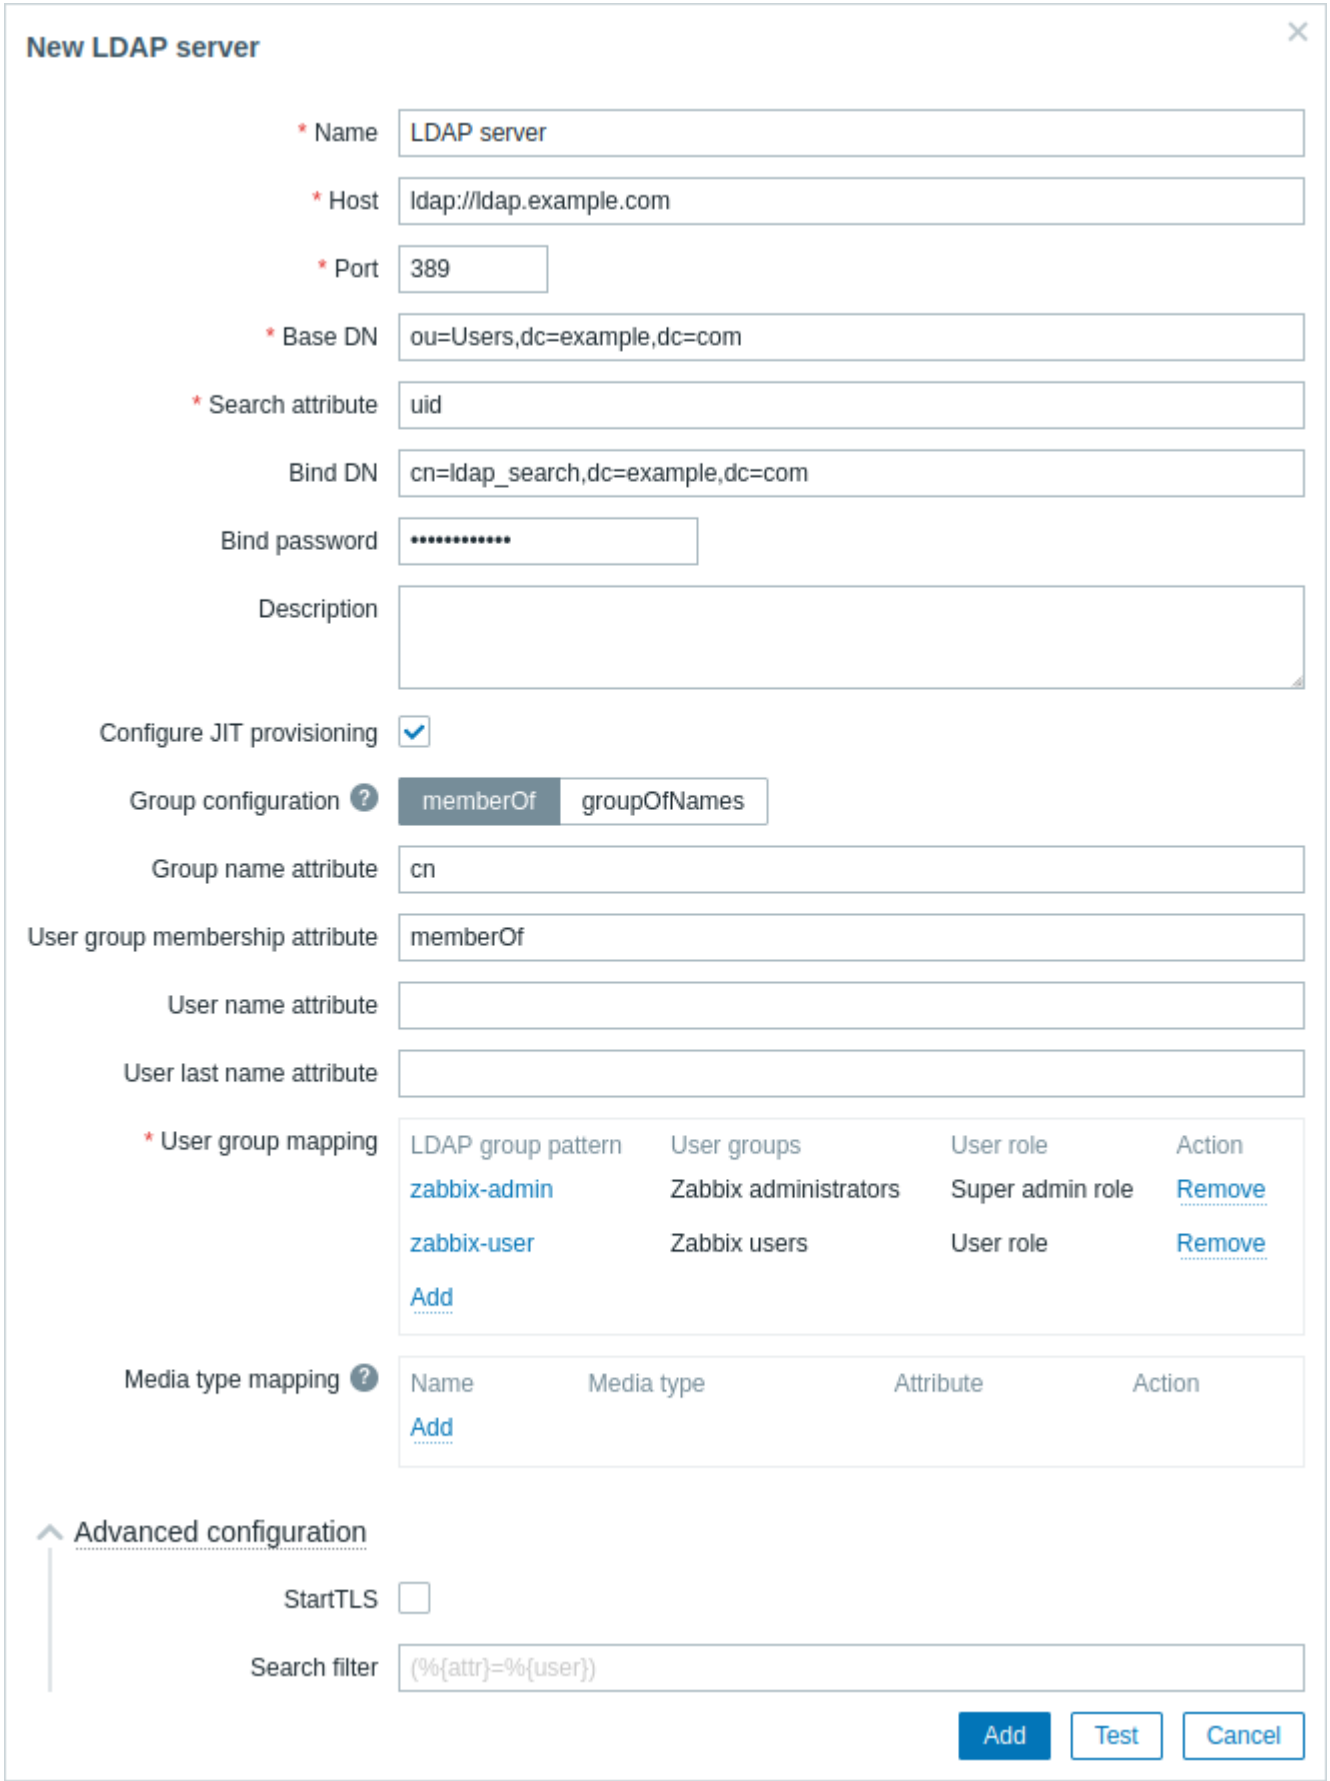
Configuration



Configuration parameters:

Parameter	Description
Enable LDAP authentication	Mark the checkbox to enable LDAP authentication.
Enable JIT provisioning	Mark the checkbox to enable JIT provisioning.
Servers	Click *Add* to configure an LDAP server (see LDAP server configuration below).
Case-sensitive login	Unmark the checkbox to disable case-sensitive login for usernames (enabled by default). Disabling case-sensitive login allows, for example, logging in as ”admin” even if the Zabbix user is ”Admin” or ”ADMIN”. Please note that if case-sensitive login is disabled and there are multiple Zabbix users with similar usernames (e.g., Admin and admin), the login for those users will always be denied with the following error message: ”Authentication failed: supplied credentials are not unique.”
Provisioning period	Set the provisioning period, i.e. how often the logged in user will be provisioned while working with the frontend.

LDAP server configuration



LDAP server configuration parameters:

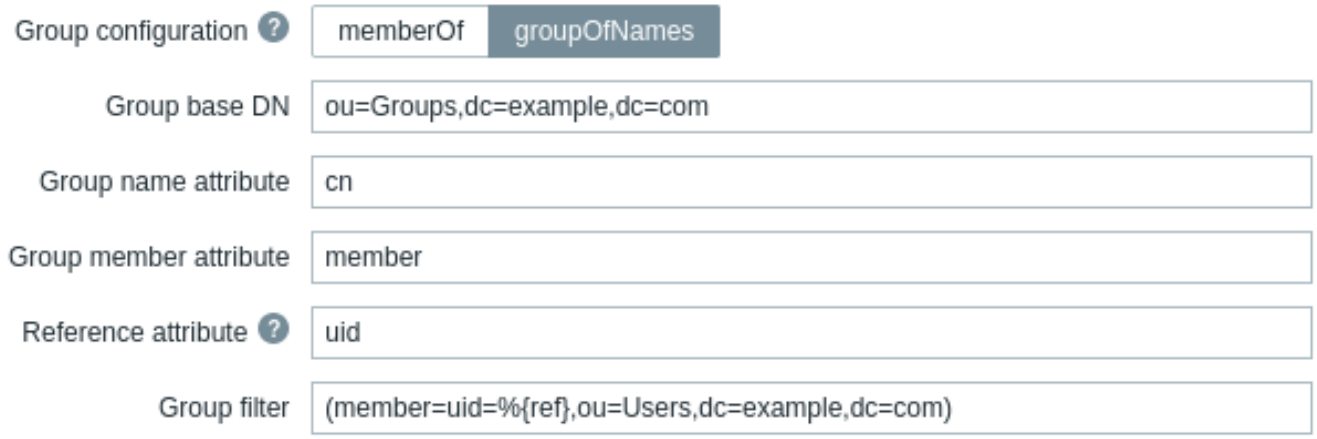
Parameter	Description
Name	Name of the LDAP server in Zabbix configuration.

Parameter	Description
Host	Hostname, IP or URI of the LDAP server. Examples: ldap.example.com, 127.0.0.1, ldap://ldap.example.com For secure LDAP server, use *ldaps* protocol and hostname. Example: ldaps://ldap.example.com With OpenLDAP 2.x.x and later, a full LDAP URI of the form ldap://hostname:port or ldaps://hostname:port may be used.
Port	Port of the LDAP server. Default is 389. For secure LDAP connection port number is normally 636. Not used when using full LDAP URIs.
Base DN	Base path to user accounts in LDAP server: ou=Users,ou=system (for OpenLDAP), DC=company,DC=com (for Microsoft Active Directory) uid=%{user},dc=example,dc=com (for direct user binding, see a note below)
Search attribute	LDAP account attribute used for search: uid (for OpenLDAP), sAMAccountName (for Microsoft Active Directory)
Bind DN	LDAP account for binding and searching over the LDAP server, examples: uid=ldap_search,ou=system (for OpenLDAP), CN=ldap_search,OU=user_group,DC=company,DC=com (for Microsoft Active Directory) Anonymous binding is also supported. Note that anonymous binding potentially opens up domain configuration to unauthorized users (information about users, computers, servers, groups, services, etc.). For security reasons, disable anonymous binds on LDAP hosts and use authenticated access instead.
Bind password	LDAP password of the account for binding and searching over the LDAP server.
Description	Description of the LDAP server.
Configure JIT provisioning	Mark this checkbox to show options related to JIT provisioning.
Group configuration	Select the group configuration method: **memberOf** - by searching users and their group membership attribute **groupOfNames** - by searching groups through the member attribute Note that memberOf is preferable as it is faster; use groupOfNames if your LDAP server does not support `memberOf` or group filtering is required.
Group name attribute	Specify the attribute to get the group name from all objects in the `memberOf` attribute (see the *User group membership attribute* field) The group name is necessary for user group mapping.
User group membership attribute	Specify the attribute that contains information about the groups that the user belongs to (e.g., `memberOf`). For example, the memberOf attribute may hold information like this: `memberOf=cn=zabbix-admin,ou=Groups,dc=example,dc=com` This field is available only for the memberOf method.
User name attribute	Specify the attribute that contains the user's first name.
User last name attribute	Specify the attribute that contains the user's last name.
User group mapping	Map an LDAP user group pattern to Zabbix user group and user role. This is required to determine what user group/role the provisioned user will get in Zabbix. Click *Add* to add a mapping. The *LDAP group pattern* field supports wildcards. The group name must match an existing group. If an LDAP user matches several Zabbix user groups, the user becomes a member of all of them. If a user matches several Zabbix user roles, the user will get the one with the highest permission level among them.
Media type mapping	Map the user's LDAP media attributes (e.g., email) to Zabbix user media for sending notifications (the attribute value is used as the media *Send to* field).
Advanced configuration	Click the *Advanced configuration* header to display advanced configuration options (see below).
StartTLS	Mark the checkbox to use the StartTLS operation when connecting to LDAP server. The connection will fall if the server doesn't support StartTLS. StartTLS cannot be used with servers that use the *ldaps* protocol.

Parameter	Description
Search filter	Define a custom string when authenticating a user in LDAP. The following placeholders are supported: `%{attr}` - search attribute name (uid, sAMAccountName) `%{user}` - user username value to authenticate For example, to carry out a case-sensitive search within the case-insensitive LDAP or Microsoft Active Directory environment, the string can be defined as follows: `(%{attr}:caseExactMatch:=%{user})`. If the filter is not customized, LDAP will use the default: `(%{attr}=%{user})`.

Note:
To configure an LDAP server for **direct user binding**, append an attribute uid=%{user} to the *Base DN* parameter (for example,*uid=%{user},dc=example,dc=com*) and leave *BindDN* and *Bind password* parameters empty. When authenticating, a placeholder %{user} will be replaced by the username entered during login.
 With direct user binding Zabbix only has access to the user's credentials during the interactive login. Therefore, provisioning tasks that operate without an interactive login (for example, the *Provision now* button or the `user.provision` API method) will ignore users that authenticate using direct user binding because Zabbix cannot bind on their behalf. Use anonymous binding or a dedicated bind user (service account) if you need provisioning and LDAP searches to work from the frontend without requiring the end-user to be logged in.

The following fields are specific to "groupOfNames" as the *Group configuration* method:



Parameter	Description
Group base DN	Base path to the groups in LDAP server.
Group name attribute	Specify the attribute to get the group name in the specified base path to groups. The group name is necessary for user group mapping.
Group member attribute	Specify the attribute that contains information about the members of the group in LDAP (e.g., `member`).
Reference attribute	Specify the reference attribute for the group filter (see the *Group filter* field). Then use `%{ref}` in the group filter to get values for the attribute specified here.
Group filter	Specify the filter to retrieve the group that the user is member of. For example, `(member=uid=%{ref},ou=Users,dc=example,dc=com)` will match "User1" if the member attribute of the group is `uid=User1,ou=Users,dc=example,dc=com` and will return the group that "User1" is a member of.

Warning:
In case of trouble with certificates, to make a secure LDAP connection (ldaps) work you may need to add a `TLS_REQCERT allow` line to the `/etc/openldap/ldap.conf` configuration file. It may decrease the security of connection to the LDAP catalog.

Note:
It is recommended to create a separate LDAP account (*Bind DN*) to perform binding and searching over the LDAP server with minimal privileges in the LDAP instead of using real user accounts (used for logging in Zabbix frontend).
 Such an approach provides more security and does not require changing the *Bind password* when the user changes his own password in the LDAP server.
 In the table above it's the *ldap_search* account name.

Testing access

The *Test* button allows you to test user access:

Parameter	Description
Login	LDAP user name to test (prefilled with the current user name from Zabbix frontend). This user name must exist in the LDAP server. Zabbix will not activate LDAP authentication if it is unable to authenticate the test user.
User password	LDAP user password to test.

SAML

Overview

SAML 2.0 authentication can be used to sign in to Zabbix.

If SAML sign-in is configured without JIT user provisioning, the user must already exist in Zabbix; the user's Zabbix password is not used for authentication. After successful SAML authentication, Zabbix matches the username attribute received from SAML with an existing local Zabbix username.

User provisioning

It is possible to configure JIT (just-in-time) **user provisioning** for SAML users. In this case, it is not required that a user already exists in Zabbix. The user account can be created when the user logs into Zabbix for the first time.

Attention:
If JIT provisioning is enabled, a user group for deprovisioned users must be specified in the *Authentication* tab.

On top of JIT provisioning it is also possible to enable and configure SCIM (System for Cross-domain Identity Management) provisioning - *continuous* user account management for those users that have been created by user provisioning. SCIM provisioning requires a Zabbix API token (with Super admin permissions) for authentication into Zabbix.

For example, if a user is moved from one SAML group to another, the user will also be moved from one group to another in Zabbix; if a user is removed from a SAML group, the user will also be removed from the group in Zabbix and, if not belonging to any other group, added to the user group for deprovisioned users.

If SCIM is enabled and configured, a SAML user will be provisioned at the moment the user logs into Zabbix and continuously updated based on changes in SAML. Already existing SAML users will not be provisioned, and only provisioned users will be updated. Note that only valid media will be added to a user when the user is provisioned or updated.

If SCIM is not enabled, a SAML user will be provisioned (and later updated) at the moment the user logs into Zabbix.

Note:
If SAML authentication is enabled, users will be able to choose between logging in locally or via SAML single sign-on. If JIT provisioning is used, then only single sign-on is possible.

Setting up identity provider

In order to work with Zabbix, a SAML identity provider (onelogin.com, auth0.com, okta.com, etc.) needs to be configured in the following way:

- *Assertion Consumer URL* should be set to `<path_to_zabbix_ui>/index_sso.php?acs`
- *Single Logout URL* should be set to `<path_to_zabbix_ui>/index_sso.php?sls`

`<path_to_zabbix_ui>` examples: `https://example.com/zabbix/ui`, `http://another.example.com/zabbix`, `http://<any_public_ip_address>/zabbix`

Setting up Zabbix

Attention:
It is required to install php-openssl if you want to use SAML authentication in the frontend.

To use SAML authentication Zabbix should be configured in the following way:

1. Provide the SP private key and certificates. An IdP certificate must be provided (paste or choose file) to activate SAML via the frontend. An SP certificate and SP private key must be provided if *Sign* or *Encrypt* options are selected.

If `$SSO['CERT_STORAGE'] = 'database'` is set in the frontend configuration file (`zabbix.conf.php`), you can paste the certificate text or upload the certificate file in the frontend during SAML configuration—no files are required on the filesystem.

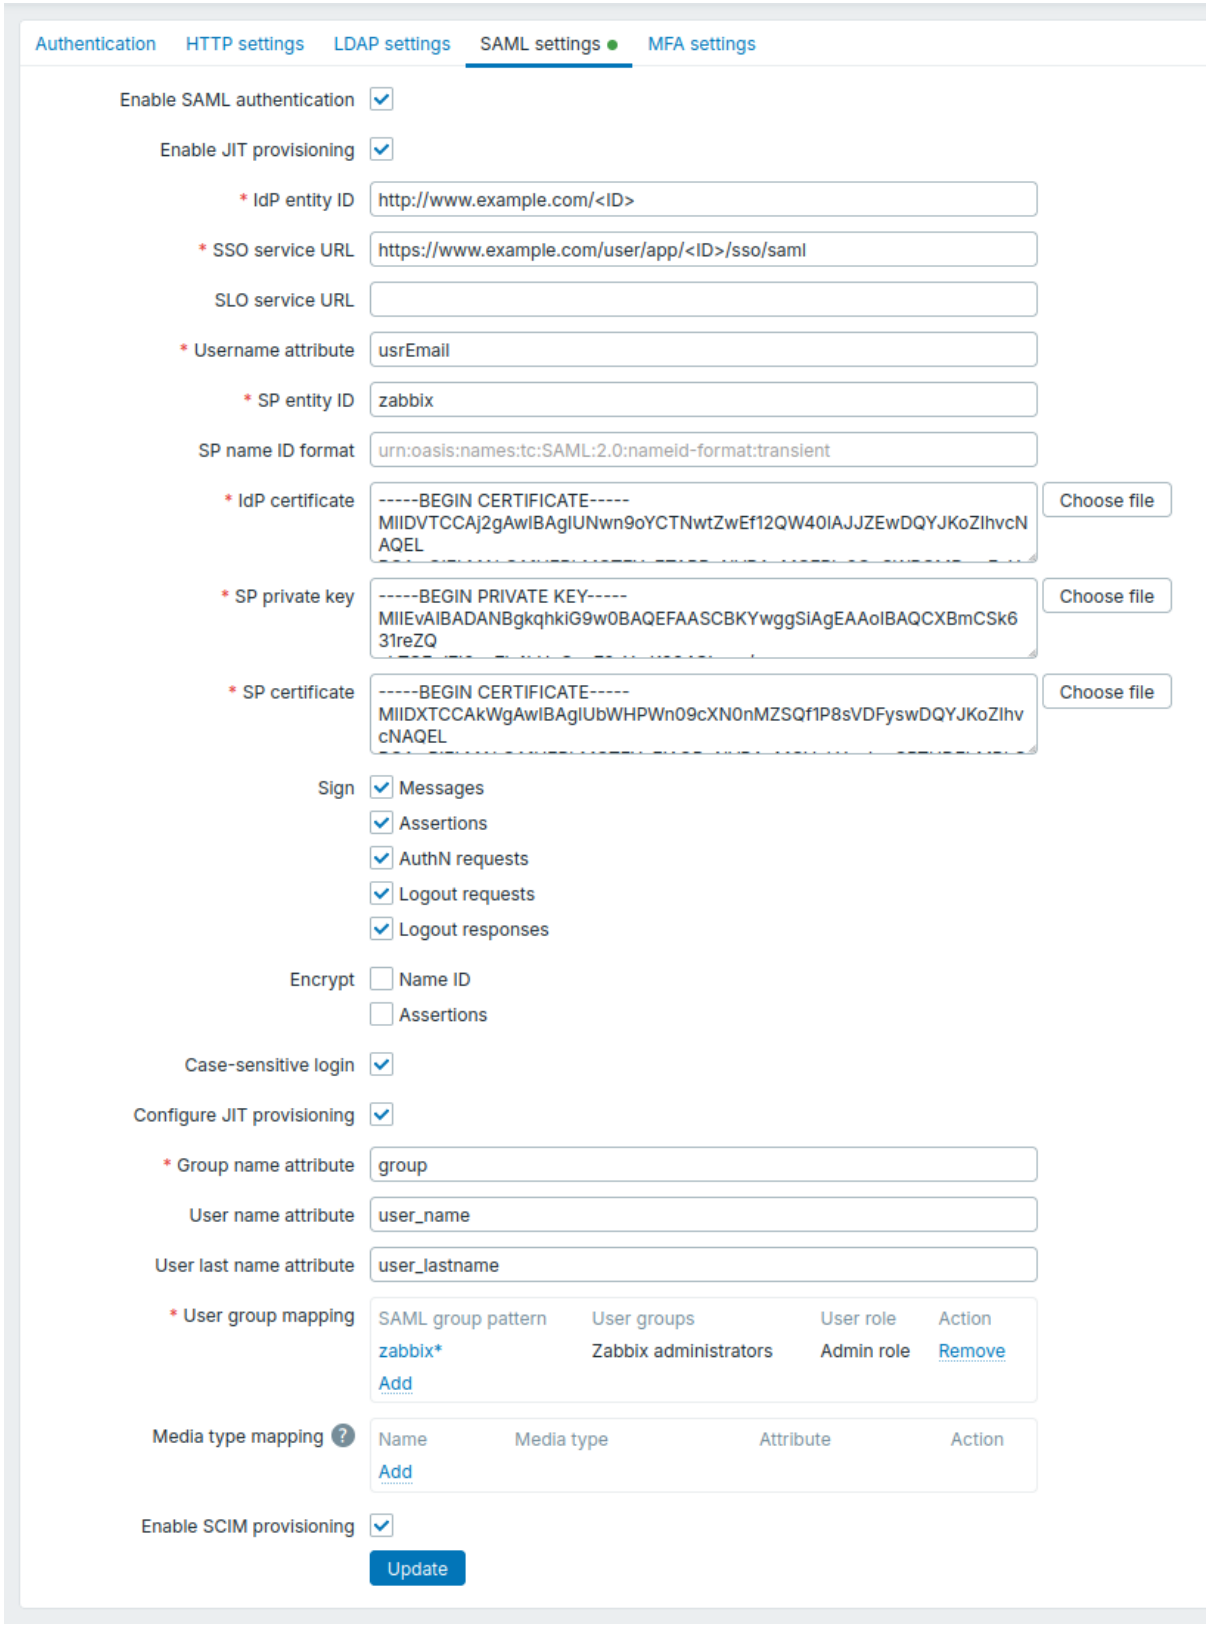
If `$SSO['CERT_STORAGE'] = 'file'` is set in the configuration file, the certificate must be made available on the filesystem (by default in `ui/conf/certs` or the path configured in the configuration file) and the frontend will not store certificates in the database. Note that if `$SSO['CERT_STORAGE']` is not set or is commented out, file storage is assumed.

By default, Zabbix will look in the following locations:

- ui/conf/certs/sp.key - SP private key file
- ui/conf/certs/sp.crt - SP cert file
- ui/conf/certs/idp.crt - IDP cert file

Attention:
Certificate material may be imported via the frontend by pasting text or uploading files. Although unencrypted imports are allowed for convenience in some environments, using encrypted transport is strongly recommended (for example, enable HTTPS/TLS for the Zabbix frontend). When SAML certificates/keys are stored in the database, enable TLS/SSL for database connections to protect traffic between the application and the database.
 Encrypt database backups that may contain SAML certificates or private keys, and limit access to backup files and to the database tables that store SAML material. Restrict database user privileges to the minimum necessary.

2. All of the most important settings can be configured in the Zabbix frontend. However, it is possible to specify additional settings in the configuration file.



Configuration parameters, available in the Zabbix frontend:

Parameter	Description
Enable SAML authentication	Mark the checkbox to enable SAML authentication.
Enable JIT provisioning	Mark the checkbox to enable JIT user provisioning.
IDP entity ID	The unique entity identifier within the SAML identity provider.
SSO service URL	The URL users will be redirected to when logging in.
SLO service URL	The URL users will be redirected to when logging out. If left empty, the SLO service will not be used.

Parameter	Description
Username attribute	SAML attribute to be used as a username when logging into Zabbix. The list of supported values is determined by the identity provider. Examples: uid userprincipalname samaccountname username userusername urn:oid:0.9.2342.19200300.100.1.1 urn:oid:1.3.6.1.4.1.5923.1.1.1.13 urn:oid:0.9.2342.19200300.100.1.44
SP entity ID	The unique service provider identifier (if not matching, the operation will be rejected). It is possible to specify a URL or any string of data.
SP name ID format	Request a particular name ID format in the response. Examples: urn:oasis:names:tc:SAML:2.0:nameid-format:persistent urn:oasis:names:tc:SAML:2.0:nameid-format:unspecified urn:oasis:names:tc:SAML:2.0:nameid-format:transient
IdP certificate	Contents of the service provider (SP) certificate for SAML single sign-on (SSO) service setup.
SP private key	Contents of the service provider (SP) Private Key for SAML single sign-on (SSO) service setup. Provides secure authentication and data exchange with the Identity Provider (IdP).
SP certificate	Contents of the service provider (SP) certificate for SAML single sign-on (SSO) service setup.
Sign	Unmark the checkboxes to disable SAML signatures for the corresponding entities: *Messages* *Assertions* *AuthN requests* *Logout requests* *Logout responses*
Encrypt	Mark the checkboxes to select entities for which SAML encryption should be enabled: *Name ID* *Assertions*
Case-sensitive login	Unmark the checkbox to disable case-sensitive login for usernames (enabled by default). Disabling case-sensitive login allows, for example, logging in as ”admin” even if the Zabbix user is ”Admin” or ”ADMIN”. Please note that if case-sensitive login is disabled and there are multiple Zabbix users with similar usernames (e.g., Admin and admin), the login for those users will always be denied with the following error message: ”Authentication failed: supplied credentials are not unique.”
Configure JIT provisioning	Mark this checkbox to show options related to JIT user provisioning.
Group name attribute	Specify the group name attribute for JIT user provisioning.
User name attribute	Specify the user name attribute for JIT user provisioning.
User last name attribute	Specify the user last name attribute for JIT user provisioning.
User group mapping	Map a SAML user group pattern to Zabbix user group and user role. This is required to determine what user group/role the provisioned user will get in Zabbix. Click *Add* to add a mapping. The *SAML group pattern* field supports wildcards. The group name must match an existing group. If a SAML user matches several Zabbix user groups, the user becomes a member of all of them. If a user matches several Zabbix user roles, the user will get the highest permission level among them.
Media type mapping	Map the user’s SAML media attributes (e.g., email) to Zabbix user media for sending notifications.
Enable SCIM provisioning	Mark this checkbox to enable SCIM 2.0 provisioning.

See examples of configuring SAML identity providers for sign-in and user provisioning into Zabbix with:

- Microsoft Azure AD
- Okta
- Onelogin

Notes on SCIM provisioning

For SCIM provisioning specify the path to the Zabbix frontend and append `api_scim.php` to it, on the identity provider side, i.e.:

`https://<path-to-zabbix-ui>/api_scim.php`

User attributes that are used in Zabbix (username, user name, user lastname and media attributes) need to be added as custom attributes and, if necessary, external namespace should be the same as user schema: `urn:ietf:params:scim:schemas:core:2.0:User`.

Advanced settings

Additional SAML parameters can be configured in the Zabbix frontend configuration file (`zabbix.conf.php`):

- `$SSO['SP_KEY']`
- `$SSO['SP_CERT']`
- `$SSO['IDP_CERT']`
- `$SSO['SETTINGS']`

Configuration examples

Example 1: Connecting to a Zabbix instance with a load balancer using TLS/SSL

If a load balancer is used to connect to a Zabbix instance, and the load balancer uses TLS/SSL while Zabbix does not, specify the `baseurl`, `strict`, and `use_proxy_headers` parameters as follows:

```
$SSO['SETTINGS'] = [
    'strict' => false,
    'baseurl' => 'https://zabbix.example.com/zabbix/',
    'use_proxy_headers' => true
];
```

Example 2: Configuring custom SAML signature and digest algorithms

```
$SSO['SETTINGS'] = [
    'security' => [
        'signatureAlgorithm' => 'http://www.w3.org/2001/04/xmldsig-more#rsa-sha384'
        'digestAlgorithm' => 'http://www.w3.org/2001/04/xmldsig-more#sha384',
        // ...
    ],
    // ...
];
```

Example 3: Configuring SSO with Kerberos authentication using ADFS

```
$SSO['SETTINGS'] = [
    'security' => [
        'requestedAuthnContext' => [
            'urn:oasis:names:tc:SAML:2.0:ac:classes:Kerberos',
        ],
        'requestedAuthnContextComparison' => 'exact'
    ]
];
```

In this case, in the SAML configuration *SP name ID* field set:

```
urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified
```

MFA

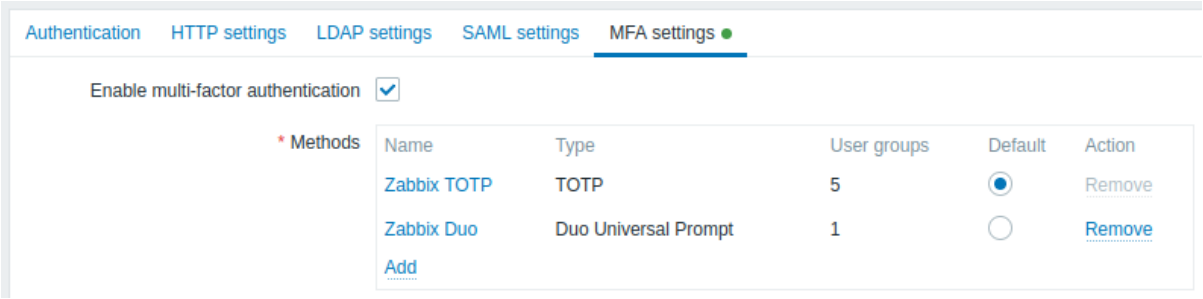
Overview

Multi-factor authentication (MFA) can be used to sign in to Zabbix, providing an additional layer of security beyond just a username and password.

With MFA, the user must exist in Zabbix, must provide Zabbix credentials when logging in, and must also prove their identity by other means, usually, a code generated by an authenticator app on the user's phone.

Multiple MFA methods are available, allowing users to choose the option that best fits their security requirements and preferences. These methods are Time-Based One-Time Password (TOTP) and Duo Universal Prompt.

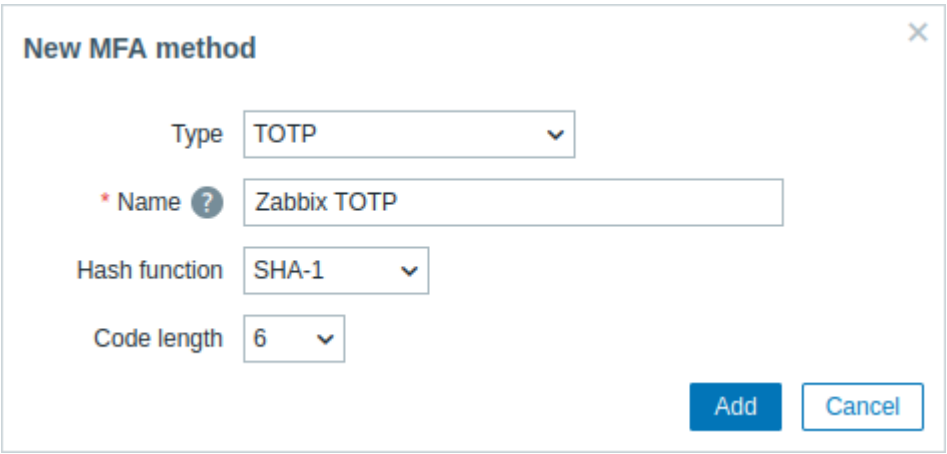
Configuration



Configuration parameters:

Parameter	Description
Enable multi-factor authentication	Mark the checkbox to enable multi-factor authentication.
Methods	Click *Add* to configure an MFA method (see method configuration below).

Method configuration



Method configuration parameters:

Parameter	Description
Type	Select the type of MFA method: **TOTP** - use an authenticator app to generate time-based one-time passwords; **Duo Universal Prompt** - use the Duo authentication service to provide multi-factor authentication.
Name	Enter a name that is displayed as the account name to all MFA users in authenticator apps (for example, ”Zabbix”).
Hash function	Select the hash function (SHA-1, SHA-256, or SHA-512) for generating TOTP codes. This parameter is available if the MFA method type is set to ”TOTP”. Note that choosing SHA-256 or SHA-512 may significantly limit compatibility, as many applications currently do not support these functions.
Code length	Select the verification code length (6 or 8). This parameter is available if MFA method type is set to ”TOTP”.
API hostname	Enter the API hostname provided by the Duo authentication service. This parameter is available if MFA method type is set to ”Duo Universal Prompt”.
Client ID	Enter the client ID provided by the Duo authentication service. This parameter is available if MFA method type is set to ”Duo Universal Prompt”.
Client secret	Enter the client secret provided by the Duo authentication service. This parameter is available if MFA method type is set to ”Duo Universal Prompt”.

Configuration examples

This section provides examples of configuring MFA using Time-Based One-Time Password (TOTP) and Duo Universal Prompt.

TOTP

For TOTP, users must verify their identity using an authenticator app (for example, the Google Authenticator app).

1. Go to the MFA settings in Zabbix under *Users* > *Authentication* and enable multi-factor authentication.

2. Add a new MFA method with the following configuration:

    - Type: TOTP
    - Name: Zabbix TOTP
    - Hash function: SHA-1
    - Code length: 6

3. Click the *Add* button and then the *Update* button.

4. Go to *Users* > *User groups* and create a new user group with the following configuration:

    - Group name: TOTP group
    - Users: Admin
    - Multi-factor authentication: Default (or ”Zabbix TOTP” if it is not set as default)

5. Log out of Zabbix and log back in using your credentials. Upon successful login, you will be prompted to enroll in MFA, displaying a QR code and a secret key.



6. Scan the QR code or enter the secret key into the Google Authenticator app. The app will generate a verification code which you should enter to complete the login process.

7. For subsequent logins, retrieve the verification code from the Google Authenticator app and enter it during login.

Duo Universal Prompt

For Duo Universal Prompt, users must verify their identity using the Duo Mobile authenticator app.

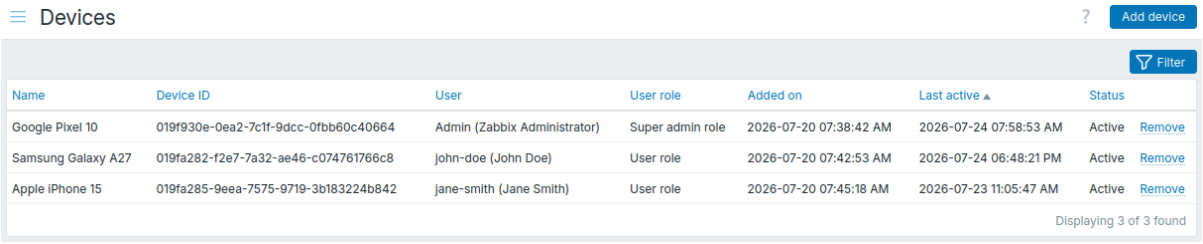
Attention:
The Duo Universal Prompt MFA method requires the installation of the php-curl extension, access to Zabbix over HTTPS, and permission for outbound connections to Duo servers. Moreover, if you have enabled Content Security Policy (CSP) on the web server, make sure to add "duo.com" to the CSP directive in your virtual host's configuration file.

1. Signup for a free Duo admin account at Duo Signup.
2. Open the Duo Admin Panel, go to *Applications* > *Protect an Application*, search for the *Web SDK* application, and click *Protect*.
3. Note the credentials (Client ID, Client secret, API hostname) required for configuring the MFA method in Zabbix.
4. Go to MFA settings in Zabbix under *Users* > *Authentication* and enable multi-factor authentication.
5. Add a new MFA method with the following configuration:
   - Type: Duo Universal Prompt
   - Name: Zabbix Duo
   - API hostname: (use API hostname from Duo)
   - Client ID: (use Client ID from Duo)
   - Client secret: (use Client secret from Duo)
6. Click the *Add* button and then the *Update* button.
7. Go to *Users* > *User groups* and create a new user group with the following configuration:
   - Group name: Duo group
   - Users: Admin
   - Multi-factor authentication: Default (or "Zabbix Duo" if it is not set as default)
8. Log out of Zabbix and log back in using your credentials. Upon successful login, you will be prompted to enroll in MFA and redirected to Duo. Complete the Duo setup and verify your user with your phone's Duo app to log in.
9. For subsequent logins, use the appropriate MFA method provided by the Duo app (such as retrieving a verification code, responding to push notifications, or using hard keys), and enter the required information during login.

Devices

This section lets Super admins manage mobile devices to be used with Zabbix mobile app.

This section is displayed if `EnableMobileDevices` is enabled.



To add a device:

1. Click *Add device*.
2. Select a user, then click *Add*. The user must have access to devices enabled.
3. A QR code appears. The user must scan this code in their Zabbix mobile app.

Note:
When you add a device for a user, Zabbix also adds the push notification media type to their user media. When you remove the device, that user media is also removed. If you have multiple push notification media types configured, Zabbix adds the one that is enabled and that comes first alphabetically.

Added devices appear in a list, which you can filter.

You can also remove devices. When you remove a device, its access credentials are revoked and removed from Zabbix, after which it can no longer access Zabbix.

The list shows these details:

Column	Description
Name	Device name.
Device ID	Device ID.
User	User that the device was added for, or *Inaccessible user* if this user has been deleted.
User role	User role.
Added on	Date and time when the device was added.
Last active	Date and time when the device was last active.
Status	Device status: **New** - device was added without scanning the QR code, and will need to be added again. **New (Token expired)** - device was added without scanning the QR code, and the token in the QR code has expired; the device will need to be added again. **Active** - device was added, the QR code was scanned, and it may use Zabbix mobile app. **Unassigned** - device was added, but its user was deleted, and it may not use Zabbix mobile app.

Users with enough permissions can also manage their own devices in *User settings* > *Devices*.

Administration

Overview

The Administration menu is for administrative functions of Zabbix. This menu is available to SuperAdmin user type users only.

Data source

Overview

The *Administration* > *Data source* section lets you specify the telemetry data source, that is, the database that holds telemetry data (traces, metrics, and logs). Telemetry data is used for application performance monitoring (APM), which helps you see how requests flow through your applications and services, measure their performance, and find the cause of slow or failed operations.

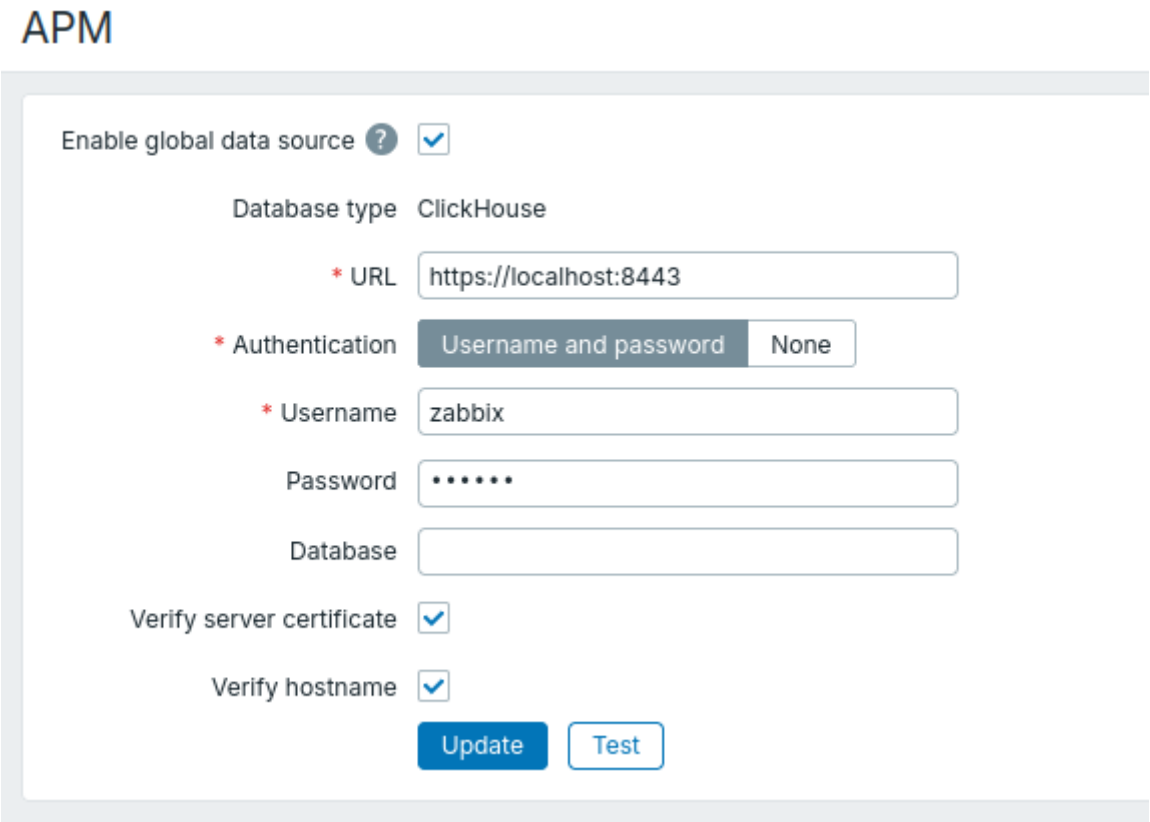
The telemetry data source must support OpenTelemetry.

The telemetry data source specified in this section is called the global data source, because it is shared by Zabbix server, proxy, and frontend. However, you can override it for any of these components by specifying a different telemetry data source in the server, proxy, or frontend configuration.

The telemetry query item queries this data source, unless a different one is specified in the component configuration.

APM

To specify the global data source, mark the *Enable global data source* checkbox and fill in the parameters described below. Before saving, you can click *Test* to check the connection to this data source (requires the php-curl extension).



Parameter	Description
Database type	The database type is hardcoded to *ClickHouse* and cannot be changed. See also: Requirements and ClickHouse setup.
URL	Enter the HTTP or HTTPS URL used to connect to the database.
Authentication	Select the authentication method used to log in to the database: *Username and password* or *None*.
Username	Specify the username used to log in to the database. This parameter is available if *Authentication* is set to *Username and password*.
Password	Specify the password used to log in to the database, or leave empty if the database user has no password. The password is reset when you change the URL. This parameter is available if *Authentication* is set to *Username and password*.
Database	Enter the name of the database that holds the telemetry data. If left empty, Zabbix uses the default database configured in ClickHouse for the user.
Verify server certificate	Mark the checkbox to verify the server certificate. If marked, the related TLS options must be set in the server, proxy, or frontend configuration. This parameter is available if *URL* contains an HTTPS URL.
Verify hostname	Mark the checkbox to verify that the hostname in *URL* matches the hostname in the server certificate. If marked, the related TLS options must be set in the server, proxy, or frontend configuration. This parameter is available if *URL* contains an HTTPS URL.

General

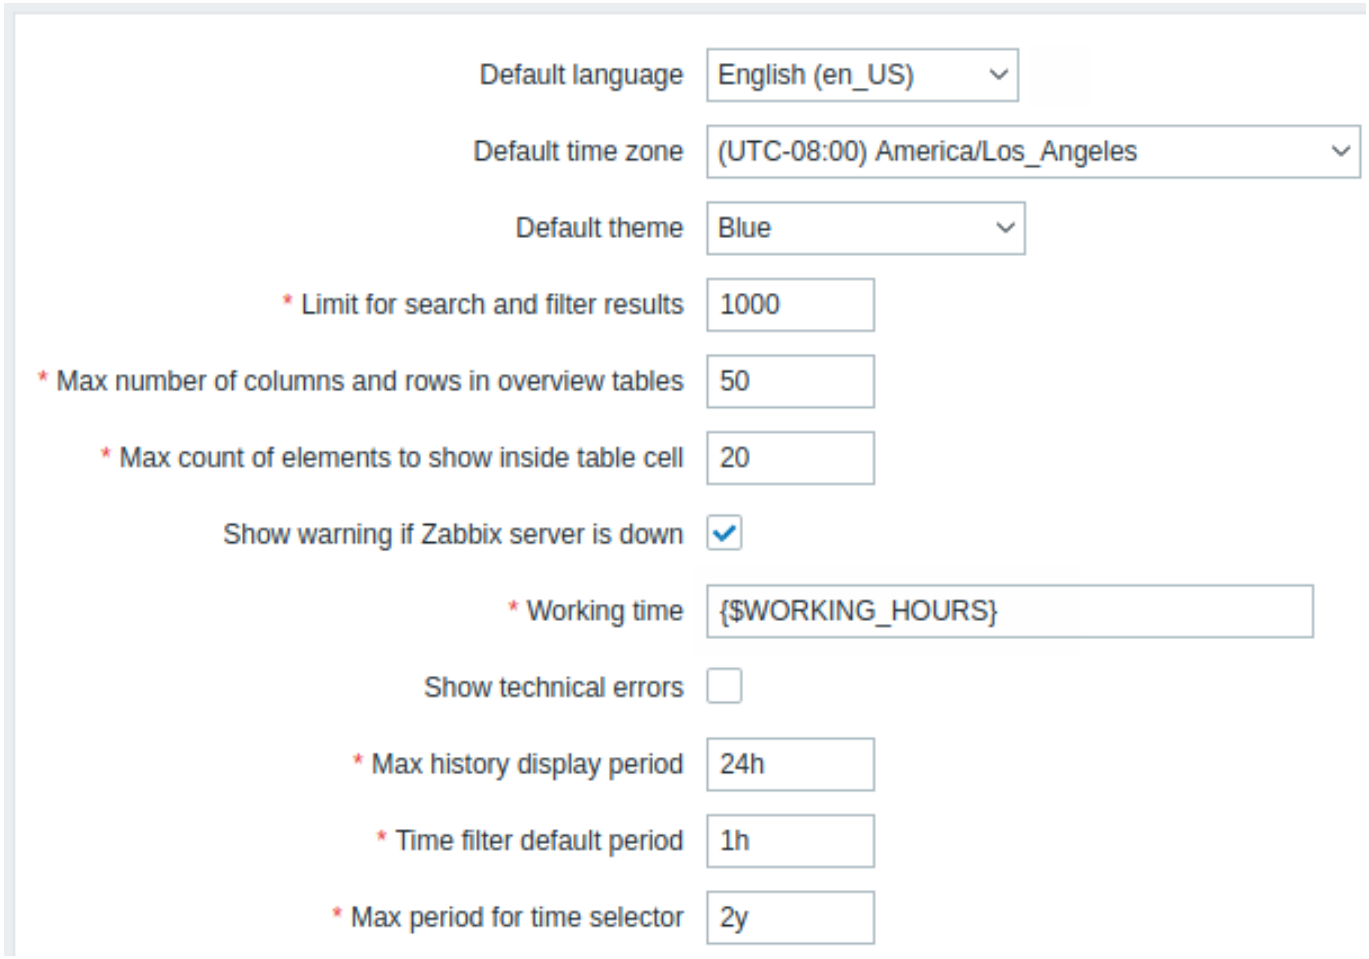
Overview

The *Administration* > *General* section contains a number of subsections for setting frontend-related defaults and customizing Zabbix.

The list of available subsections appears upon clicking *General* in the *Administration* menu section. It is also possible to switch between subsections by using the title dropdown in the upper-left corner.

GUI

This section provides customization of several frontend-related defaults.



Configuration parameters:

Parameter	Description
Default language	Default language for users who have not specified a language in their profiles and guest users. For more information, see Installation of additional frontend languages.
Default time zone	Default time zone for users who have not specified a time zone in their profiles and guest users.
Default theme	Default theme for users who have not specified a theme in their profiles and guest users.
Limit for search and filter results	Maximum amount of elements (rows) that will be displayed in a web-interface list, for example, in *Data collection > Hosts*. *Note*: If set to, for example, '50', only the first 50 elements will be displayed in all affected frontend lists. If some list contains more than fifty elements, the indication of that will be the '+' sign in *"Displaying 1 to 50 of **50+** found"*. Also, if filtering is used and still there are more than 50 matches, only the first 50 will be displayed. Note that increasing the value of this parameter may lead to decreased performance and increased memory consumption on the web server side.
Max number of columns and rows in overview tables	Maximum number of columns and rows to display in the *Trigger overview* dashboard widget. The same limit applies to both columns and rows. If more rows and/or columns than shown exist, the system will display a warning at the bottom of the table: "Not all results are displayed. Please provide more specific search criteria."
Max count of elements to show inside table cell	For entries that are displayed in a single table cell, no more than configured here will be shown.
Show warning if Zabbix server is down	This parameter enables a warning message to be displayed in a browser window if the Zabbix server cannot be reached (possibly down). The message remains visible even if the user scrolls down the page. When hovered over, the message is temporarily hidden to reveal the contents underneath it.
Working time	This system-wide parameter defines working hours. In graphs, working time is displayed as a white background and non-working time is displayed as gray. See Time period specification page for description of the time format. User macros are supported.

Parameter	Description
Show technical errors	If checked, all registered users will be able to see technical errors (PHP/SQL). If unchecked, the information is only available to Zabbix Super Admins and users belonging to the user groups with enabled debug mode.
Max history display period	Maximum time period for which to display historical data in *Monitoring* > *Latest data*, host web scenario information in *Monitoring* > *Hosts*, and in the *Top items* dashboard widget. Allowed range: 24 hours (default) - 1 week. Time suffixes, e.g., 1w (one week), 36h (36 hours), are supported.
Time filter default period	Time period to be used in graphs and dashboards by default. Allowed range: 1 minute - 10 years (default: 1 hour). Time suffixes, e.g., 10m (ten minutes), 5w (five weeks), are supported. Note: when a user changes the time period while viewing a graph, this time period is stored as user preference, replacing the global default or a previous user selection.
Max period for time selector	Maximum available time period for graphs and dashboards. Users will not be able to visualize older data. Allowed range: 1 year - 10 years (default: 2 years). Time suffixes, e.g., 1y (one year), 365w (365 weeks), are supported.

Autoregistration

In this section, you can configure the encryption level for active agent autoregistration.



Parameters marked with an asterisk are mandatory.

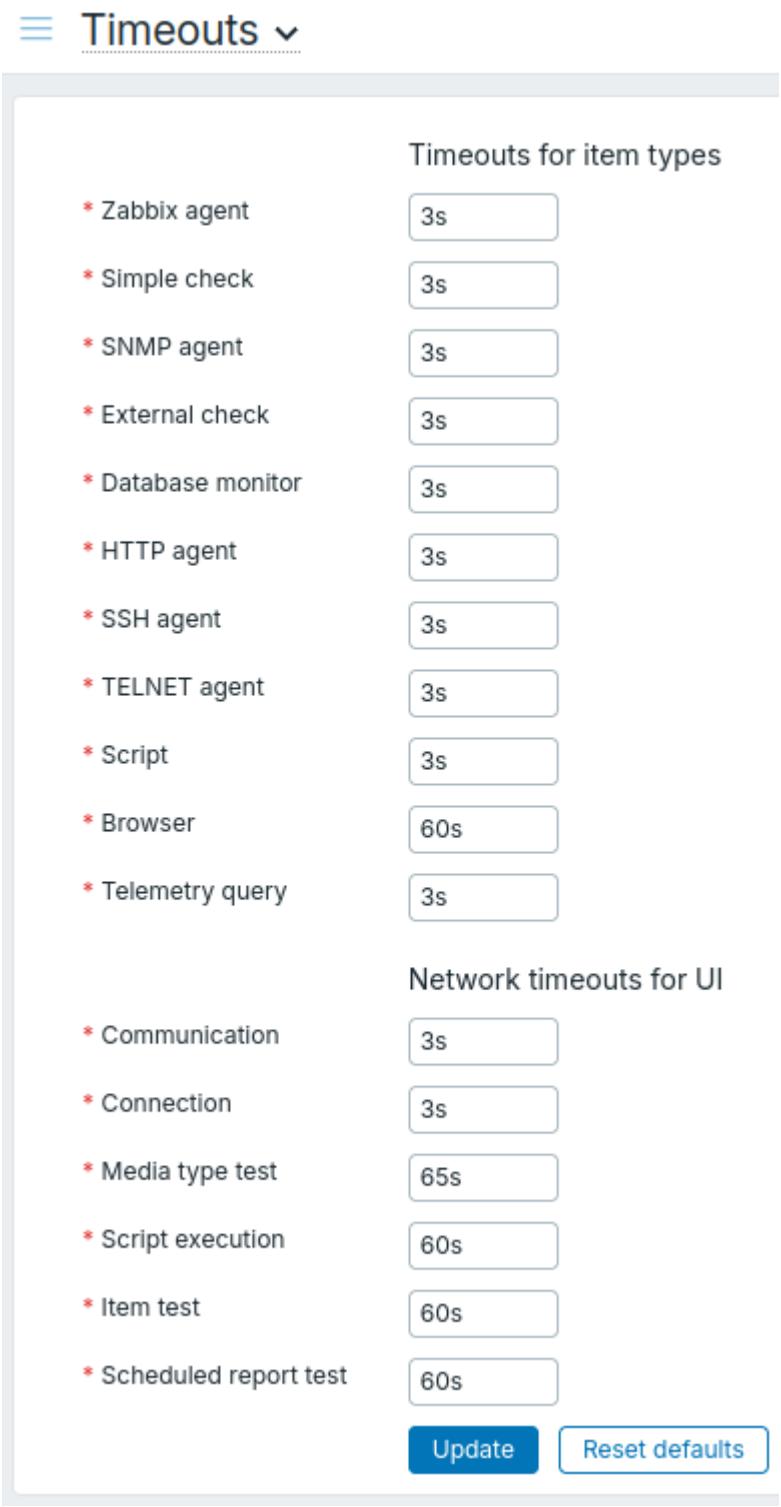
Configuration parameters:

Parameter	Description
Encryption level	Select one or both options for encryption level: **No encryption** - unencrypted connections are allowed **PSK** - TLS encrypted connections with a pre-shared key are allowed
PSK identity	Enter the pre-shared key identity string. This field is only available if 'PSK' is selected as *Encryption level*. Do not put sensitive information in the PSK identity, it is transmitted unencrypted over the network to inform a receiver which PSK to use.
PSK	Enter the pre-shared key (an even number of hexadecimal characters). Maximum length: 512 hex-digits (256-byte PSK) if Zabbix uses GnuTLS or OpenSSL library, 64 hex-digits (32-byte PSK) if Zabbix uses mbed TLS (PolarSSL) library. Example: 1f87b595725ac58dd977beef14b97461a7c1045b9a1c963065002c5473194952 This field is only available if 'PSK' is selected as *Encryption level*.

See also: Secure autoregistration

Timeouts

In this section, it is possible to set global item-type timeouts and network timeouts. All fields in this form are mandatory.

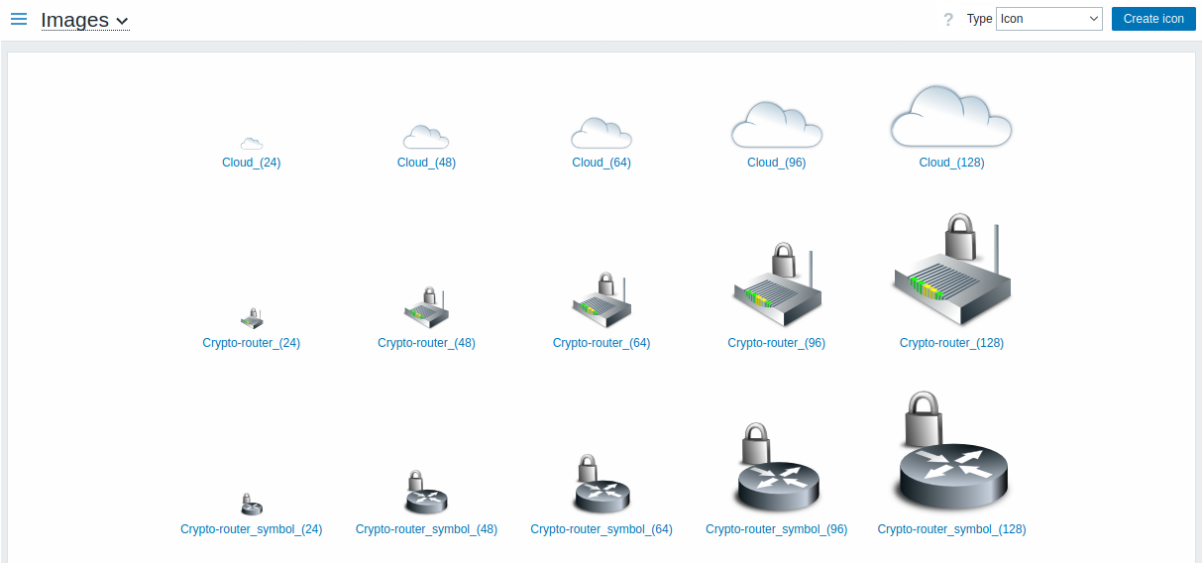


Parameter	Description
Timeouts for item types	How many seconds Zabbix should wait before aborting the check (based on its type). For additional details and limitations, see item timeout. Allowed range: 1 - 600s (default: 3s; default for browser items: 60s). Time suffixes, e.g., 30s, 1m, and user macros are supported. Note that proxy-level timeouts, if configured, override these global settings. Furthermore, if timeouts for specific items are configured, they override both global and proxy settings.
Network timeouts for UI	
Communication	How many seconds to wait before closing an idle socket (if a connection to Zabbix server has been established earlier, but frontend cannot finish data reading/sending operation during this time, the connection will be dropped). Allowed range: 1 - 300s (default: 3s).

Parameter		Description
	Connection	How many seconds to wait before stopping an attempt to connect to Zabbix server. Allowed range: 1 - 30s (default: 3s).
	Media type test	How many seconds to wait for a response when testing a media type. Allowed range: 1 - 300s (default: 65s).
	Script execution	How many seconds to wait for a response when executing a script. Allowed range: 1 - 300s (default: 60s). This timeout is for the entire script execution chain, which can be of various length. For example, if the script is executed on the agent, it is a roundtrip through server (possibly also proxy) to the agent and back. For scripts that are executed on the agent, the server or proxy timeout applies. For scripts that are executed on an active agent only, it is likely that the default server/proxy timeout must be raised. The server/proxy timeout must be higher than the active check refresh frequency; otherwise, the server/proxy timeout may be exceeded before the active agent even receives the script. See also: Script timeout.
	Item test	How many seconds to wait for returned data when testing an item. Allowed range: 1 - 600s (default: 60s).
	Scheduled report test	How many seconds to wait for returned data when testing a scheduled report. Allowed range: 1 - 300s (default: 60s).

Images

The Images section displays all the images available in Zabbix. Images are stored in the database.



The *Type* dropdown allows you to switch between icon and background images:

- Icons are used to display network map elements.
- Backgrounds are used as background images of network maps.

Adding image

You can add your own image by clicking the *Create icon* or *Create background* button in the upper-right corner.

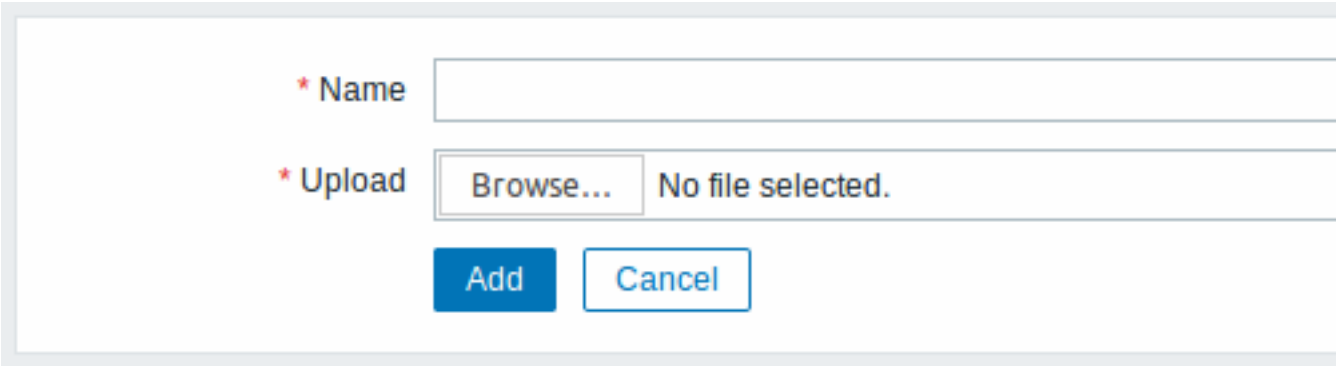


Image attributes:

Parameter	Description
Name	Unique name of an image.
Upload	Select the file (PNG, JPEG, GIF, or WebP) from a local system to be uploaded to Zabbix. *Note* that it may be possible to upload other formats that will be converted to PNG during upload. GD library is used for image processing, therefore formats that are supported depend on the library version used (2.0.28 or higher is required by Zabbix).

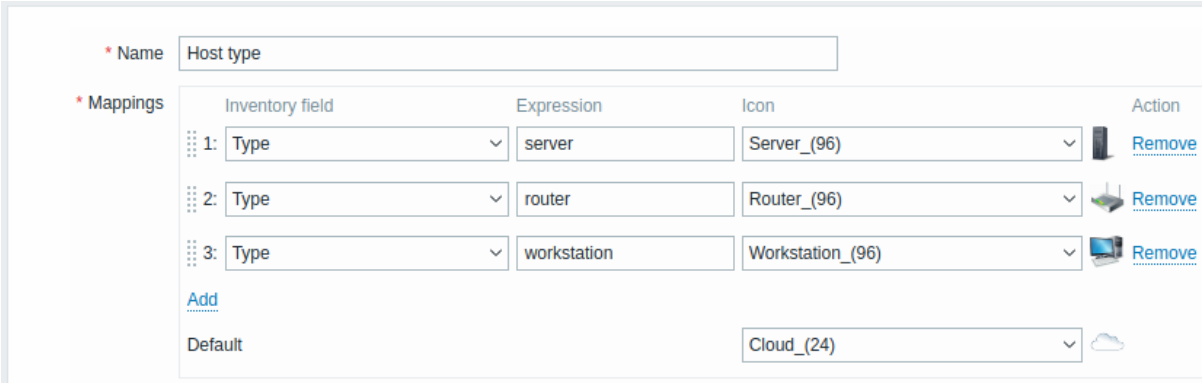
Note:
Maximum size of the upload file is limited by the value of `ZBX_MAX_IMAGE_SIZE` that is 1024x1024 bytes or 1 MB.
The upload of an image may fail if the image size is close to 1 MB and the `max_allowed_packet` MySQL configuration parameter is at a default of 1MB. In this case, increase the `max_allowed_packet` parameter.

Icon mapping

This section allows creating the mapping of certain hosts with certain icons. Host inventory field information is used to create the mapping.

The mappings can then be used in network map configuration to assign appropriate icons to matching hosts automatically.

To create a new icon map, click *Create icon map* in the upper-right corner.



Configuration parameters:

Parameter	Description
Name	Unique name of icon map.
Mappings	A list of mappings. The order of mappings determines which one will have priority. You can move mappings up and down the list with drag-and-drop.
Inventory field	Host inventory field that will be looked into to seek a match.
Expression	Regular expression describing the match.
Icon	Icon to use if a match for the expression is found.
Default	Default icon to use.

Regular expressions

This section allows creating custom regular expressions that can be used in several places in the frontend. See Regular expressions section for details.

Trigger displaying options

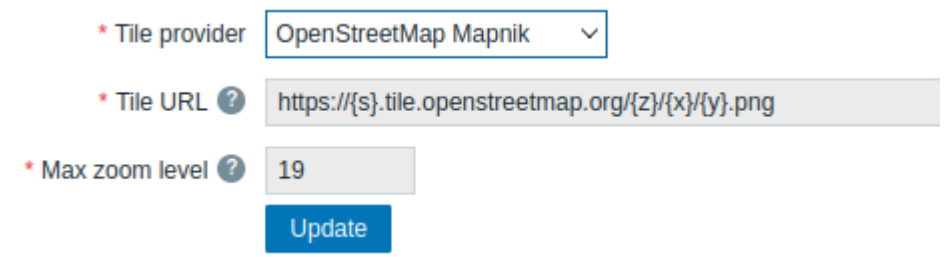
This section allows customizing how trigger state is displayed in the frontend and trigger severity names and colors.



Parameter	Description
Use custom event status colors	Checking this parameter turns on the customization of colors for acknowledged/unacknowledged problems.
Unacknowledged PROBLEM events, *Acknowledged PROBLEM events*, *Unacknowledged RESOLVED events*, *Acknowledged RESOLVED events*	Enter new color code or click the color to select a new one from the provided palette. If *blinking* checkbox is marked, triggers will blink for some time upon the change of state to become more visible.
Display OK triggers for	Time period for displaying OK triggers. Allowed range: 0 - 24 hours. Time suffixes, e.g., 5m, 2h, 1d, are supported.
On status change triggers blink for	Length of trigger blinking. Allowed range: 0 - 24 hours. Time suffixes, e.g., 5m, 2h, 1d, are supported.
Not classified, *Information*, *Warning*, *Average*, *High*, *Disaster*	Custom severity names and/or colors to display instead of system default. Enter new color code or click the color to select a new one from the provided palette. Note that custom severity names entered here will be used in all locales. If you need to translate them to other languages for certain users, see Customizing trigger severities page.

Geographical maps

This section allows selecting geographical map tile service provider and configuring service provider settings for the Geomap dashboard widget. To provide visualization using the geographical maps, Zabbix uses open-source JavaScript interactive maps library Leaflet. Please note that Zabbix has no control over the quality of images provided by third-party tile providers, including the predefined tile providers.



Parameter	Description
Tile provider	Select one of the available tile service providers or select *Other* to add another tile provider or self-hosted tiles (see Using a custom tile service provider).
Tile URL	The URL template (up to 2048 characters) for loading and displaying the tile layer on geographical maps. This field is editable only if *Tile provider* is set to *Other*. The following placeholders are supported: *{s}* represents one of the available subdomains; *{z}* represents zoom level parameter in the URL; *{x}* and *{y}* represent tile coordinates; *{r}* can be used to add "@2x" to the URL to load retina tiles. Example: `https://{s}.example.com/{z}/{x}/{y}{r}.png`
Attribution text	Tile provider attribution text to be displayed in a small text box on the map. This field is visible only if *Tile provider* is set to *Other*.
Max zoom level	Maximum zoom level of the map. This field is editable only if *Tile provider* is set to *Other*.

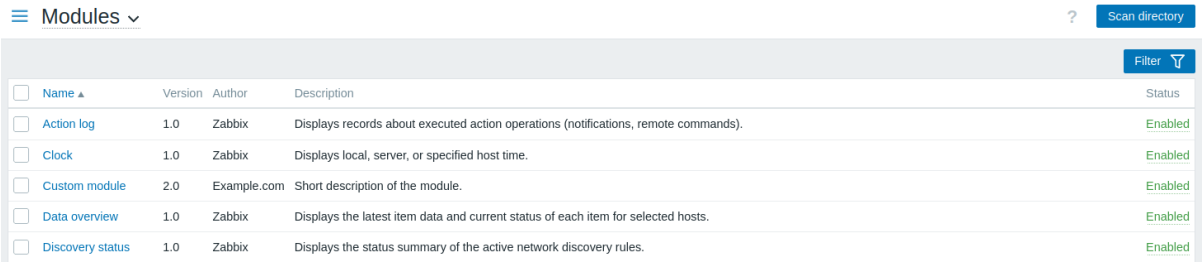
Using a custom tile service provider

The Geomap widget is capable to load raster tile images from a custom self-hosted or a third-party tile provider service. To use a custom third-party tile provider service or a self-hosted tile folder or server, select *Other* in the *Tile provider* field and specify the custom URL in the *Tile URL* field using proper placeholders.

Modules

This section allows you to administer custom, as well as built-in frontend modules.

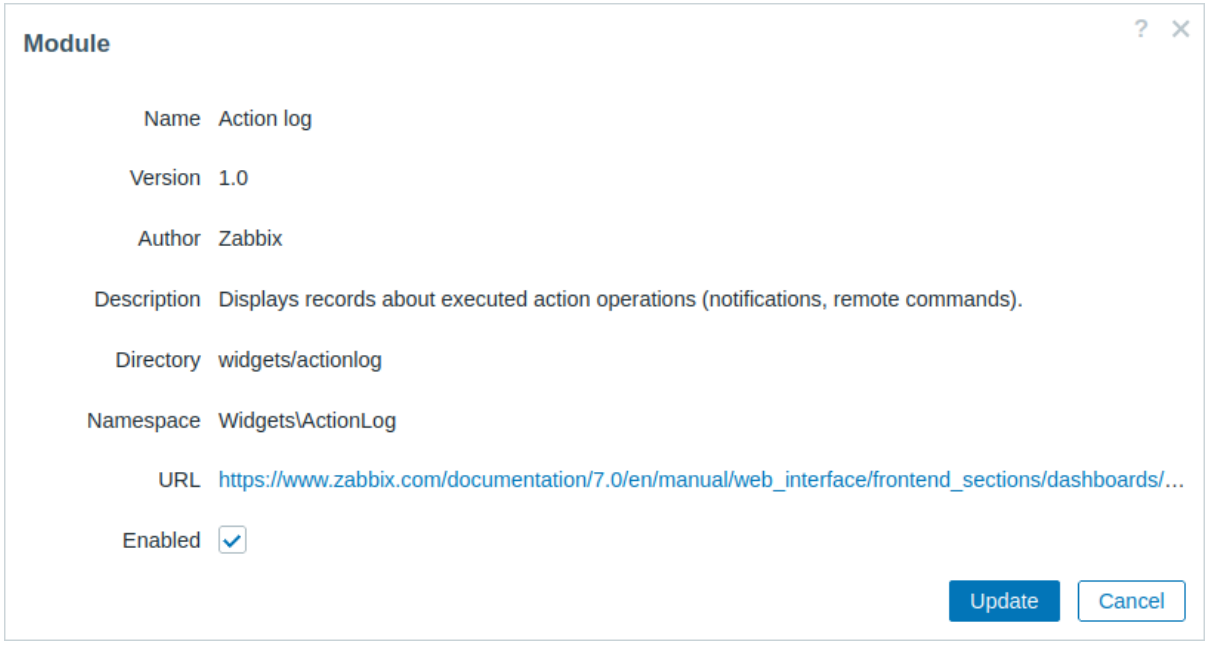
Note:
If the `$ZBX_FEATURE_FLAGS['modules_config_enabled']` feature flag is set to `false` in the frontend configuration file, access to this section is disabled.



Click *Scan directory* to register/unregister any custom modules. Registered modules will appear in the list; unregistered modules will be removed from the list.

Click the module status in the list to enable/disable a module. You may also mass enable/disable modules by selecting them in the list and then clicking the *Enable/Disable* buttons below the list.

Click the module name in the list to view its details in a pop-up window.

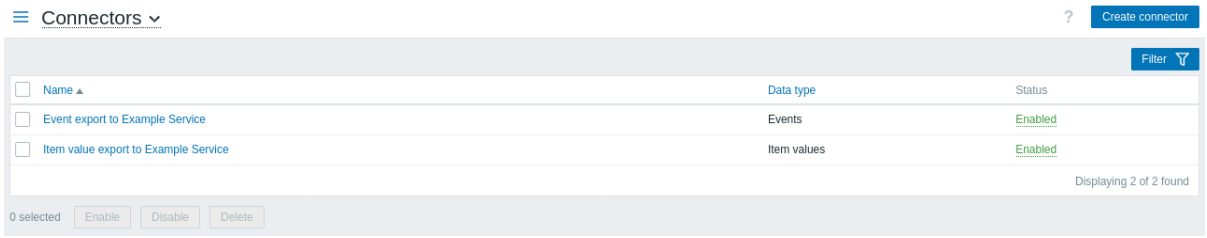


Module status can also be updated in the module details pop-up window; to do this, mark/unmark the *Enabled* checkbox and then click *Update*.

You may filter modules by name or status (enabled/disabled).

Connectors

This section allows you to configure connectors for Zabbix data streaming to external systems over HTTP.



Click *Create connector* to configure a new connector.

You may filter connectors by name or status (enabled/disabled). Click the connector status in the list to enable/disable a connector. You may also mass enable/disable connectors by selecting them in the list and then clicking the *Enable/Disable* buttons below the list.

Other

This section allows configuring miscellaneous other frontend parameters.



Parameter	Description
Frontend URL	URL (up to 2048 characters) to Zabbix web interface. This parameter is used by Zabbix web service for communication with frontend and should be specified to enable scheduled reports.
Group for discovered hosts	Hosts discovered by network discovery and agent autoregistration will be automatically placed in the host group, selected here.
Default host inventory mode	Default mode for host inventory. It will be followed whenever a new host or host prototype is created by server or frontend unless overridden during host discovery/autoregistration by the *Set host inventory mode* operation.

Parameter	Description
User group for database down message	User group for sending the database down alarm message. Zabbix server depends on the availability of the backend database. It cannot work without a database. If the database is down, selected users can be notified by Zabbix. Notifications will be sent to the user group set here using the enabled user media entries. Notifications are transmitted using the following media types (when enabled): email; SMS; custom alert scripts. Even if a webhook media entry is configured and enabled, it will not trigger notifications. Zabbix server will not stop; it will wait until the database is back again to continue processing. Notification consists of the following content: `[MySQL|PostgreSQL] database <DB Name> [on <DB Host>:<DB Port>] is not available: <error message depending on the type of DBMS (database)>` <DB Host> is not added to the message if it is defined as an empty value and <DB Port> is not added if it is the default value ("0"). The alert manager (a special Zabbix server process) tries to establish a new connection to the database every 10 seconds. If the database is still down the alert manager repeats sending alerts, but not more often than every 15 minutes.
Log unmatched SNMP traps	Log SNMP trap if no corresponding SNMP interfaces have been found.

Authorization

Parameter	Description
Login attempts	Number of unsuccessful login attempts before the possibility to log in gets blocked.
Login blocking interval	Period of time for which logging in will be prohibited when *Login attempts* limit is exceeded. Allowed range: 30 - 3600 seconds. Time suffixes, e.g., 90s, 5m, 1h, are supported.

Storage of secrets

Parameter	Description
Vault provider	Select the secret management software for storing user macro values - *HashiCorp Vault* (default) or *CyberArk Vault*.
Resolve secret vault macros by	Resolve secret vault macros by: **Zabbix server** - secrets are retrieved by Zabbix server and forwarded to proxies when needed (default); **Zabbix server and proxy** - secrets are retrieved by both Zabbix server and proxies, allowing them to resolve macros independently.

See also: Storage of secrets.

Security

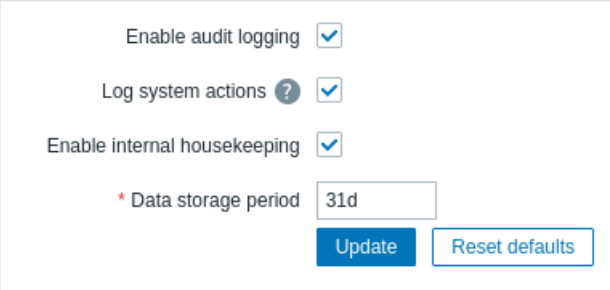
Parameter	Description
Validate URI schemes	Unmark this checkbox to disable URI scheme validation (enabled by default). If marked, you can specify a comma-separated list of allowed URI schemes (default: http,https,ftp,file,mailto,tel,ssh). Applies to all fields in the frontend where URIs are used (for example, map element URLs).
Use X-Frame-Options HTTP header	Unmark this checkbox to disable the HTTP X-Frame-options header (not recommended). If marked, you can specify the value of the HTTP X-Frame-options header. Supported values: **SAMEORIGIN** (default) or **'self'** (must be single-quoted) - the page can only be displayed in a frame on the same origin as the page itself; **DENY** or **'none'** (must be single-quoted) - the page cannot be displayed in a frame, regardless of the site attempting to do so; **a string of space-separated hostnames**; adding **'self'** (must be single-quoted) to the list allows the page to be displayed in a frame on the same origin as the page itself. Note that using **'self'** or **'none'** without single quotes will result in them being regarded as hostnames.

Parameter	Description
Use iframe sandboxing	Unmark this checkbox to disable putting the retrieved URL content into sandbox (not recommended). If marked, you can specify the iframe sandboxing exceptions; unspecified restrictions will still be applied. If this field is empty, all sandbox attribute restrictions apply. For more information, see the description of the `sandbox` attribute.

Audit log

Overview

This section allows configuring audit log settings.



The following parameters are available:

Parameter	Description
Enable audit logging	Enable (default) or disable audit logging.
Log system actions	Enable (default) or disable audit logging of low-level discovery, network discovery and autoregistration activities performed by the server (System user).
Enable internal housekeeping	Enable (default) or disable internal housekeeping for audit log records.
Data storage period	Amount of days audit log records should be kept before being removed by the housekeeper. Mandatory if housekeeping is enabled. Default: 31 days.

The audit log compression parameters are available in the *Administration > Housekeeping* section, within the *History, trends and audit log compression* block, which becomes visible when TimescaleDB is used.

Housekeeping

Overview

The *Administration > Housekeeping* section lets you configure Zabbix internal housekeeping that automatically removes outdated or user-deleted data from the database. This prevents the database from growing indefinitely, which would consume unnecessary disk space and degrade database performance.

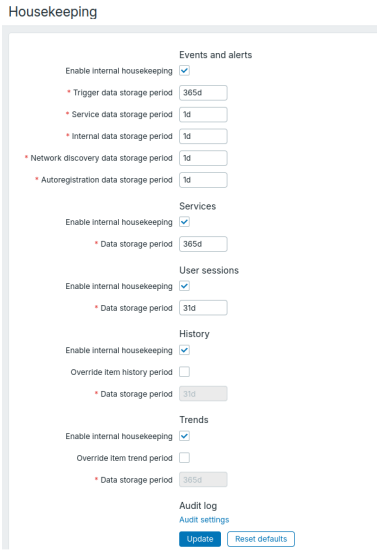
Housekeeping can be enabled and configured for the following types of data:

- Events and alerts from triggers, services, internal data, network discovery, and autoregistration
- Services
- User sessions
- History and trends for items

Note:
Audit log housekeeping is configured separately.

Configuration

The following parameters are available:



Parameter	Description
Enable internal housekeeping	Enable or disable internal housekeeping (enabled by default). If enabled, the housekeeping procedure automatically removes data from the database that exceeds the *Data storage period*.
Data storage period	Specify how long to keep data before it is removed by the housekeeper. Mandatory if internal housekeeping is enabled. Range: 1 day (1 hour for history) - 25 years; or "0". Time suffixes (e.g., 1d, 1w) are supported. For *Events and alerts*, the data storage period is set separately for triggers, services, internal data, network discovery, and autoregistration. The housekeeper removes only those events that are not associated with problems. For example, a problem/recovery event older than *Data storage period* will not be removed if it's associated with an unresolved problem. When the housekeeper removes outdated entities, it first removes problems, then events. For *History* and *Trends*, the data storage periods also determine how long data remains visible in *Monitoring* > *Latest data*, even if internal housekeeping is disabled. Note that when an item or trigger is deleted, its problems and associated events are also deleted.
Override item history period	If enabled, the history/trend storage period specified in the item configuration is overridden by the *Data storage period* setting (except for items with the *Do not store* option enabled). This option can be used even when internal housekeeping is disabled and an external housekeeper is used.
Override item trend period	

If using **ClickHouse**, the *ClickHouse data storage period* section becomes available. It shows the data storage period for each item value type. To change the period, see Configuring ClickHouse data storage period. Note that the housekeeper does not delete data from ClickHouse.

If using **Elasticsearch**, the *Elasticsearch data storage period* section becomes available. It shows the item value types saved in Elasticsearch, with the data storage period displayed as *Unknown*. To control how long data is kept, see Index Lifecycle Management. Note that the housekeeper does not delete data from Elasticsearch.

If using **TimescaleDB**, the *History, trends and audit log compression* section becomes available.

Note:
For TimescaleDB, enable *Override item history period*, *Override item trend period*, and *Enable internal housekeeping* for history and trends to fully benefit from automatic partitioning. If these options are disabled, data kept in history and trends tables will still be partitioned, but the housekeeper will not drop outdated partitions, and configuration warnings will appear. When dropping of outdated partitions is enabled, Zabbix server and frontend will no longer track deleted items, and history for those items will be cleared when an outdated partition is deleted.

The housekeeper is executed automatically every hour by default. You can change the execution frequency by setting the `HousekeepingFrequency` Zabbix server configuration parameter.

The *Reset defaults* button lets you revert any changes made.

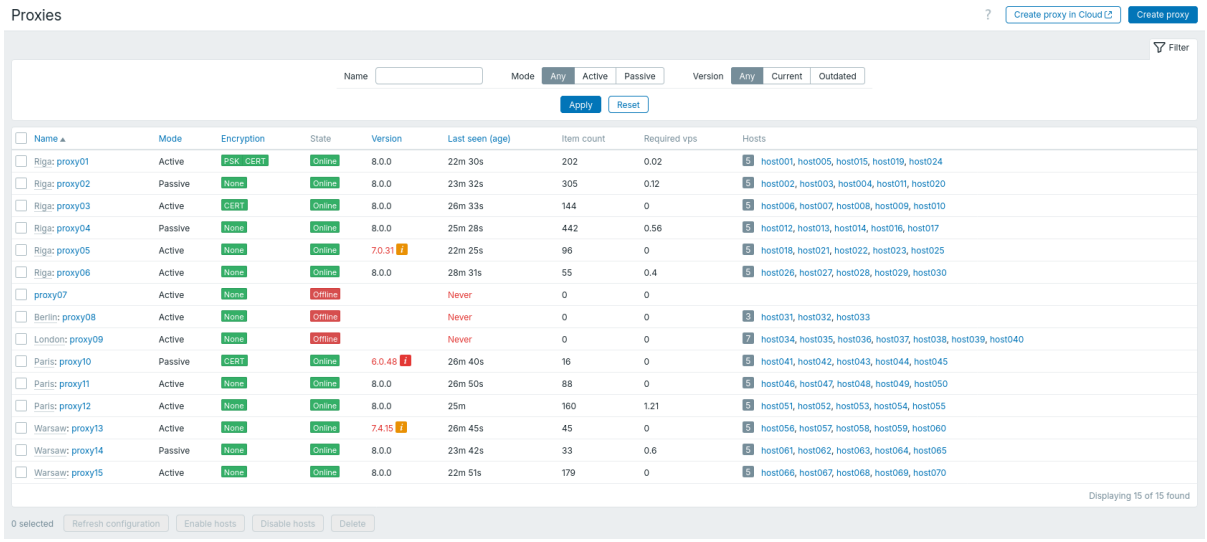
Proxies

Overview

In the *Administration* > *Proxies* section proxies for distributed monitoring can be configured in the Zabbix frontend.

Proxies

A listing of existing proxies with their details is displayed.



Displayed data:

Column	Description
Name	Name of the proxy. Clicking the proxy name opens the proxy configuration form. If the proxy belongs to a proxy group, the group name is displayed before the proxy name as a gray link. Clicking the group name opens the proxy group configuration form.
Mode	Proxy mode - *Active* or *Passive*.
Encryption	Encryption status for connections from the proxy: **None** - no encryption; **PSK** - using pre-shared key; **Cert** - using certificate.
State	State of the proxy: **Unknown** - proxy was created while Zabbix server was down, or server has not yet updated the state; **Online** - proxy has communicated with Zabbix server (passive proxy responded to server request; active proxy sent request) within the failover period; **Offline** - proxy has not communicated with Zabbix server within the failover period.
Version	Proxy version (three digit version number). If proxy is outdated or unsupported, version number is highlighted (red) and info status icon (yellow or red) is displayed. Hover over the icon for details.
Last seen (age)	The time when the proxy was last seen by the server.
Item count	The number of enabled items on enabled hosts assigned to the proxy.
Required vps	Required proxy performance (the number of values that need to be collected per second).
Hosts	Count of enabled hosts assigned to the proxy and a list of hosts monitored by the proxy. Clicking the host name opens the host configuration form.

To configure a new proxy, click the *Create proxy* button in the upper-right corner. To deploy a new proxy in Zabbix Cloud, click the *Create proxy in Cloud* button next to it. This opens Zabbix Cloud in a new browser tab, where you can create a proxy node. Once the proxy node is created, configure the proxy using the *Create proxy* button. For more information, see Set up a Zabbix Cloud proxy.

Mass editing options

Buttons below the list offer some mass-editing options:

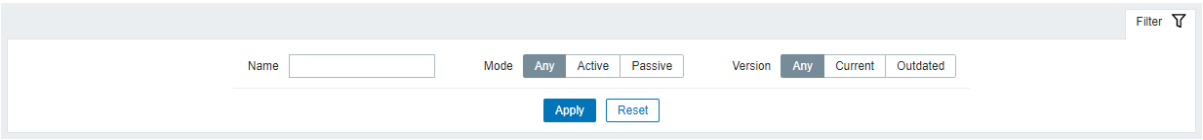
- *Refresh configuration* - refresh configuration of the proxies;
- *Enable hosts* - change the status of hosts monitored by the proxy to *Monitored*;
- *Disable hosts* - change the status of hosts monitored by the proxy to *Not monitored*;
- *Delete* - delete the proxies.

To use these options, mark the checkboxes before the respective proxies, then click the required button.

Using filter

You can use the filter to display only the proxies you are interested in. For better search performance, data is searched with macros unresolved.

The *Filter* link is available above the list of proxies. If you click it, a filter becomes available where you can filter proxies by name, mode and version. Note that the filter option *Outdated* displays both outdated (partially supported) and unsupported proxies.



Proxy groups

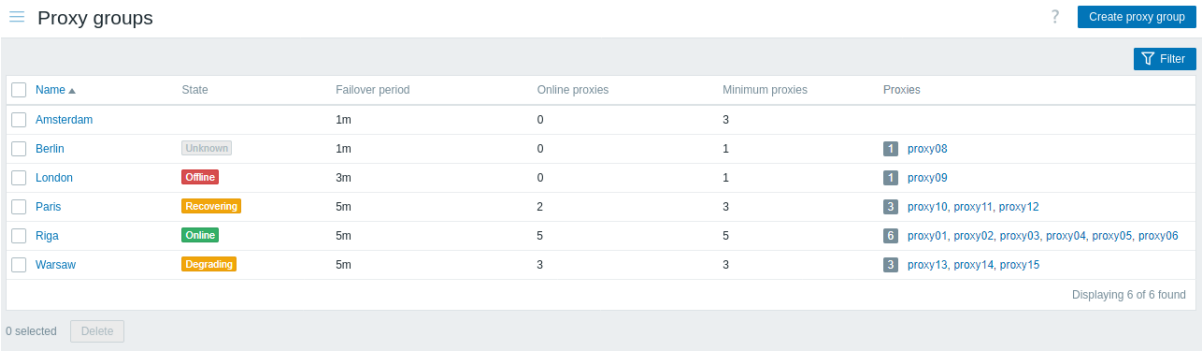
Overview

In *Administration* > *Proxy groups* it is possible to configure proxy groups.

Proxy groups are used in proxy load balancing with automated distribution of hosts between proxies and high availability between proxies.

Proxy groups

A listing of existing proxy groups with their details is displayed.



Displayed data:

Column	Description
Name	Name of the proxy group. Clicking the proxy group name opens the proxy group configuration form.
State	State of the proxy group (displayed if group contains at least one proxy): **Unknown** - group was created while Zabbix server was down, or server has not yet updated the state; **Online** - minimum number of proxies have communicated with Zabbix server within the failover period; **Degrading** - group is going offline shortly due to online proxies falling below the *Minimum proxies* threshold; **Offline** - fewer than minimum number of proxies have communicated with Zabbix server within the failover period; **Recovering** - group is coming back online shortly due to online proxies reaching the *Minimum proxies* threshold. Proxy group state can be monitored with zabbix[proxy group,<name>,state] and zabbix[proxy group,discovery] items.

Column	Description
Failover period	Period during which a proxy in the proxy group must communicate with Zabbix server to be considered online.
Online proxies	Number of online proxies (displayed in red if below the group minimum).
Minimum proxies	Minimum number of online proxies required to keep the proxy group online.
Proxies	Count of proxies in the group and a list of proxies in the group, with links to the proxy configuration form. The maximum number of listed proxies is limited by the *Max count of elements to show inside table cell* value.

To configure a new proxy group, click the *Create proxy groups* button in the upper-right corner.

Mass editing options

Buttons below the list offer some mass-editing options:

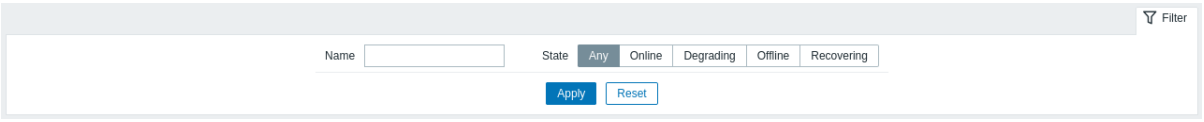
- *Delete* - delete the proxy groups.

To use these options, mark the checkboxes before the respective proxy groups, then click the required button.

Using filter

You can use the filter to display only the proxy groups you are interested in. For better search performance, data is searched with macros unresolved.

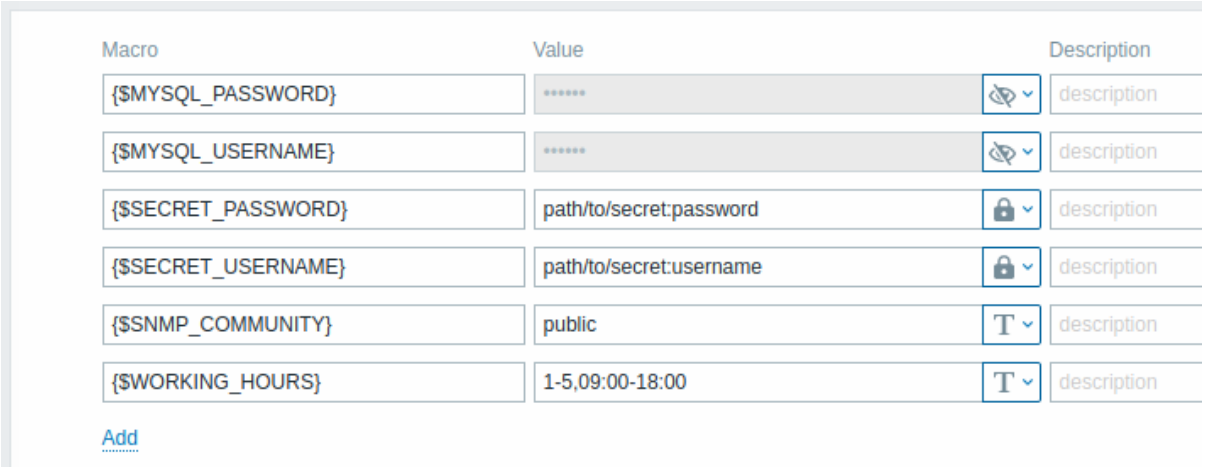
The *Filter* link is available above the list of proxy groups. If you click it, a filter becomes available where you can filter proxy groups by name and status.



Macros

Overview

This section allows you to define system-wide user macros as name-value pairs. Note that macro values can be kept as plain text, secret text or Vault secret. Adding a description is also supported.



Queue

Overview

In the *Administration* > *Queue* section items that are waiting to be updated are displayed.

Ideally, when you open this section it should all be ”green” meaning no items in the queue. If all items are updated without delay, there are none waiting. However, due to lacking server performance, some items may get delayed and the information is displayed in this section. For more details, see the Queue section.

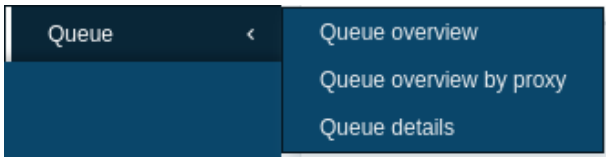
Note:
The Zabbix proxy is not aware of maintenance periods; see Calculation of queues during maintenance for details.

Note:
The queue is available only if Zabbix server is running. Items are not counted in the queue if the item interface becomes unavailable due to connection problems or agent not working properly.

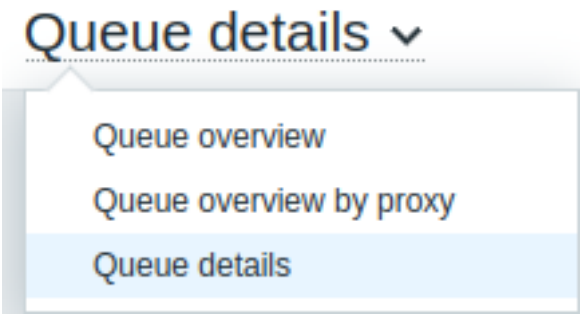
The *Administration* > *Queue* section contains the following pages:

- Queue overview - displays queue by item type;
- Queue overview by proxy - displays queue by proxy;
- Queue details - displays a list of delayed items.

The list of available pages appears upon clicking *Queue* in the *Administration* menu section. It is also possible to switch between pages by using a title dropdown in the upper-left corner.



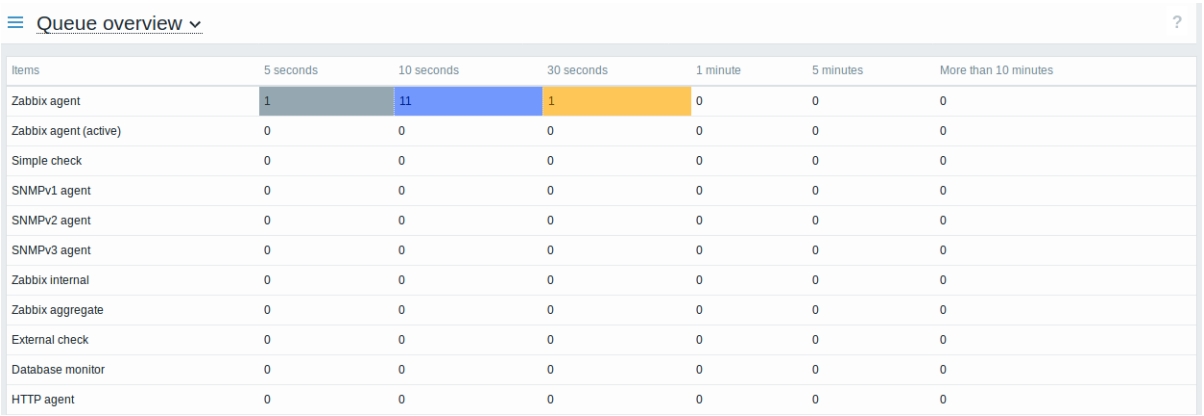
Third-level menu.



Title dropdown.

Overview by item type

In this screen it is easy to locate if the problem is related to one or several item types.

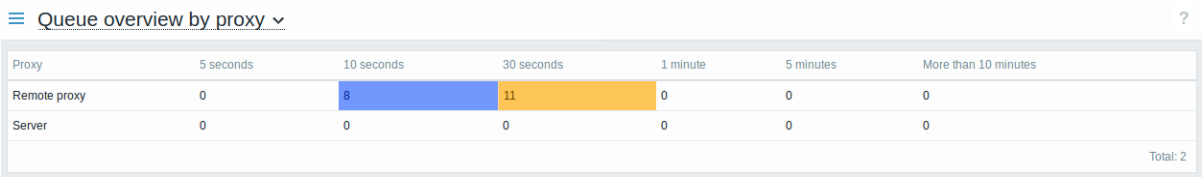

Queue overview

Items	5 seconds	10 seconds	30 seconds	1 minute	5 minutes	More than 10 minutes
Zabbix agent	1	11	1	0	0	0
Zabbix agent (active)	0	0	0	0	0	0
Simple check	0	0	0	0	0	0
SNMPv1 agent	0	0	0	0	0	0
SNMPv2 agent	0	0	0	0	0	0
SNMPv3 agent	0	0	0	0	0	0
Zabbix internal	0	0	0	0	0	0
Zabbix aggregate	0	0	0	0	0	0
External check	0	0	0	0	0	0
Database monitor	0	0	0	0	0	0
HTTP agent	0	0	0	0	0	0

Each line contains an item type. Each column shows the number of waiting items - waiting for 5-10 seconds/10-30 seconds/30-60 seconds/1-5 minutes/5-10 minutes or over 10 minutes respectively.

Overview by proxy

In this screen it is easy to locate if the problem is related to one of the proxies or the server.


Queue overview by proxy

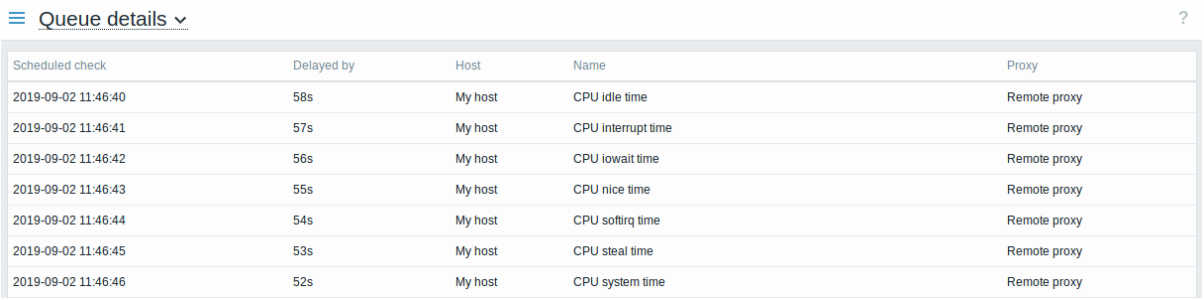
Proxy	5 seconds	10 seconds	30 seconds	1 minute	5 minutes	More than 10 minutes
Remote proxy	0	8	11	0	0	0
Server	0	0	0	0	0	0

Total: 2

Each line contains a proxy, with the server last in the list. Each column shows the number of waiting items - waiting for 5-10 seconds/10-30 seconds/30-60 seconds/1-5 minutes/5-10 minutes or over 10 minutes respectively.

List of waiting items

In this screen, each waiting item is listed.



Displayed data:

Column	Description
Scheduled check	The time when the check was due is displayed.
Delayed by	The length of the delay is displayed.
Host	Host of the item is displayed.
Name	Name of the waiting item is displayed.
Proxy	The proxy name is displayed, if the host is monitored by proxy.

Possible error messages

You may encounter a situation when no data is displayed and the following error message appears:



Error message in this case is the following:

```
Cannot display item queue. Permission denied
```

This happens when the PHP configuration parameters in the frontend configuration file file - `$ZBX_SERVER` or both `$ZBX_SERVER` and `$ZBX_SERVER_PORT` - point to an existing Zabbix server that uses a different database.

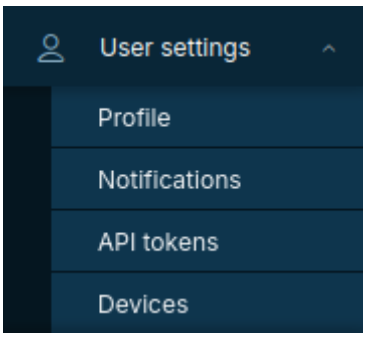
User settings

Overview

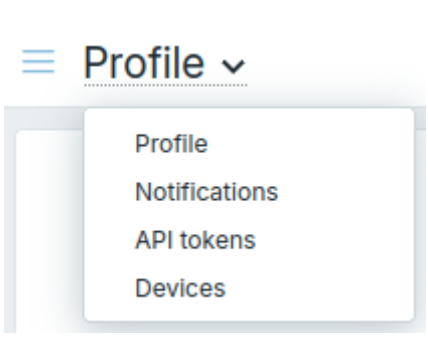
Depending on your user role, the *User settings* section may contain the following pages:

- *Profile* or *User profile* - for customizing certain Zabbix frontend features;
- *Notifications* - for customizing current user notifications;
- *API tokens* - for managing API tokens assigned to the current user;
- *Devices* - for managing mobile devices belonging to the current user.

The list of available pages appears upon clicking the user icon near the bottom of the Zabbix menu (not available for the *guest* user). It is also possible to switch between pages by using the title dropdown in the upper-left corner.



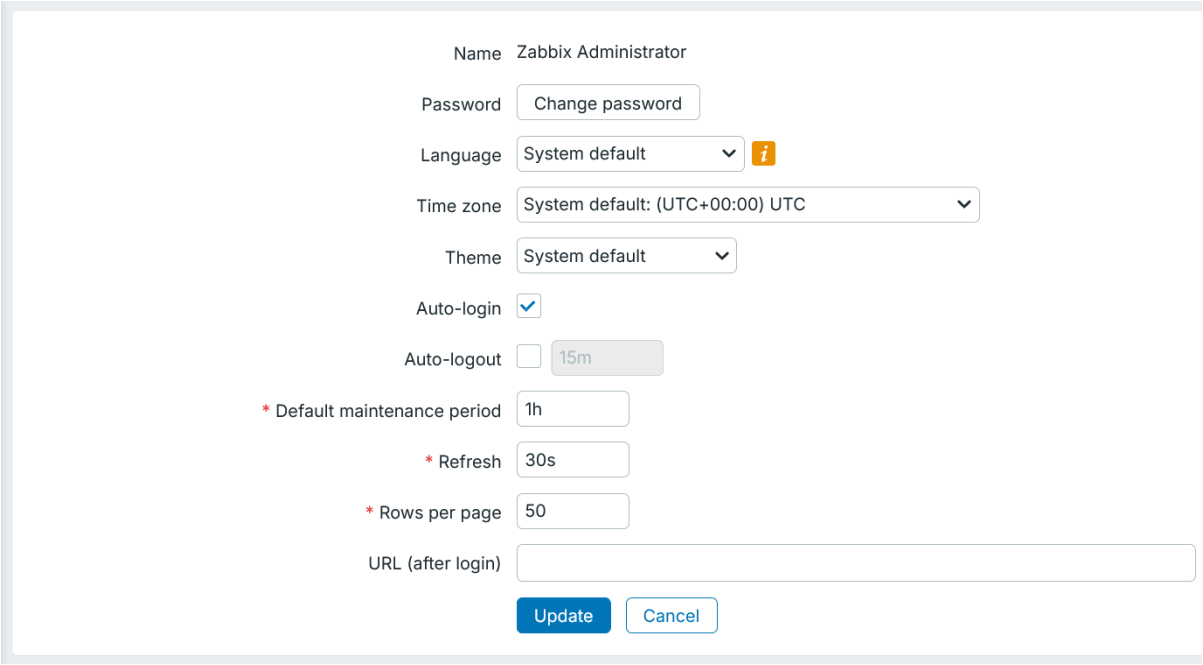
Third-level menu.



Title dropdown.

User profile

The *User profile* section provides options to set a custom interface language, color theme, number of rows displayed in the lists, etc. The changes made here will be applied to the current user only.



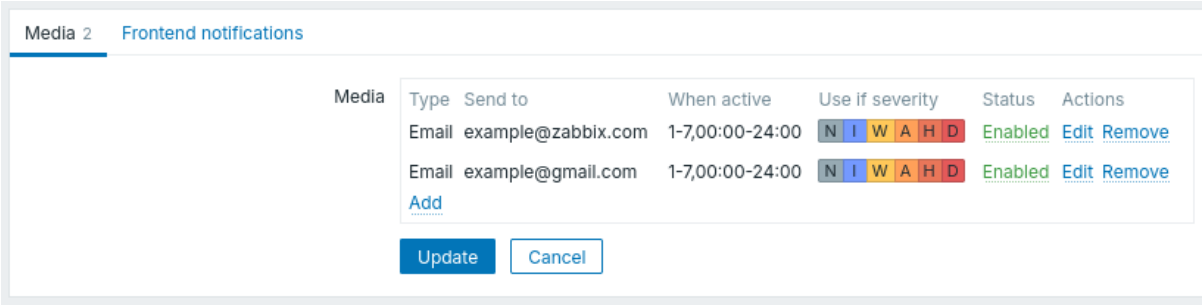
Parameter	Description
Name	User name and surname are displayed. In the absence of name and surname, the username is displayed.
Password	Click the *Change password* button to open three fields: *Old password*, *New password*, *New password (once again)*. On a successful password change, the user will be logged out of all active sessions. Note that the password can only be changed for users using Zabbix internal authentication.
Language	Select the interface language of your choice or select **System default** to use default system settings. Selecting *English (en_US)* will also enable the US time/date format in the frontend. For more information, see Installation of additional frontend languages.
Time zone	Select the time zone to override the global time zone on the user level or select **System default** to use global time zone settings.
Theme	Select a color theme specifically for your profile: **System default** - use default system settings; **Blue** - standard blue theme; **Blue (classic)** - standard blue theme using legacy fonts; **Dark** - alternative dark theme; **Dark blue** - dark blue theme; **Dark (classic)** - alternative dark theme using legacy fonts; **High-contrast light** - light theme with high contrast; **High-contrast dark** - dark theme with high contrast.
Auto-login	Mark this checkbox to make Zabbix remember you and log you in automatically for 30 days. When logging in with *Remember for 30 days* enabled: - Auto-logout timer is reset (session persists until the 30-day period ends). - Auto-login is enabled on subsequent visits within 30 days. When logging in without *Remember for 30 days*: - Auto-login parameter is cleared. - Auto-logout follows the configured timeout as usual. Browser cookies are used for this.

Parameter	Description
Auto-logout	With this checkbox marked, you will be logged out automatically after the set amount of seconds (minimum 90 seconds, maximum 1 day). Time suffixes are supported, for example: 90s, 5m, 2h, 1d. Note that this option will not work in the following cases: * When the *Monitoring* menu pages perform background information refreshes. In case pages that are refreshing data in a specific time interval (dashboards, graphs, latest data, etc.) are left open, the session lifetime is extended, respectively disabling the auto-logout feature. * If logging in with the *Remember me for 30 days* option checked; in this case the auto-logout timer is overridden, keeping your session active for the full period. * When authentication is performed via an SSO Identity Provider (IdP), auto-logout ends only the Zabbix session. The SSO session with the IdP remains active, so the user may be able to log back in without entering a password. *Auto-logout* can also accept "0", meaning that the auto-logout feature becomes disabled after profile settings update.
Default maintenance period	Set the default maintenance period for ad-hoc maintenance.
Refresh	Set how often the information on the *Monitoring* menu pages will be refreshed (minimum 0 seconds, maximum 1 hour). Time suffixes are supported, for example: 30s, 90s, 1m, 1h.
Rows per page	Set how many rows will be displayed per page in the lists. Fewer rows (and fewer records to display) result in faster loading times.
URL (after login)	Set a specific URL to be displayed after login. Instead of the default *Dashboards*, it can be, for example, the URL of *Monitoring* > *Triggers*. This field is disabled, if the URL after login is restricted to a specific page for your user role by a *Super admin* user. If the field is disabled or the user does not specify any value, the default is displayed below the field, e.g. `zabbix.php?action=dashboard.view`

Notifications

The *Notifications* section provides options for customizing current user notifications.

The *Media* tab allows you to specify media details for the user, such as media types and addresses to use and when to use them to deliver notifications.



If the media type has been disabled:

- a yellow info icon is displayed after the name;
- "Disabled" is displayed in the *Status* column.

Note:
User permissions to change media details for themselves can be granted/revoked based on their user role (see *Create and edit own media* option). Super admin user permissions to change media details for others can also be granted/revoked based on their user role (see *Create and edit user media* option)

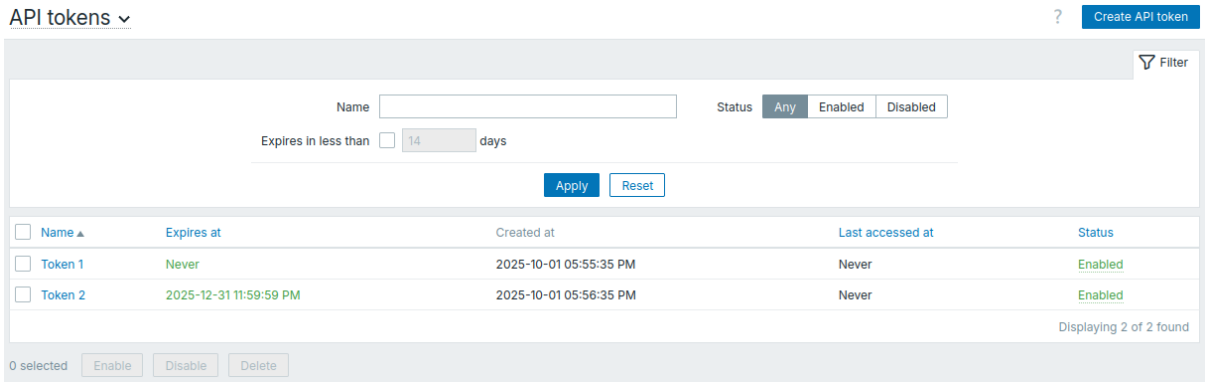
Note that for provisioned users:

- provisioned user media cannot be deleted;
- provisioned user media can be disabled/enabled;
- such provisioned user media fields as *When active*, *Use if severity* and *Enabled* can be manually edited;
- additional user media can be added (for example, additional email addresses) for provisioned users manually;
- manually added user media can be deleted.

The *Frontend notifications* tab allows you to set global notifications.

API tokens

The *API tokens* section allows you to view tokens assigned to the user, edit token details and create new tokens. This section is available to a user only if the *Manage API tokens* action is allowed in the user role settings.



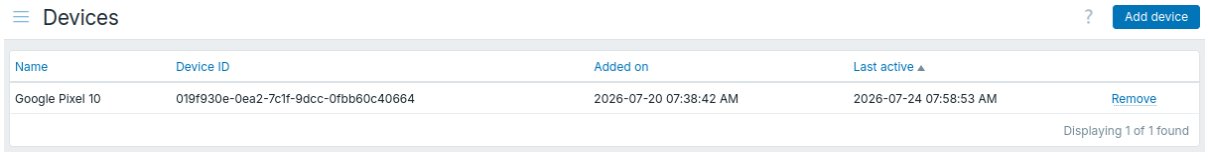
You can filter API tokens by name, expiry date, or status (*Enabled/Disabled*). Click the token status in the list to quickly enable/disable a token. You can also enable/disable multiple tokens at once by selecting them in the list and then clicking the *Enable/Disable* buttons below the list.

Attention:
Users cannot view the *Auth token* value of the tokens assigned to them in Zabbix. The *Auth token* value is displayed only once - immediately after creating a token. If the token has been lost, it has to be regenerated.

Devices

The *Devices* section lets you add or remove your mobile devices to be used with Zabbix mobile app.

This section is displayed if `EnableMobileDevices` is enabled.



To add a device:

1. Click *Add device*.
2. Scan the QR code in your Zabbix mobile app.

Note:
When you add a device, Zabbix also adds the push notification media type to your user media. When you remove the device, that user media is also removed. If you have multiple push notification media types configured, Zabbix adds the one that is enabled and that comes first alphabetically.

Added devices appear in a list. It shows their name, ID, and the dates and times when they were added and last active.

You can also remove devices. When you remove a device, its access credentials are revoked and removed from Zabbix, after which it can no longer access Zabbix.

Super admins can also manage other user devices in *Users* > *Devices*.

Global notifications

Overview

Global notifications provide a way to display real-time issues directly on your current screen within Zabbix frontend.

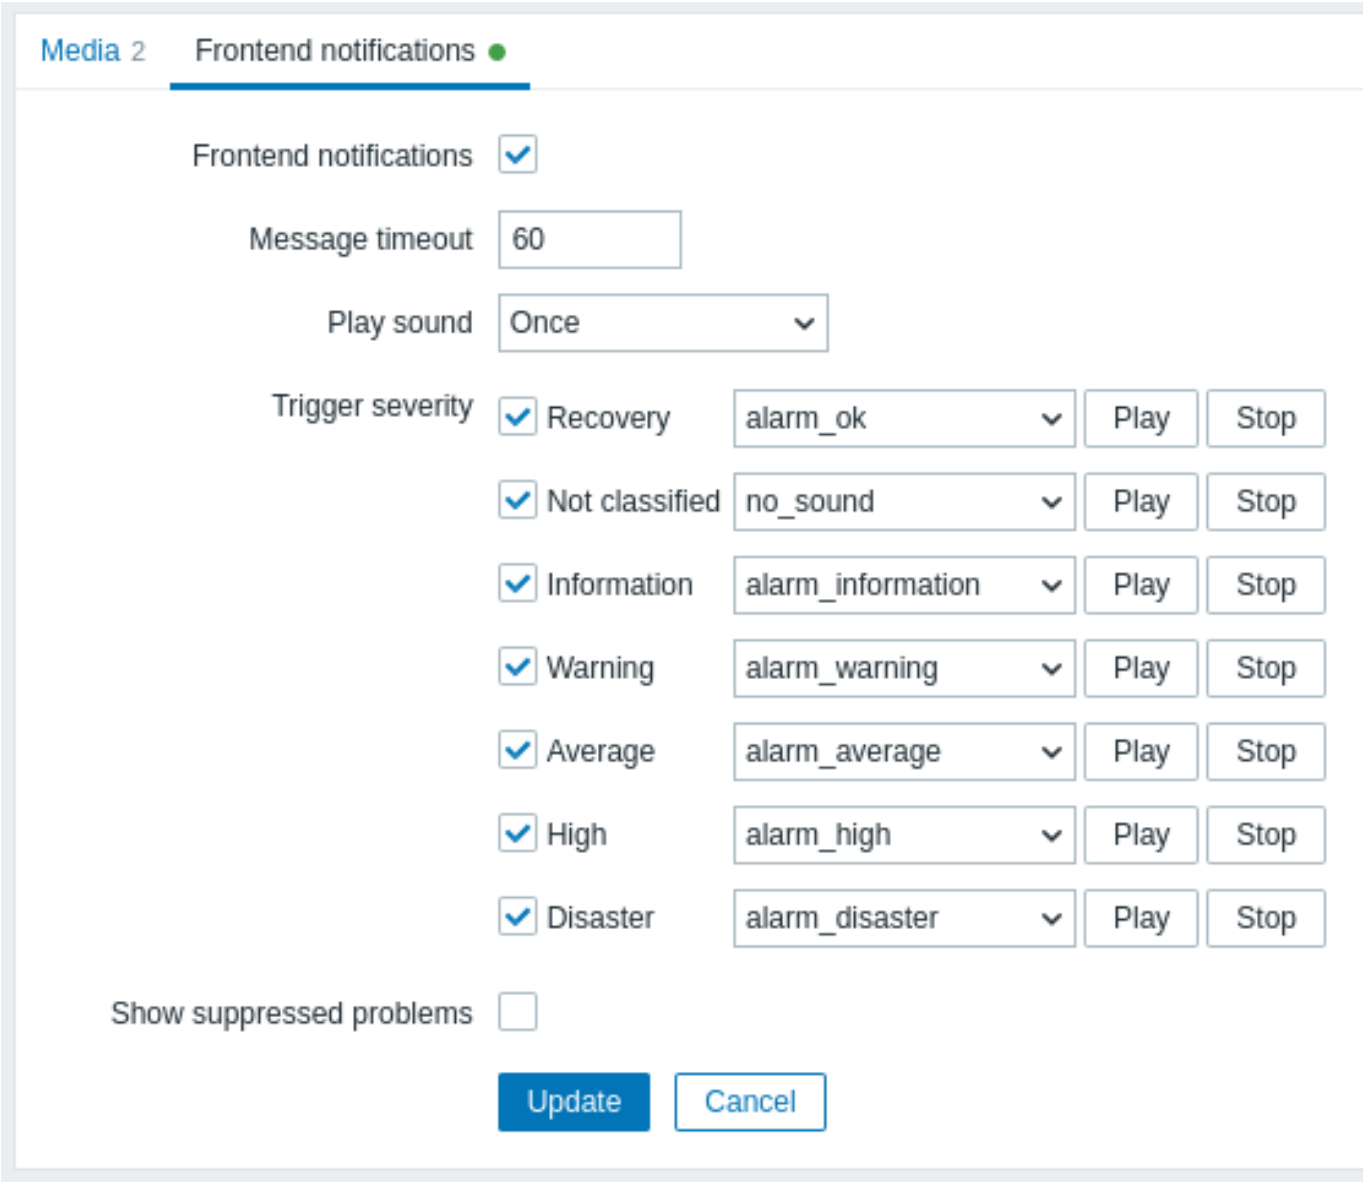
Without global notifications, when working outside the *Problems* or *Dashboard* sections, you would not receive any information about current issues. Global notifications ensure that this information is displayed, regardless of your current location within the Zabbix frontend.

Global notifications include both displaying a message and playing a sound.

Attention:
The autoplay of sounds might be disabled (by default) in recent browser versions. In such cases, you need to enable this setting manually.

Configuration

Global notifications can be enabled per user in the **Frontend notifications** tab of Notifications section.

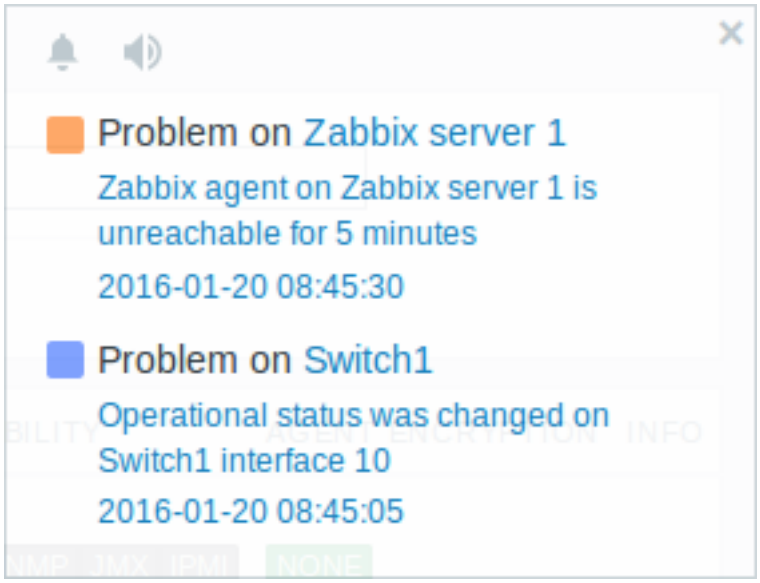


Parameter	Description
Frontend notifications	Mark the checkbox to enable global notifications.
Message timeout	Set the duration for which the message will be displayed. By default, messages remain on the screen for 60 seconds. Time suffixes are supported, for example: 30s, 5m, 2h, 1d.
Play sound	Set the duration for which the sound will be played. **Once** - sound is played once and fully; **10 seconds** - sound is repeated for 10 seconds; **Message timeout** - sound is repeated while the message is visible.
Trigger severity	Set the trigger severities for which global notifications and sounds will be activated. You can also select sounds appropriate for various severities. If no severity is marked, no messages will be displayed. Additionally, recovery messages will only be displayed for marked severities. For instance, if *Recovery* and *Disaster* are marked, global notifications will be displayed for problems and recoveries of *Disaster* severity triggers.

Parameter	Description
Show suppressed problems	Mark the checkbox to display notifications for problems that would otherwise be suppressed (not shown) due to host maintenance.

Global messages displayed

As messages arrive, they are displayed in a floating section on the right-hand side. You can freely reposition this section by dragging the section header.



For this section, several controls are available:

- **Snooze** button silences the currently active alarm sound;
- **Mute/Unmute** button switches between playing and not playing the alarm sounds at all.

Sound in browsers

Overview

Sound is used in global notifications.

For the sounds to be played in Zabbix frontend, *Frontend notifications* must be enabled in the user profile's *Frontend notifications* tab, with all trigger severities checked. Additionally, sounds should be enabled in the global notification pop-up window.

If, for any reason, audio cannot be played on the device, the button in the global notification pop-up window will remain permanently in the "mute" state, accompanied by the message "Cannot support notification audio for this device" upon hovering over the button.

Sounds, including the default audio clips, are supported in MP3 format only.

The sounds of Zabbix frontend have been successfully tested in recent Firefox and Opera browsers on Linux, and in Chrome, Firefox, Microsoft Edge, and Opera browsers on Windows.

Attention:
The autoplay of sounds might be disabled (by default) in recent browser versions. In such cases, you need to enable this setting manually.

Global search

It is possible to search Zabbix frontend for hosts, host groups, templates and template groups.

The search input box is located below the Zabbix logo in the menu. The search can be started by pressing *Enter* or clicking the search icon.



If there is a host that contains the entered string in any part of the name, a dropdown will appear, listing all such hosts (with the matching part highlighted in orange). The dropdown will also list a host if that host's visible name is a match to the technical name entered as a search string; the matching host will be listed, but without any highlighting.

Searchable attributes

Hosts can be searched by the following properties:

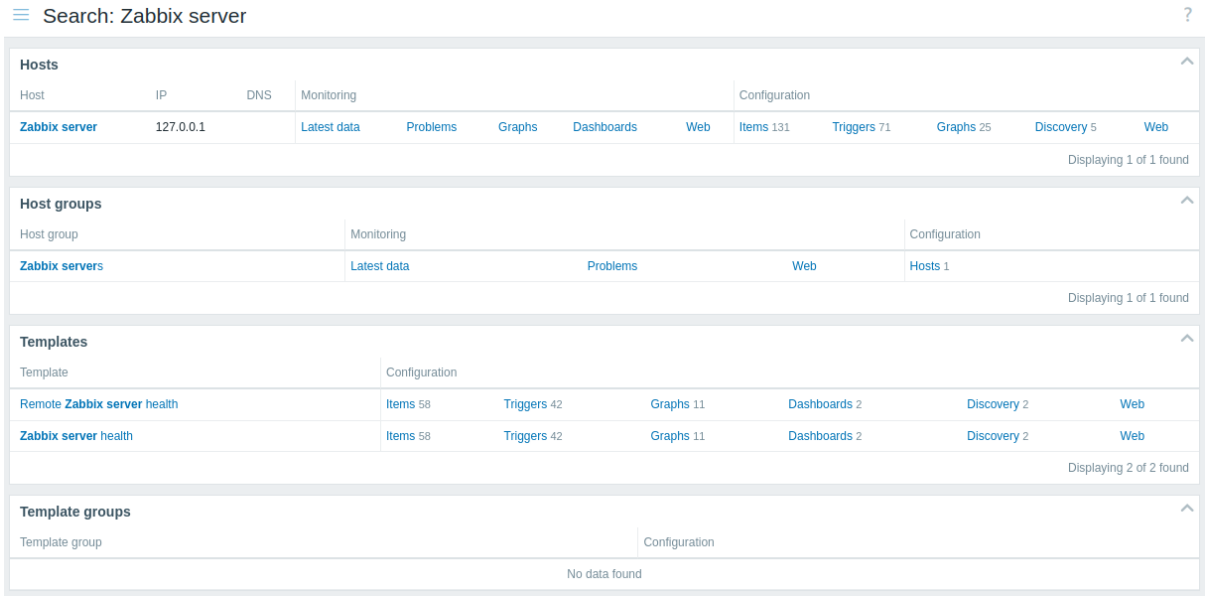
- Host name
- Visible name
- IP address
- DNS name

Templates can be searched by name or visible name. If you search by a name that is different from the visible name (of a template/host), in the search results it is displayed below the visible name in parentheses.

Host and template groups can be searched by name. Specifying a parent group implicitly selects all nested groups.

Search results

Search results consist of four separate blocks for hosts, host groups, templates and template groups.



It is possible to collapse/expand each individual block. The entry count is displayed at the bottom of each block, for example, *Displaying 13 of 13 found*. If there are no entries, the entry count is not displayed. Total entries displayed within one block are limited to 100.

Each entry provides links to monitoring and configuration data. See the full list of links.

For all configuration data (such as items, triggers, graphs) the amount of entities found is displayed by a number next to the entity name, in gray. **Note** that if there are zero entities, no number is displayed.

Enabled hosts are displayed in blue, disabled hosts in red.

Links available

For each entry the following links are available:

- Hosts
  - Monitoring
    - Latest data
    - Problems
    - Graphs
    - Host dashboards
    - Web scenarios
  - Configuration
    - Items
    - Triggers
    - Graphs
    - Discovery rules
    - Web scenarios
- Host groups
  - Monitoring
    - Latest data
    - Problems
    - Web scenarios
  - Configuration
    - Hosts
- Templates
  - Configuration
    - Items
    - Triggers
    - Graphs
    - Template dashboards
    - Discovery rules
    - Web scenarios
- Template groups
  - Configuration
    - Templates

Frontend maintenance mode

Overview

It is possible to temporarily disable Zabbix frontend to restrict access. This is useful for protecting Zabbix database from user-initiated changes, preserving its integrity.

While Zabbix frontend is in maintenance mode, you can safely stop the database and perform maintenance tasks.

Users from defined IP addresses will be able to interact with the frontend normally during maintenance mode.

Configuration

To enable maintenance mode, open the `maintenance.inc.php` file (located in `/conf` of the Zabbix HTML document directory on the web server) and uncomment the following lines:

```
// Maintenance mode.
define('ZBX_DENY_GUI_ACCESS', 1);

// Array of IP addresses allowed to connect to the frontend (optional).
$ZBX_GUI_ACCESS_IP_RANGE = array('127.0.0.1');

// Message shown on warning screen (optional).
$ZBX_GUI_ACCESS_MESSAGE = 'We are upgrading MySQL database till 15:00. Stay tuned...';
```

Note:
In most cases, the `maintenance.inc.php` file is located in `/conf` of the Zabbix HTML document directory on the web server. However, some operating systems and web servers may use a different location.
For example, the location for:

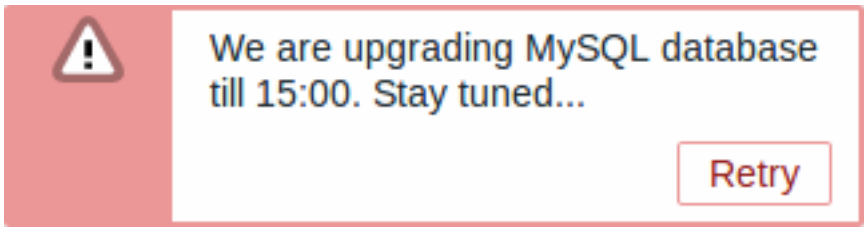
- SUSE and RedHat is `/etc/zabbix/web/maintenance.inc.php`.
- Debian-based systems is `/usr/share/zabbix/conf/`.

See also Copying PHP files.

Parameter	Details
ZBX_DENY_GUI_ACCESS	If defined with any value, maintenance mode will be enabled. To disable maintenance mode, comment out or delete.
ZBX_GUI_ACCESS_IP_RANGE	Array of IP addresses allowed to connect to the frontend (optional). For example: `array('192.168.1.1', '192.168.1.2')`
ZBX_GUI_ACCESS_MESSAGE	A message to inform users about the maintenance (optional). If undefined, the default message *'Zabbix is under maintenance'* will be used.

Display

Users will see the following screen when trying to access Zabbix frontend while in maintenance mode. The screen is refreshed every 30 seconds in order to return to a normal state without user intervention when the maintenance is over.



IP addresses defined in *ZBX_GUI_ACCESS_IP_RANGE* will be able to access the frontend as usual.

Page parameters

Overview

Most Zabbix web interface pages support various HTTP GET parameters that control what will be displayed. They may be passed by specifying parameter=value pairs after the URL, separated from the URL by a question mark (?) and from each other by ampersands (&).

Monitoring > Problems

The following parameters are supported:

Parameter	Description	Example
show	Filter option *Show*. Possible values: 1 - recent problems; 2 - all; 3 - in problem state.	show=1
name	Filter option *Problem*: freeform string.	name=Zabbix agent

Parameter	Description	Example
severities	Filter option *Severity*: array of selected severities in the format `severities[*]=*` (replace `*` with severity level). Possible values: 0 - not classified; 1 - information; 2 - warning; 3 - average; 4 - high; 5 - disaster.	severities[3]=3
inventory	Filter option *Host inventory*: array of inventory fields `[field],[value]`	inventory[0][field]=type&invento
evaltype	Filter option *Tags*: tag evaluation method. Possible values: 0 - and/or; 2 - or.	evaltype=0
tags	Filter option *Tags*: array of defined tags `[tag], [operator], [value]` Possible operator values: 0 - contains; 1 - equals; 2 - does not contain; 3 - does not equal; 4 - exists; 5 - does not exist.	tags[0][tag]=target&tags[0][ope
show_tags	Filter option *Show tags*. Possible values: 0 - none; 1 - one; 2 - two; 3 - three.	show_tags=3
tag_name_format	Filter option *Tag name*. Possible values: 0 - full name; 1 - shortened name; 2 - none.	tag_name_format=1
tag_priority	Filter option *Tag display priority*: comma-separated string of tag display priority	tag_priority=customer, target
show_suppressed	Filter option *Show suppressed problems*. Possible values: 1 - show; 2 - do not show.	show_suppressed=1
acknowledgement_status	Filter option *Acknowledgement status*. Possible values: 0 - all; 1 - unacknowledged; 2 - acknowledged.	acknowledgement_status=0
acknowledged_by_me	Filter option *By me*; supported only with `acknowledgement_status=2`. Possible values: 0 - disabled; 1 - enabled.	acknowledged_by_me=1
compact_view	Filter option *Compact view*. Possible values: 0 - disabled; 1 - enabled.	compact_view=1

Parameter	Description	Example
highlight_row	Filter option *Highlight whole row*. Possible values: 0 - disabled; 1 - enabled.	highlight_row=1
filter_name	Filter properties option *Name*: freeform string	filter_name=Databases
filter_show_counter	Filter properties option *Show number of records*. Possible values: 0 - disabled; 1 - enabled.	filter_show_counter=1
filter_custom_time	Filter properties option *Set custom time period*. Possible values: 0 - disabled; 1 - enabled.	filter_custom_time=1
sort	Column to sort by. Possible values: clock - sort by *Time* column; host - sort by *Host* column; severity - sort by *Severity* column; name - sort by *Problem* column.	sort=clock
sortorder	Sort order of results. Possible values: DESC - descending; ASC - ascending.	sortorder=DESC
age_state	Filter option *Age less than*; supported only with `show`=3. Possible values: 0 - disable age parameter; 1 - enable age parameter.	age_state=1
age	Filter option *Age less than*: integer, number of days; supported only with `age_state`=1 and `show`=3.	age=7
groupids	Filter option *Host groups*: array of host group IDs.	groupids[]=4
hostids	Filter option *Hosts*: array of host IDs	hostids[]=10084
triggerids	Filter option *Triggers*: array of trigger IDs	triggerids[]=22382
show_timeline	Filter option *Show timeline*; not supported with `compact_view`=1. Possible values: 0 - do not show; 1 - show.	show_timeline=1
details	Filter option *Show details*. Possible values: 0 - do not show; 1 - show.	details=1
from	Date range start, can be relative (e.g., now-1m); supported only with `filter_custom_time`=1	from=now-2h
to	Date range end, can be relative (e.g., now-1m); supported only with `filter_custom_time`=1	to=now

See also: Problems page filter options.

Kiosk mode

The kiosk mode in supported frontend pages can be activated using URL parameters. For example, in dashboards:

- `/zabbix.php?action=dashboard.view&kiosk=1` - activate kiosk mode
- `/zabbix.php?action=dashboard.view&kiosk=0` - activate normal mode

Slideshow

It is possible to activate a slideshow in the dashboard:

- `/zabbix.php?action=dashboard.view&slideshow=on` - activate slideshow

Definitions

Overview

While many things in the frontend can be configured using the frontend itself, some customizations are currently only possible by editing a definitions file.

This file is `defines.inc.php` located in /include of the Zabbix HTML document directory.

Parameters

Parameters in this file that could be of interest to users:

- ZBX_MIN_PERIOD

Minimum graph period, in seconds. One minute by default.

- GRAPH_YAXIS_SIDE_DEFAULT

Default location of Y axis in simple graphs and default value for drop down box when adding items to custom graphs. Possible values: 0 - left, 1 - right.

Default: 0

- ZBX_SESSION_NAME

String used as the name of the Zabbix frontend session cookie.

Default: zbx_sessionid

- ZBX_DATA_CACHE_TTL

TTL timeout in seconds used to invalidate data cache of Vault response. Set 0 to disable Vault response caching.

Default: 60

- SUBFILTER_VALUES_PER_GROUP

Number of subfilter values per group (For example, in the latest data subfilter).

Default: 1000

- ZBX_MAX_WIDGET_LINES

Maximum number of widget lines to display.

Default: 1000

Additional frontend languages

Overview In order to use any other language than English in Zabbix web interface, its locale should be installed on the web server. Additionally, the PHP gettext extension is required for the translations to work.

Installing locales To list all installed languages, run:

```
locale -a
```

If some languages that are needed are not listed, open the `/etc/locale.gen` file and uncomment the required locales. As Zabbix uses UTF-8 encoding, you need to select locales with UTF-8 charset.

Now run:

```
locale-gen
```

Restart the web server. Note that if PHP-FPM is installed and used by the web server, it also needs to be restarted along with the web server.

The locales should now be installed. It may be required to reload Zabbix frontend page in browser using `Ctrl + F5` for new languages to appear.

Installing Zabbix If installing Zabbix directly from Zabbix git repository, translation files should be generated manually. To generate translation files, run:

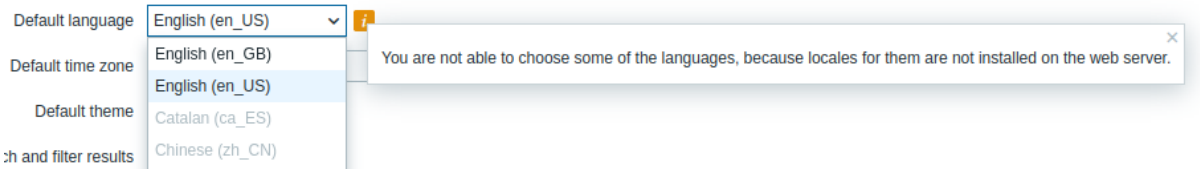
```
make gettext
locale/make_mo.sh
```

This step is not needed when installing Zabbix from packages or source `tar.gz` files.

Selecting a language There are several ways to select a language in Zabbix web interface:

- When installing web interface - in the frontend installation wizard. Selected language will be set as system default.
- After the installation, system default language can be changed in the *Administration* > *General* > *GUI*.
- Language for a particular user can be changed in the user profile.

If a locale for a language is not installed on the machine, this language will be greyed out in Zabbix language selector. An orange icon is displayed next to the language selector if at least one locale is missing. Upon clicking this icon the following message will be displayed: "You are not able to choose some of the languages, because locales for them are not installed on the web server."



Creating your own theme

Overview

By default, Zabbix provides a number of predefined frontend themes. To create a custom theme, follow the step-by-step procedure provided here.

Step 1: Create your theme

To create your own theme:

1. Create a CSS file from scratch or copy an existing theme (such as `blue-theme.css`, `dark-theme.css`, or another one) from the `assets/styles/` directory and modify it.
2. Save your theme (for example, `custom-theme.css`) in the `assets/styles/` directory.

Step 2: Enable your theme in the Theme list

To make your custom theme available in the frontend, you must add it to the list of themes returned by the `APP::getThemes()` method. This is done by overriding the `ZBase::getThemes()` method inside the APP class as follows:

1. Open the `include/classes/core/APP.php` file and locate the APP class:

```
class APP extends ZBase {

}
```

2. Inside the class, before its final closing brace }, add the following method:

```
public static function getThemes() {
        return array_merge(parent::getThemes(), [
            'custom-theme' => _('Custom theme')
        ]);
}
```

where:

- `custom-theme` - the internal identifier of your theme and must match the name of your CSS file without the `.css` extension;
- `Custom theme` - the display name you will see in the frontend.

To add multiple themes, list them in the array, separating each entry with a comma. A trailing comma after the last entry is optional.

For example:

```
      public static function getThemes() {
          return array_merge(parent::getThemes(), [
              'custom-theme' => _('Custom theme'),
              'anothertheme' => _('Another theme'),
              'onemoretheme' => _('One more theme')
          ]);
      }
```

Step 3: Change graph colors

To change graph colors, a record for the new theme must be added to the `graph_theme` database table.

Example of adding a record to a MySQL/MariaDB database:

```
    mysql -u zabbix -p
    # Enter password:
```

```
mysql> USE zabbix;
mysql> INSERT INTO graph_theme (
    graphthemeid,
    theme,
    backgroundcolor,
    graphcolor,
    gridcolor,
    maingridcolor,
    gridbordercolor,
    textcolor,
    highlightcolor,
    leftpercentilecolor,
    rightpercentilecolor,
    nonworktimecolor,
    colorpalette
)
SELECT
    MAX(graphthemeid) + 1,
    'custom-theme',
    'FFFFFF',
    'FFFFFF',
    'CCD5D9',
    'ACBBC2',
    'ACBBC2',
    '1F2C33',
    'CC745E',
    '5ECCAB',
    'CC745E',
    'EBEBEB',
    '329F7E,C2583D,346D91,B26E44,CC6C91,7A6DC2,C4AA56,7C2731,BF75B8,73A350,B04833,633A60,879CCC,7FAD6C,324
FROM graph_theme;
```

Step 4: Use the new theme

In the Zabbix frontend, you can either set this theme as the default theme or select it in a user profile.

Debug mode

Overview

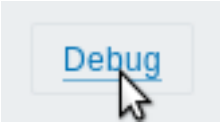
Debug mode may be used to diagnose performance problems with frontend pages.

Configuration

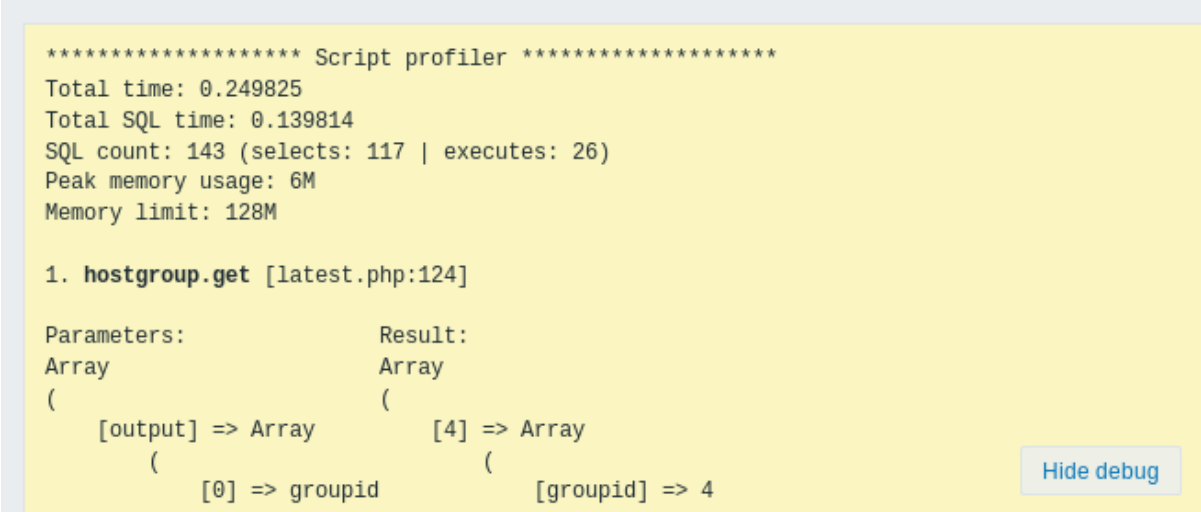
Debug mode can be activated for individual users who belong to a user group:

- when configuring a user group;
- when viewing configured user groups.

When *Debug mode* is enabled for a user group, its users will see a *Debug* button in the lower-right corner of the browser window:



Clicking the *Debug* button opens a new window below the page contents which contains the SQL statistics of the page, along with a list of API calls and individual SQL statements:



In case of performance problems with the page, this window may be used to search for the root cause of the problem.

Warning:
Enabled *Debug mode* negatively affects frontend performance.

Cookies used by Zabbix

Overview

This page provides a list of cookies used by Zabbix.

Name	Description	Values	Expires/Max-Age	HttpOnly[1]	Secure[2]
ZBX_SESSION_ID	Zabbix frontend session data, stored as JSON encoded by base64		Session (expires when the browsing session ends)	+	+ (only if HTTPS is enabled on a web server)
tab	Active tab number; this cookie is only used on pages with multiple tabs (e.g., *Host*, *Trigger* or *Action* configuration page) and is created, when a user navigates from a primary tab to another tab (such as *Tags* or *Dependencies* tab). 0 is used for the primary tab.	Example: 1	Session (expires when the browsing session ends)	-	-
browserwarning_ignore	Whether a warning about using an outdated browser should be ignored.	yes	Session (expires when the browsing session ends)	-	-

Name	Description	Values	Expires/Max-Age	HttpOnly[1]	Secure[2]
system-message-ok	A message to show as soon as page is reloaded.	Plain text message	Session (expires when the browsing session ends) or as soon as page is reloaded	+	-
system-message-error	An error message to show as soon as page is reloaded.	Plain text message	Session (expires when the browsing session ends) or as soon as page is reloaded	+	-

Note:
Forcing 'HttpOnly' flag on Zabbix cookies by a webserver directive is not supported.

Footnotes

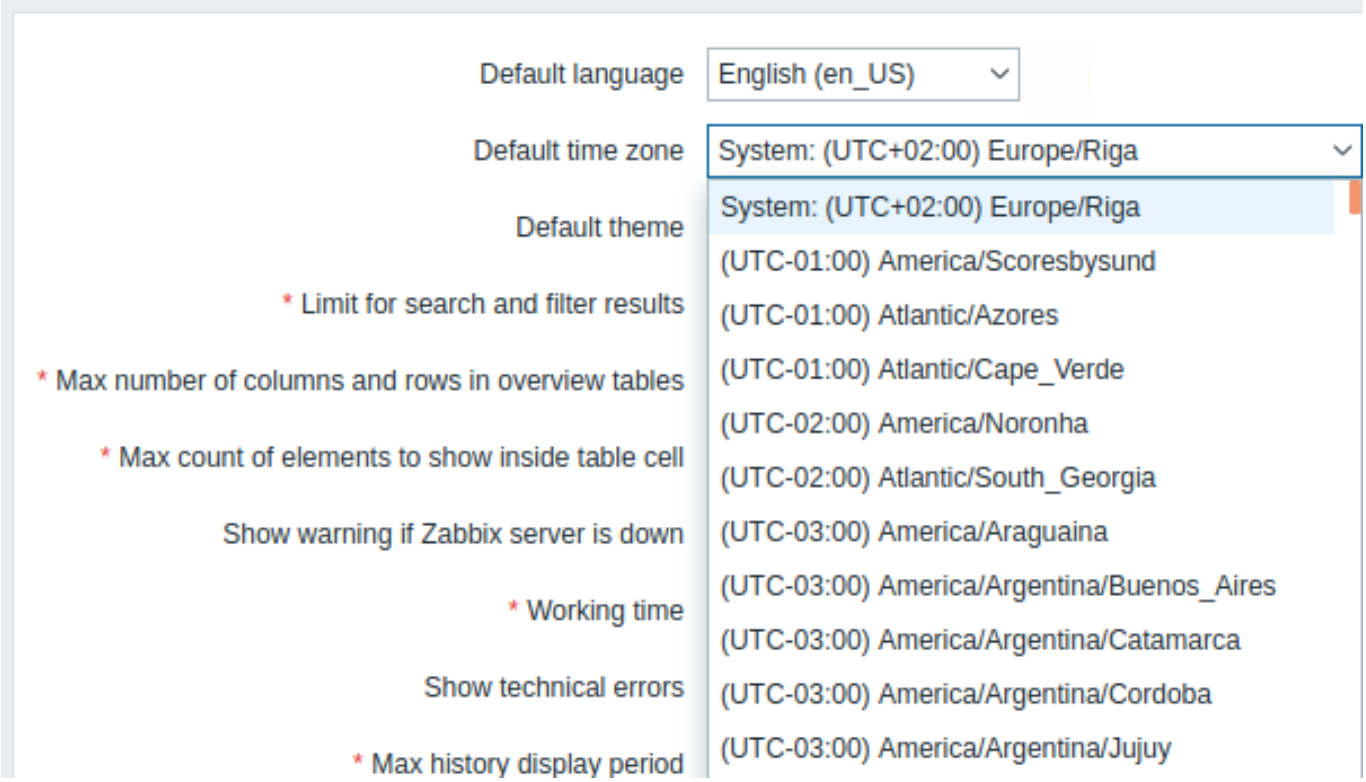
[1] When `HttpOnly` is 'true' the cookie will be made accessible only through the HTTP protocol. This means that the cookie won't be accessible by scripting languages, such as JavaScript. This setting can effectively help to reduce identity theft through XSS attacks (although it is not supported by all browsers).

[2] `Secure` indicates that the cookie should only be transmitted over a secure HTTPS connection from the client. When set to 'true', the cookie will only be set if a secure connection exists.

Time zones

Overview

The frontend time zone can be set globally in the frontend and adjusted for individual users.



If *System* is selected, the web server time zone will be used for the frontend (including the value of `date.timezone` in `php.ini`, if set), while Zabbix server will use the time zone of the machine it is running on.

Note:
Zabbix server will only use the specified global/user time zone when expanding macros in notifications (e.g., `{EVENT.TIME}` can expand to a different time zone per user) and for the time limit when notifications are sent (see "When active" setting in the user media configuration).

The choice of time zone does not affect the frontend time/date format. Instead, you may adjust the interface language (either at installation or under user settings) - selecting *English (en_US)* will also enable the US time/date format in the frontend.

Configuration

The global time zone:

- Can be set manually when installing the frontend.
- Can be modified in *Administration* > *General* > *GUI*.

User-level time zone:

- Can be set when configuring/updating a user.
- Can be set by each user in their user profile.

See also: Aligning time zones when using scheduling intervals.

Resetting password

Overview This section describes the steps for resetting user passwords in Zabbix.

Steps Turn to your Zabbix administrator if you have forgotten your Zabbix password and cannot log in.

A Super administrator user can change passwords for all users in the user configuration form.

If a Super administrator has forgotten their password and cannot log in, the following SQL query must be run to apply the default password to the Super admin user (replace 'Admin' with the appropriate Super admin username):

```
UPDATE users SET passwd = '$2a$10$ZXIvHAEP2ZM.dLXTm6uPHOMVlARXX7cqjbhM6Fn0cANzkCQBWpMrS' WHERE username =
```

After running this query, the user password will be set to *zabbix*. Make sure to change the default password on the first login.

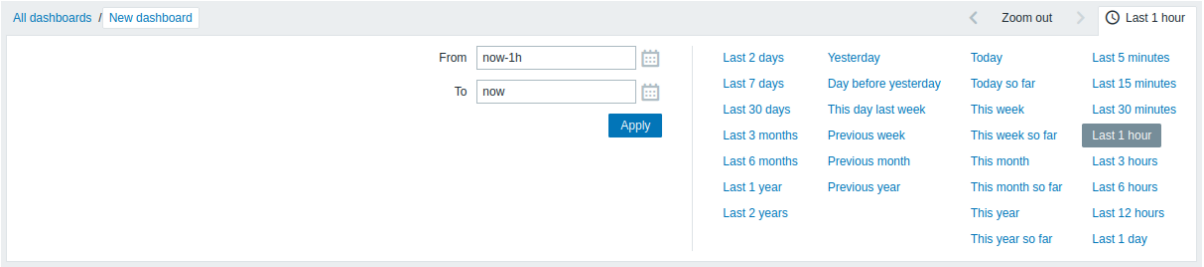
Time period and host selectors

Overview

This page describes the time period and host selectors available on dashboards. These selectors allow the displayed data to be dynamically adjusted across all compatible dashboard widgets without editing individual widget settings.

Time period selector

The time period selector allows you to select often required periods with one mouse click. It can be expanded or collapsed by clicking the time period selector button in the upper-right corner of the dashboard.



Options such as *Today*, *This week*, etc., display the whole period, including the hours/days in the future. Options such as *Today so far*, *This week so far*, etc., display only the hours passed.

Once a period is selected, it can be moved back and forth in time by clicking the arrow buttons. The *Zoom out* button allows you to zoom out the period by 50% in each direction.

> **Note:**
> For graphs (except those with the *Custom* time period), selecting the displayed time period is also possible by highlighting an area in the graph with the left mouse button. Once you release the left mouse button, the graph will zoom into the highlighted area. Zooming out is also possible by double-clicking in the graph.

The *From/To* fields display the selected period in either absolute time syntax (in format `Y-m-d H:i:s`) or relative time syntax. A relative time period can contain one or several mathematical operations (- or +), for example, `now-1d` or `now-1d-2h+5m`.

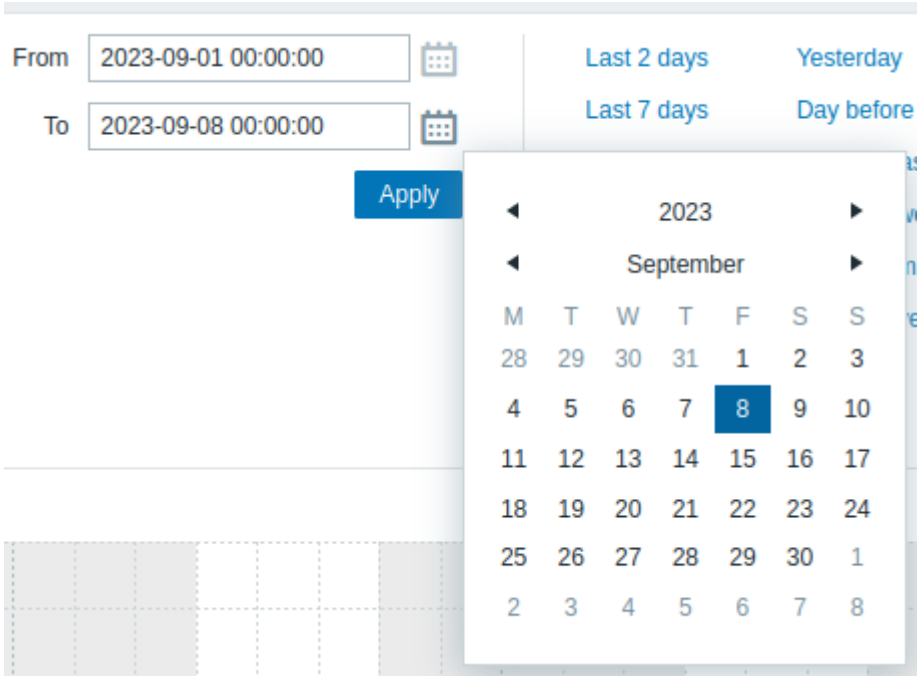
The following relative time abbreviations are supported:

- now
- s (seconds)
- m (minutes)
- h (hours)
- d (days)
- w (weeks)
- M (months)
- y (years)

Precision is supported in the time period selector (for example, /M in now-1d/M). Details of precision:

Precision	From	To
m	Y-m-d H:m:00	Y-m-d H:m:59
h	Y-m-d H:00:00	Y-m-d H:59:59
d	Y-m-d 00:00:00	Y-m-d 23:59:59
w	Monday of the week 00:00:00	Sunday of the week 23:59:59
M	First day of the month 00:00:00	Last day of the month 23:59:59
y	1st of January of the year 00:00:00	31st of December of the year 23:59:59

It is also possible to select a time period using the date picker. To open it, click the calendar icon next to the *From/To* fields.



Note:
Within the date picker, you can navigate between year/month/date using Tab, Shift+Tab, and keyboard arrow buttons. Pressing Enter confirms the selection.

Examples:

From	To	Selected period
now/d	now/d	00:00 - 23:59 today
now/d	now/d+1d	00:00 today - 23:59 tomorrow
now/w	now/w	Monday 00:00:00 - Sunday 23:59:59 this week
now-1y/w	now-1y/w	The week of Monday 00:00:00 - Sunday 23:59:59 one year ago

Attention:
Using "now/M+1M" for the *To* parameter may add 31 days, which can result in the date shifting by 1-3 days depending on the number of days in the month. For example, if used in January, the result may be 02 March instead of the expected 28 February. To avoid this issue, use "now/M-3d+1M/M", which adjusts for month length accurately. Similarly, if configuring *From* to go backward, use "now/M+3d-1M/M".

Host selector

The host selector lets you select a host to have its data displayed in compatible dashboard widgets.

It becomes available after editing and saving a dashboard, if at least one widget is configured with the *Override host* parameter set to use the dashboard as a data source.



Configuration parameters: Zabbix frontend

Overview Configuration parameters allow customizing Zabbix frontend behavior.

Specify parameter values in the frontend configuration file (`zabbix.conf.php`).

This page is structured as:

- Parameter overview (to see full details of a parameter, click its name)
- Parameter details

Parameter	Description
$DB['TYPE']	Database type.
$DB['SERVER']	Database host name or IP address.
$DB['PORT']	Database port.
$DB['DATABASE']	Database name used by the Zabbix frontend to connect to the configured database.
$DB['USER']	Username used by the Zabbix frontend to connect to the database.
$DB['PASSWORD']	Password used by the Zabbix frontend to connect to the database.
$DB['SCHEMA']	Database schema name used for PostgreSQL.
$DB['ENCRYPTION']	Enable or disable TLS encryption for connections between the Zabbix frontend and the database.
$DB['KEY_FILE']	Path to the client private key file used for the database TLS connection.
$DB['CERT_FILE']	Path to the client certificate file used for the database TLS connection.
$DB['CA_FILE']	Path to the CA certificate file used for the database TLS connection.
$DB['VERIFY_HOST']	Enable or disable verification of the database server hostname in the TLS certificate.
$DB['CIPHER_LIST']	List of allowed TLS ciphers for the database connection.
$DB['VAULT']	Specify the Vault secrets manager used to store database credentials.
$DB['VAULT_URL']	URL of the Vault secrets manager used to retrieve database credentials.
$DB['VAULT_PREFIX']	Prefix for the Vault path or query used to retrieve database credentials.
$DB['VAULT_DB_PATH']	Path to the database credentials in the Vault secrets manager.
$DB['VAULT_TOKEN']	Authentication token used to access the Vault secrets manager.
$DB['VAULT_APP_ROLE_ID']	AppRole role ID used to authenticate to HashiCorp Vault.
$DB['VAULT_APP_SECRET_ID']	AppRole secret ID used to authenticate to HashiCorp Vault.
$DB['VAULT_CERT_FILE']	Path to the client certificate file used to authenticate to the Vault secrets manager.
$DB['VAULT_KEY_FILE']	Path to the client private key file used to authenticate to the Vault secrets manager.
$DB['VAULT_CACHE']	Enables or disables caching of credentials retrieved from the Vault secrets manager.
$ZBX_SERVER	Specify a value to override the Zabbix hostname/IP address.
$ZBX_SERVER_PORT	Specify a value to override the Zabbix server port.
$ZBX_SERVER_NAME	Zabbix server name.
$IMAGE_FORMAT_DEFAULT	Default image format.
$HISTORY_PROVIDERS[]	Configuration of history storage providers for Elasticsearch or ClickHouse.
$SSO['SETTINGS']	Configuration settings for SAML sign-in.
$SSO['CERT_STORAGE']	Specify where SAML certificates are stored.
$SSO['SP_KEY']	Path to the SAML service provider (SP) private key file when the private key is stored on the filesystem.
$SSO['SP_CERT']	Path to the SAML service provider (SP) certificate file when the certificate is stored on the filesystem.
$SSO['IDP_CERT']	Path to the SAML identity provider certificate file when the certificate is stored on the filesystem.
$ZBX_FEATURE_FLAGS['banners_enabled']	Enable or disable support for banners.
$ZBX_FEATURE_FLAGS['http_auth_enabled']	Enable or disable user HTTP authentication.
$ZBX_FEATURE_FLAGS['modules_enabled']	Enable or disable access to modules.
$ZBX_FEATURE_FLAGS['media_type_edit_disabled']	Disable editing of specific pre-configured media types.
$ZBX_SERVER_TLS['ACTIVE']	Specifies whether TLS is enabled for connections to the Zabbix server.

Parameter	Description
$ZBX_SERVER_TLS['CA_FILE']	Path to the Certificate Authority (CA) certificate file used to verify the server certificate.
$ZBX_SERVER_TLS['KEY_FILE']	Path to the client private key file corresponding to the client certificate.
$ZBX_SERVER_TLS['CERT_FILE']	Path to the client certificate file when mutual TLS authentication is required.
$ZBX_SERVER_TLS['CERTIFICATE_ISSUER']	Distinguished Name (DN) of the issuer to match against the server certificate.
$ZBX_SERVER_TLS['CERTIFICATE_SUBJECT']	Distinguished Name (DN) of the subject to match against the server certificate.
$TELEMETRY_PROVIDERS[]	Configuration of a telemetry data source.

Parameter details

$DB['TYPE']

Database type.

Values: `MYSQL` - MySQL, `POSTGRESQL` - PostgreSQL.

$DB['SERVER']

Database host name or IP address.

If `$DB['SERVER']` is not defined, the default Unix-domain socket is used. For MySQL, setting `$DB['SERVER']` to `localhost` also results in a connection via the default Unix socket.
 For PostgreSQL, `$DB['SERVER']` can contain multiple `host:port` entries separated by commas. If a port is not defined for a host in the comma-separated list, the value defined in `$DB['PORT']` is used as the default port. If a socket path is specified in `$DB['SERVER']` (for example, `/var/run/pgbouncer`), that Unix-domain socket is used.

Examples:

```
$DB['SERVER'] = 'localhost';
$DB['SERVER'] = 127.0.0.1;
```

$DB['PORT']

Database port.

If not defined, the value 0 is used.

$DB['DATABASE']

Database name used by the Zabbix frontend to connect to the configured database.

Example:

```
$DB['DATABASE'] = 'zabbix';
```

$DB['USER']

Username used by the Zabbix frontend to connect to the database.

$DB['PASSWORD']

Password used by the Zabbix frontend to connect to the database.

$DB['SCHEMA']

Database schema name used for PostgreSQL.

$DB['ENCRYPTION']

Enable or disable encryption for connections between the Zabbix frontend and the database.

Values: `true` - enabled, `false` - disabled.

$DB['KEY_FILE']

Path to the client private key file used for the database TLS connection.

Example:

```
// For MySQL:
$DB['KEY_FILE'] = '/etc/ssl/mysql/client-key.pem';
```

$DB['CERT_FILE']

Path to the client certificate file used for the database TLS connection.

Example:

```
// For MySQL:
$DB['CERT_FILE'] = '/etc/ssl/mysql/client-cert.pem';
```

$DB['CA_FILE']

Path to the CA certificate file used for the database TLS connection.

Examples:

```
// For MySQL:
$DB['CA_FILE'] = '/etc/ssl/mysql/ca.pem';
// For PostgreSQL:
$DB['CA_FILE'] = '/etc/ssl/pgsql/root.crt';
```

$DB['VERIFY_HOST']

Enable or disable verification of the database server hostname in the TLS certificate.

Values: `true` - enabled, `false` - disabled.

$DB['CIPHER_LIST']

List of allowed TLS ciphers for the database connection.

Examples:

```
// For MySQL:
$DB['CIPHER_LIST'] = 'TLS_AES_256_GCM_SHA384:TLS_CHACHA20_POLY1305_SHA256:TLS_AES_128_GCM_SHA256:TLS_AES_1
```

$DB['VAULT']

Specify the Vault secrets manager used to store database credentials.

Values: `CyberArk` or `HashiCorp`.

Examples:

```
// For CyberArk:
$DB['VAULT'] = 'CyberArk';
// For HashiCorp:
$DB['VAULT'] = 'HashiCorp';
```

$DB['VAULT_URL']

URL of the Vault secrets manager used to retrieve database credentials.

Examples:

```
// For CyberArk:
$DB['VAULT_URL'] = 'https://127.0.0.1:1858';
// For HashiCorp:
$DB['VAULT_URL'] = 'https://localhost:8200';
```

$DB['VAULT_PREFIX']

Prefix for the Vault path or query used to retrieve database credentials.

Example:

```
// For HashiCorp:
$DB['VAULT_PREFIX'] = '/v1/secret/data/zabbix/';
```

$DB['VAULT_DB_PATH']

Path to the database credentials in the Vault secrets manager.

Examples:

```
// For CyberArk:
$DB['VAULT_DB_PATH'] = 'AppID=foo&Query=Safe=bar;Object=buzz';
// For HashiCorp:
$DB['VAULT_DB_PATH'] = 'database';
```

$DB['VAULT_TOKEN']

Authentication token used to access the Vault secrets manager.

Example:

```
// For HashiCorp:
$DB['VAULT_TOKEN'] = 'hvs.CAESIIG_PILmULFYOsEyWHxkZ2mF2a8VPKNLE8eHqd4autYGGh4KHGh2cy5aeTYONFNSaUp3ZnpWbDF1
```

$DB['VAULT_APP_ROLE_ID']

VaultAppSecretID role ID used to authenticate to HashiCorp Vault.

Example:

```
$DB['VAULT_APP_ROLE_ID'] = '6f1567f1-538d-7d70-e8e4-5d3e59471b2e';
```

$DB['VAULT_APP_SECRET_ID']

VaultAppSecretID used to authenticate to HashiCorp Vault.

Example:

```
$DB['VAULT_APP_SECRET_ID'] = '89dac8b1-f7d6-96a7-2ecf-f94d2de3d6f1';
```

$DB['VAULT_CERT_FILE']

Path to the client certificate file used to authenticate to the Vault secrets manager.

Example:

```
// For CyberArk:
$DB['VAULT_CERT_FILE'] = 'conf/certs/cyberark-cert.pem';
```

$DB['VAULT_KEY_FILE']

Path to the client private key file used to authenticate to the Vault secrets manager.

Example:

```
// For CyberArk:
$DB['VAULT_KEY_FILE'] = 'conf/certs/cyberark-key.pem';
```

$DB['VAULT_CACHE']

Enables or disables caching of credentials retrieved from the Vault secrets manager.

Values: `true` - enabled, `false` - disabled.

$ZBX_SERVER

Specify a value to override the Zabbix hostname/IP address.
 If $ZBX_SERVER is specified without $ZBX_SERVER_PORT, the default port (10051) is used.

$ZBX_SERVER_PORT

Specify a value to override the Zabbix server port.
 This parameter is used together with $ZBX_SERVER.

$ZBX_SERVER_NAME

Zabbix server name.

$IMAGE_FORMAT_DEFAULT

Default image format.

Default: IMAGE_FORMAT_PNG
 Values: - IMAGE_FORMAT_PNG - PNG format - IMAGE_FORMAT_JPEG - JPEG format - IMAGE_FORMAT_GIF - GIF format

$HISTORY_PROVIDERS[]

Configuration of history storage providers for Elasticsearch or ClickHouse.

Supported configuration parameters: - `types` - array of data types to be stored in the external storage; - `provider` - history provider type: `elasticsearch` or `clickhouse`; - `url` - history provider URL; - `db` - database name (used for ClickHouse); - `username` - database user (used for ClickHouse); - `password` - database password (used for ClickHouse).

Examples:

```
// For ClickHouse:
$HISTORY_PROVIDERS[] = [
    'types' => ['uint', 'dbl', 'str', 'log', 'text', 'json'],
    'provider' => 'clickhouse',
    'url' => 'http://localhost:8123',
    'db' => 'zabbix',
    'username' => 'zabbix',
    'password' => 'zabbix'
];
// For Elasticsearch:
$HISTORY_PROVIDERS[] = [
    'types' => ['uint', 'dbl', 'str', 'log', 'text', 'json'],
    'provider' => 'elasticsearch',
    'url' => 'http://localhost:9200'
];
```

$SSO['SETTINGS']

Configuration settings for SAML sign-in. The `$SSO['SETTINGS']` array must follow the structure expected by the *SAML PHP Toolkit* library (supplied with Zabbix).

Only the following options can be specified in `$SSO['SETTINGS']`:

- strict
- baseurl
- compress
- contactPerson
- organization
- sp (only options specified in this list)
    - attributeConsumingService
    - x509certNew
- idp (only options specified in this list)
    - singleLogoutService (only one option)
        * responseUrl
    - certFingerprint
    - certFingerprintAlgorithm
    - x509certMulti
- security (only options specified in this list)
    - signMetadata
    - wantNameId
    - requestedAuthnContext
    - requestedAuthnContextComparison
    - wantXMLValidation
    - relaxDestinationValidation
    - destinationStrictlyMatches
    - rejectUnsolicitedResponsesWithInResponseTo
    - signatureAlgorithm
    - digestAlgorithm
    - lowercaseUrlencoding

All other options are taken from the database and cannot be overridden. The debug option will be ignored.

If the Zabbix UI is behind a proxy or load balancer, the custom `use_proxy_headers` option can be used:

- `false` (default) - ignore the `X-Forwarded-*` HTTP headers;
- `true` - use the `X-Forwarded-*` HTTP headers to build the base URL.

For a full description of available configuration options, refer to the official library documentation.

Example:

```
$SSO['SETTINGS'] = [
    'security' => [
        'signatureAlgorithm' => 'http://www.w3.org/2001/04/xmldsig-more#rsa-sha384',
        'digestAlgorithm' => 'http://www.w3.org/2001/04/xmldsig-more#sha384',
        // ...
    ],
```

```
    // ...
];
```

$SSO['CERT_STORAGE']

Specify where SAML certificates are stored.

Default: `file`
 Values: - `database` - the private key and certificates are stored in the database; - `file` - the private key and certificates are stored on the filesystem.

Example:

```
$SSO['CERT_STORAGE'] = 'database';
```

$SSO['SP_KEY']

Path to the SAML service provider (SP) private key file when the private key is stored on the filesystem.

Example:

```
$SSO['SP_KEY'] = 'conf/certs/sp.key';
```

$SSO['SP_CERT']

Path to the SAML service provider (SP) certificate file when the certificate is stored on the filesystem.

Example:

```
$SSO['SP_CERT'] = 'conf/certs/sp.crt';
```

$SSO['IDP_CERT']

Path to the SAML identity provider certificate file when the certificate is stored on the filesystem.

Example:

```
$SSO['IDP_CERT'] = 'conf/certs/idp.crt';
```

$ZBX_FEATURE_FLAGS['banners_enabled']

Enable or disable support for banners.

Default: `true`
 Values: `true` - enabled, `false` - disabled.

$ZBX_FEATURE_FLAGS['http_auth_enabled']

Enable or disable user HTTP authentication.

Default: `true`
 Values: `true` - enabled, `false` - disabled.

$ZBX_FEATURE_FLAGS['modules_config_enabled']

Enable or disable access to modules.

Default: `true`
 Values: `true` - enabled, `false` - disabled.

$ZBX_FEATURE_FLAGS['media_type_denylist']

Disable editing of specific pre-configured media types.

Values: `email`, `script`, `sms`, `webhook`, `push`.

Example:

```
$ZBX_FEATURE_FLAGS['media_type_denylist'] = ['webhook','push'];
```

$ZBX_SERVER_TLS['ACTIVE']

Specifies whether a TLS connection between the Zabbix frontend and Zabbix server is enabled.

Default: `false`
 Values: `true` - enabled, `false` - disabled.

$ZBX_SERVER_TLS['CA_FILE']

Path to the Certificate Authority (CA) certificate file used to verify the server certificate.

Example:

```
$ZBX_SERVER_TLS['CA_FILE'] = '/etc/zabbix/ssl/ca.crt';
```

$ZBX_SERVER_TLS['KEY_FILE']

Path to the client private key file corresponding to the client certificate used for mutual TLS authentication.

Example:

```
$ZBX_SERVER_TLS['KEY_FILE'] = '/etc/zabbix/ssl/frontend.key';
```

$ZBX_SERVER_TLS['CERT_FILE']

Path to the client certificate file when mutual TLS authentication is required.

Example:

```
$ZBX_SERVER_TLS['CERT_FILE'] = '/etc/zabbix/ssl/frontend.crt';
```

$ZBX_SERVER_TLS['CERTIFICATE_ISSUER']

Distinguished Name (DN) of the issuer to match against the server certificate.

Example:

```
$ZBX_SERVER_TLS['CERTIFICATE_ISSUER']  = 'CN=MyZabbixCA';
```

$ZBX_SERVER_TLS['CERTIFICATE_SUBJECT']

Distinguished Name (DN) of the subject to match against the server certificate.

Example:

```
$ZBX_SERVER_TLS['CERTIFICATE_SUBJECT'] = 'CN=zabbix-server.example.com';
```

$TELEMETRY_PROVIDERS[]

Configuration of a telemetry data source.

Supported configuration parameters for ClickHouse: - `url` - set the HTTP or HTTPS URL used to connect to the telemetry data source. - `db` - set the name of the database. - `username` - set the username used to log in to the database. - `password` - set the password used to log in to the database. - `ssl_cert_file`, `ssl_key_file`, `ssl_key_password`, `ssl_verify_peer`, `ssl_verify_host`, `ssl_ca_location`, `ssl_ca_file` - set these options to configure a secure (TLS/SSL) connection to the database. - `vault_path` - set the path to a Vault secret holding the `username` and `password` values. Do not use together with the `username` and `password` options.

Example:

```
$TELEMETRY_PROVIDERS[] = [
    'provider' => 'clickhouse',
    'url' => 'http://localhost:8123',
    'db' => 'zabbix',
    'username' => 'zabbix',
    'password' => 'zabbix'
];
```

API

Overview The Zabbix API allows you to programmatically retrieve and modify configuration of Zabbix and provides access to historical data. It is widely used to:

- Create new applications to work with Zabbix.
- Integrate Zabbix into a third-party software.
- Automate routine tasks.

The Zabbix API is an HTTP-based API, and it is shipped as a part of the web frontend. It uses the JSON-RPC 2.0 protocol, which means two things:

- The API consists of a set of separate methods.
- Requests and responses between the clients and the API are encoded using the JSON format.

For more information about the protocol and JSON, see the JSON-RPC 2.0 specification and the JSON format homepage.

For more information about integrating Zabbix functionality into your Python applications, see Python library for Zabbix.

Note:
User access in Zabbix, including both configuration and historical data, depends on the user type, the assigned user role, and user groups.

Structure The API consists of a number of methods that are nominally grouped into separate APIs. Each of the methods performs one specific task. For example, the `host.create` method belongs to the *host* API and is used to create new hosts. Historically, APIs are sometimes referred to as "classes".

Most APIs contain at least four methods: `get`, `create`, `update` and `delete` for retrieving, creating, updating and deleting data respectively, but some APIs may provide a totally different set of methods.

Performing requests Once you have set up the frontend, you can use remote HTTP requests to call the API. To do that, you need to send HTTP POST requests to the `api_jsonrpc.php` file located in the frontend directory.
 For example, if your Zabbix frontend is installed under `https://example.com/zabbix`, an HTTP request to call the `apiinfo.version` method may look like this:

```
curl --request POST \
  --url 'https://example.com/zabbix/api_jsonrpc.php' \
  --header 'Content-Type: application/json-rpc' \
  --data '{"jsonrpc":"2.0","method":"apiinfo.version","params":{},"id":1}'
```

The request must have the `Content-Type` header set to one of these values: `application/json-rpc`, `application/json` or `application/jsonrequest`.

The request object must contain the following properties:

- `jsonrpc` - the version of the JSON-RPC protocol used by the API (Zabbix API implements JSON-RPC version 2.0);
- `method` - the API method being called;
- `params` - the parameters that will be passed to the API method;
- `id` - an arbitrary identifier of the request (if omitted, the API treats the request as a notification).

If the request is correct, the response returned by the API should look like this:

```
{
    "jsonrpc": "2.0",
    "result": "8.2.0",
    "id": 1
}
```

The response object, in turn, contains the following properties:

- `jsonrpc` - the version of the JSON-RPC protocol;
- `result` - the data returned by the method;
- `id` - an identifier of the corresponding request.

Authentication To access any data in Zabbix, you need to either:

- Use an existing API token created in the Zabbix frontend, *Users* > *API tokens* section, or created using the Token API.
- Use an authentication token obtained with the `user.login` method.

For example, if you wanted to obtain a new authentication token by logging in as a standard *Admin* user, then a JSON request would look like this:

```
curl --request POST \
  --url 'https://example.com/zabbix/api_jsonrpc.php' \
  --header 'Content-Type: application/json-rpc' \
  --data '{"jsonrpc":"2.0","method":"user.login","params":{"username":"Admin","password":"zabbix"},"id":1}
```

If you provided the credentials correctly, the response returned by the API should contain the user authentication token:

```
{
    "jsonrpc": "2.0",
    "result": "0424bd59b807674191e7d77572075f33",
    "id": 1
}
```

Authorization methods By "Authorization" header

All API requests require an authentication or an API token. You can provide the credentials by using the Authorization header in the request:

```
curl --request POST \
  --url 'https://example.com/zabbix/api_jsonrpc.php' \
  --header 'Authorization: Bearer 0424bd59b807674191e7d77572075f33'
```

> **Attention:**
> If you are experiencing authentication issues, see Authorization header forwarding.

Zabbix API accepts headers in a case-insensitive way (e.g., `authorization`, `Authorization`, and `AUTHORIZATION` are treated the same).

The Authorization header is supported in cross-origin requests (CORS).

By Zabbix cookie

A *"zbx_session"* cookie is used to authorize an API request from Zabbix UI performed using JavaScript (from a module or a custom widget).

Example workflow The following section walks you through several usage examples in greater detail.

Retrieving hosts Now you have a valid user authentication token (represented as a variable in the following examples) that can be used to access the data in Zabbix. For example, you can use the `host.get` method to retrieve the IDs, host names and interfaces of all the configured hosts:

Request:

```
curl --request POST \
  --url 'https://example.com/zabbix/api_jsonrpc.php' \
  --header 'Authorization: Bearer ${AUTHORIZATION_TOKEN}' \
  --header 'Content-Type: application/json-rpc' \
  --data @data.json
```

> **Note:**
> `data.json` is a file that contains a JSON query. Instead of a file, you can pass the query in the `--data` argument.

`data.json`

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": [
            "hostid",
            "host"
        ],
        "selectInterfaces": [
            "interfaceid",
            "ip"
        ]
    },
    "id": 2
}
```

The response object will contain the requested data about the hosts:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostid": "10084",
            "host": "Zabbix server",
            "interfaces": [
                {
```

```
                    "interfaceid": "1",
                    "ip": "127.0.0.1"
                }
            ]
        }
    ],
    "id": 2
}
```

Note:
For performance reasons it is always recommended to list the object properties you want to retrieve. Thus, you will avoid retrieving everything.

Creating a new item Now, create a new item on the host "Zabbix server" using the data you have obtained from the previous `host.get` request. This can be done using the `item.create` method:

```
curl --request POST \
  --url 'https://example.com/zabbix/api_jsonrpc.php' \
  --header 'Authorization: Bearer ${AUTHORIZATION_TOKEN}' \
  --header 'Content-Type: application/json-rpc' \
  --data '{"jsonrpc":"2.0","method":"item.create","params":{"name":"Free disk space on /home/joe/","key_":
```

A successful response will contain the ID of the newly created item, which can be used to reference the item in the following requests:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "24759"
        ]
    },
    "id": 3
}
```

Note:
The `item.create` method as well as other *create methods* can also accept arrays of objects and create multiple items with one API call.

Creating multiple triggers Thus, if *create methods* accept arrays, you can add multiple triggers, for example, this one:

```
curl --request POST \
  --url 'https://example.com/zabbix/api_jsonrpc.php' \
  --header 'Authorization: Bearer ${AUTHORIZATION_TOKEN}' \
  --header 'Content-Type: application/json-rpc' \
  --data '{"jsonrpc":"2.0","method":"trigger.create","params":[{"description":"Processor load is too high
```

The successful response will contain the IDs of the newly created triggers:

```
{
    "jsonrpc": "2.0",
    "result": {
        "triggerids": [
            "17369",
            "17370"
        ]
    },
    "id": 4
}
```

Updating an item Enable an item by setting its status to 0:

```
curl --request POST \
  --url 'https://example.com/zabbix/api_jsonrpc.php' \
  --header 'Authorization: Bearer ${AUTHORIZATION_TOKEN}' \
  --header 'Content-Type: application/json-rpc' \
  --data '{"jsonrpc":"2.0","method":"item.update","params":{"itemid":"10092","status":0},"id":5}'
```

The successful response will contain the ID of the updated item:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "10092"
        ]
    },
    "id": 5
}
```

Note:
The `item.update` method as well as other *update methods* can also accept arrays of objects and update multiple items with one API call.

Updating multiple triggers Enable multiple triggers by setting their status to 0:

```
curl --request POST \
  --url 'https://example.com/zabbix/api_jsonrpc.php' \
  --header 'Authorization: Bearer ${AUTHORIZATION_TOKEN}' \
  --header 'Content-Type: application/json-rpc' \
  --data '{"jsonrpc":"2.0","method":"trigger.update","params":[{"triggerid":"13938","status":0},{"triggeri
```

The successful response will contain the IDs of the updated triggers:

```
{
    "jsonrpc": "2.0",
    "result": {
        "triggerids": [
            "13938",
            "13939"
        ]
    },
    "id": 6
}
```

Error handling Up to the present moment, everything you have tried has worked fine. But what would happen if you tried making an incorrect call to the API? Try to create another host by calling `host.create` but omitting the mandatory `groups` parameter:

```
curl --request POST \
  --url 'https://example.com/zabbix/api_jsonrpc.php' \
  --header 'Authorization: Bearer ${AUTHORIZATION_TOKEN}' \
  --header 'Content-Type: application/json-rpc' \
  --data '{"jsonrpc":"2.0","method":"host.create","params":{"host":"Linux server","interfaces":[{"type":1,
```

The response will then contain an error message:

```
{
    "jsonrpc": "2.0",
    "error": {
        "code": -32602,
        "message": "Invalid params.",
        "data": "No groups for host \"Linux server\"."
    },
    "id": 7
}
```

If an error has occurred, instead of the `result` property, the response object will contain the `error` property with the following data:

- `code` - an error code;
- `message` - a short error summary;
- `data` - a more detailed error message.

Errors can occur in various cases, such as, using incorrect input values, a session timeout or trying to access non-existing objects. Your application should be able to gracefully handle these kinds of errors.

API versions To simplify API versioning, since Zabbix 2.0.4, the version of the API matches the version of Zabbix itself. You can use the `apiinfo.version` method to find out the version of the API you are working with. This can be useful for adjusting your application to use version-specific features.

Zabbix guarantees feature backward compatibility inside a major version. When making backward incompatible changes between major releases, Zabbix usually leaves the old features as deprecated in the next release, and only removes them in the release after that. Occasionally, Zabbix may remove features between major releases without providing any backward compatibility. It is important that you never rely on any deprecated features and migrate to newer alternatives as soon as possible.

> **Note:**
> You can follow all the changes made to the API in the API changelog.

Further reading Now, you have enough knowledge to start working with the Zabbix API, however, do not stop here. For further reading you are advised to have a look at the list of available APIs.

Method reference

This section provides an overview of the functions provided by the Zabbix API and will help you find your way around the available classes and methods.

Monitoring The Zabbix API allows you to access history and other data gathered during monitoring.

Dashboards

Manage dashboards and make scheduled reports based on them.

Dashboard API | Template dashboard API | Report API

High availability cluster

Retrieve a list of server nodes and their status.

High availability cluster API

History

Retrieve historical values gathered by Zabbix monitoring processes for presentation or further processing.

History API

Trends

Retrieve trend values calculated by Zabbix server for presentation or further processing.

Trend API

Events

Retrieve events generated by triggers, network discovery and other Zabbix systems for more flexible situation management or third-party tool integration.

Event API

Problems

Retrieve problems according to the given parameters.

Problem API

Maps

Configure maps to create detailed dynamic representations of your IT infrastructure.

Map API

Tasks

Interact with Zabbix server task manager, creating tasks and retrieving response.

Task API

APM The Zabbix API allows you to access telemetry data collected for application performance monitoring (APM).

Traces

Retrieve traces and the spans they consist of.

APM trace API | APM span API

Metrics

Retrieve metric data points.

APM metric API

Logs

Retrieve log records.

APM log API

Services The Zabbix API allows you to access data gathered during service monitoring.

Service Level Agreement

Define Service Level Objectives (SLO), retrieve detailed Service Level Indicators (SLI) information about service performance.

SLA API

Services

Manage services for service-level monitoring and retrieve detailed SLA information about any service.

Service API

Data collection The Zabbix API allows you to manage the configuration of your monitoring system.

Hosts and host groups

Manage host groups, hosts and everything related to them, including host interfaces, host macros and maintenance periods.

Host API | Host group API | User macro API | Value map API | Maintenance API

Items

Define items to monitor.

Item API

Triggers

Configure triggers to notify you about problems in your system. Manage trigger dependencies.

Trigger API

Graphs

Edit graphs or separate graph items for better presentation of the gathered data.

Graph API | Graph item API

Templates and template groups

Manage templates and link them to hosts or other templates.

Template API | Template group API | Value map API

Low-level discovery

Configure low-level discovery rules as well as item, trigger and graph prototypes to monitor dynamic entities.

LLD rule API | Item prototype API | Trigger prototype API | Graph prototype API | Host prototype API

Complex event processing

Create complex event processing rules.

Complex event processing API

Event correlation

Create custom event correlation rules.

Correlation API

Network discovery

Manage network-level discovery rules to automatically find and monitor new hosts. Gain full access to information about discovered services and hosts.

Discovery rule API | Discovery check API | Discovered host API | Discovered service API

Export and import

Export and import Zabbix configuration data for configuration backups, migration or large-scale configuration updates.

Configuration API

Web monitoring

Configure web scenarios to monitor your web applications and services.

Web scenario API

Alerts The Zabbix API allows you to manage the actions and alerts of your monitoring system.

Actions and alerts

Define actions and operations to notify users about certain events or automatically execute remote commands. Gain access to information about generated alerts and their receivers.

Action API | Alert API

Media types

Configure media types and multiple ways users will receive alerts.

Media type API

Scripts

Configure and execute scripts to help you with your daily tasks.

Script API

Users The Zabbix API allows you to manage users of your monitoring system.

Users and user groups

Add users that will have access to Zabbix, assign them to user groups and grant permissions. Make roles for granular management of user rights.

User API | User group API | User directory API | User role API

API Tokens

Manage authorization tokens.

Token API

Authentication

Change authentication configuration options.

Authentication API

Devices

Manage mobile devices for use with Zabbix mobile app.

Device API

Administration With the Zabbix API you can change administration settings of your monitoring system.

General

Change certain global configuration options.

Autoregistration API | Icon map API | Image API | Settings API | Regular expression API | Module API | Connector API

Audit log

Track configuration changes each user has done.

Audit log API

Housekeeping

Configure housekeeping.

Housekeeping API

Proxies and proxy groups

Manage the proxies used in your distributed monitoring setup.

Proxy API | Proxy group API

Macros

Manage macros.

User macro API

Zabbix information Retrieve the Zabbix API version or the time of the last update from Zabbix server.

API info API | Server info API

Action

This class is designed to work with actions.

Object references:

- Action
- Action operation
  - Action operation message
  - Action operation condition
- Action recovery operation
- Action update operation
- Action filter
  - Action filter condition

Available methods:

- action.create - create new actions
- action.delete - delete actions
- action.get - retrieve actions
- action.update - update actions

Action object

The following objects are directly related to the `action` API.

Action

The action object has the following properties.

Property	Type	Description
actionid	ID	ID of the action. Property behavior: - *read-only* - *required* for update operations

Property	Type	Description
esc_period	string	Default operation step duration. Must be at least 60 seconds. Accepts a number of seconds (e.g., 60, 300), a value with a time suffix (e.g., 30s, 1m, 2h), or a user macro. Property behavior: - *supported* if `eventsource` is set to "event created by a trigger", "internal event", or "event created on service status update"
eventsource	integer	Type of events that the action will handle. Refer to the event `source` property for a list of supported event types. Property behavior: - *constant* - *required* for create operations
name	string	Name of the action. Property behavior: - *required* for create operations
status	integer	Whether the action is enabled or disabled. Possible values: 0 - *(default)* enabled; 1 - disabled.
pause_symptoms	integer	Whether to pause escalation if event is a symptom event. Possible values: 0 - Don't pause escalation for symptom problems; 1 - *(default)* Pause escalation for symptom problems. Property behavior: - *supported* if `eventsource` is set to "event created by a trigger"
pause_suppressed	integer	Whether to pause escalation during maintenance periods or not. Possible values: 0 - Don't pause escalation; 1 - *(default)* Pause escalation. Property behavior: - *supported* if `eventsource` is set to "event created by a trigger"
notify_if_canceled	integer	Whether to notify when escalation is canceled. Possible values: 0 - Don't notify when escalation is canceled; 1 - *(default)* Notify when escalation is canceled. Property behavior: - *supported* if `eventsource` is set to "event created by a trigger"

Action operation

The action operation object defines an operation that will be performed when an action is executed. It has the following properties.

Property	Type	Description
operationtype	integer	Type of operation. Possible values: 0 - send message; 1 - global script; 2 - add host; 3 - remove host; 4 - add to host group; 5 - remove from host group; 6 - link template; 7 - unlink template; 8 - enable host; 9 - disable host; 10 - set host inventory mode; 13 - add host tags; 14 - remove host tags. Possible values if `eventsource` of Action object is set to ”event created by a trigger” or ”event created on service status update”: 0 - ”send message”; 1 - ”global script”. Possible values if `eventsource` of Action object is set to ”internal event”: 0 - ”send message”. Property behavior: - *required*
esc_period	string	Duration of an escalation step in seconds. Must be greater than 60 seconds. Accepts a number of seconds (e.g., 60, 300), a value with a time suffix (e.g., 30s, 1m, 2h), or a user macro. If set to 0 or 0s, the default action escalation period will be used. Default: 0s. Property behavior: - *supported* if `eventsource` of Action object is set to ”event created by a trigger”, ”internal event”, or ”event created on service status update”
esc_step_from	integer	Step to start escalation from. Default: 1. Property behavior: - *supported* if `eventsource` of Action object is set to ”event created by a trigger”, ”internal event”, or ”event created on service status update”
esc_step_to	integer	Step to end escalation at. Default: 1. Property behavior: - *supported* if `eventsource` of Action object is set to ”event created by a trigger”, ”internal event”, or ”event created on service status update”
evaltype	integer	Operation condition evaluation method. Possible values: 0 - *(default)* And/Or; 1 - And; 2 - Or.

Property	Type	Description
opcommand	object	Global script to execute. The global script must have the `scriptid` property defined. Property behavior: - *required* if `operationtype` is set to ”global script”
opcommand_grp	array	Host groups to run global scripts on. The host groups must have the `groupid` property defined. Property behavior: - *required* if `operationtype` is set to ”global script” and `opcommand_hst` is not set
opcommand_hst	array	Host to run global scripts on. The hosts must have the `hostid` property defined. Property behavior: - *required* if `operationtype` is set to ”global script” and `opcommand_grp` is not set
opconditions	array	Operation conditions used for trigger actions. The operation condition object is described in detail below.
opgroup	array	Host groups to add hosts to. The host groups must have the `groupid` property defined. Property behavior: - *required* if `operationtype` is set to ”add to host group” or ”remove from host group”
opmessage	object	Object containing the data about the message sent by the operation. The operation message object is described in detail below. Property behavior: - *required* if `operationtype` is set to ”send message”
opmessage_grp	array	User groups to send messages to. The user groups must have the `usrgrpid` property defined. Property behavior: - *required* if `operationtype` is set to ”send message” and `opmessage_usr` is not set
opmessage_usr	array	Users to send messages to. The users must have the `userid` property defined. Property behavior: - *required* if `operationtype` is set to ”send message” and `opmessage_grp` is not set
optemplate	array	Templates to link to the hosts. The templates must have the `templateid` property defined. Property behavior: - *required* if `operationtype` is set to ”link template” or ”unlink template”

Property	Type	Description
opinventory	object	Inventory mode set host to. The inventory must have the `inventory_mode` property defined. Property behavior: - *required* if `operationtype` is set to ”set host inventory mode”
optag	array	Host tags to add or remove. Tags must have `tag` property defined. The `value` property is optional. Property behavior: - *supported* if `operationtype` is set to ”add host tags” or ”remove host tags”.

Action operation message

The operation message object contains data about the message that will be sent by the operation. It has the following properties.

Property	Type	Description
default_msg	integer	Whether to use the default action message text and subject. Possible values: 0 - use the data from the operation; 1 - *(default)* use the data from the media type.
mediatypeid	ID	ID of the media type that will be used to send the message. Property behavior: - *supported* if `operationtype` of Action operation object, Action recovery operation object, or Action update operation object is set to ”send message”, or if `operationtype` of Action update operation object is set to ”notify all involved”
message	string	Operation message text. Property behavior: - *supported* if `default_msg` is set to ”use the data from the operation”
subject	string	Operation message subject. Property behavior: - *supported* if `default_msg` is set to ”use the data from the operation”

Action operation condition

The action operation condition object defines a condition that must be met to perform the current operation. It has the following properties.

Property	Type	Description
conditiontype	integer	Type of condition. Possible values: 14 - event acknowledged. Property behavior: - *required*
value	string	Value to compare with. Property behavior: - *required*

Property	Type	Description
operator	integer	Condition operator. Possible values: 0 - *(default)* =

The following operators and values are supported for each operation condition type.

Condition	Condition name	Supported operators	Expected value
14	Event acknowledged	=	Whether the event is acknowledged. Possible values: 0 - not acknowledged; 1 - acknowledged.

Action recovery operation

The action recovery operation object defines an operation that will be performed when a problem is resolved. Recovery operations are possible **only** for trigger, internal and service actions. It has the following properties.

Property	Type	Description
operationtype	integer	Type of operation. Possible values if `eventsource` of Action object is set to ”event created by a trigger” or ”event created on service status update”: 0 - send message; 1 - global script; 11 - notify all involved. Possible values if `eventsource` of Action object is set to ”internal event”: 0 - send message; 11 - notify all involved. Property behavior: - *required*
opcommand	object	Global script to execute. The global script must have the `scriptid` property defined. Property behavior: - *required* if `operationtype` is set to ”global script”
opcommand_grp	array	Host groups to run global scripts on. The host groups must have the `groupid` property defined. Property behavior: - *required* if `eventsource` of Action object is set to ”event created by a trigger”, and `operationtype` is set to ”global script”, and `opcommand_hst` is not set
opcommand_hst	array	Host to run global scripts on. The hosts must have the `hostid` property defined. Property behavior: - *required* if `eventsource` of Action object is set to ”event created by a trigger”, and `operationtype` is set to ”global script”, and `opcommand_grp` is not set

Property	Type	Description
opmessage	object	Object containing the data about the message sent by the recovery operation. The operation message object is described in detail above. Property behavior: - *required* if `operationtype` is set to ”send message”
opmessage_grp	array	User groups to send messages to. The user groups must have the `usrgrpid` property defined. Property behavior: - *required* if `operationtype` is set to ”send message” and `opmessage_usr` is not set
opmessage_usr	array	Users to send messages to. The users must have the `userid` property defined. Property behavior: - *required* if `operationtype` is set to ”send message” and `opmessage_grp` is not set

Action update operation

The action update operation object defines an operation that will be performed when a problem is updated (commented upon, acknowledged, severity changed, or manually closed). Update operations are possible **only** for trigger and service actions. It has the following properties.

Property	Type	Description
operationtype	integer	Type of operation. Possible values: 0 - send message; 1 - global script; 12 - notify all involved. Property behavior: - *required*
opcommand	object	Global script to execute. The global script must have the `scriptid` property defined. Property behavior: - *required* if `operationtype` is set to ”global script”
opcommand_grp	array	Host groups to run global scripts on. The host groups must have the `groupid` property defined. Property behavior: - *required* if `eventsource` of Action object is set to ”event created by a trigger”, and `operationtype` is set to ”global script”, and `opcommand_hst` is not set
opcommand_hst	array	Host to run global scripts on. The hosts must have the `hostid` property defined. Property behavior: - *required* if `eventsource` of Action object is set to ”event created by a trigger”, and `operationtype` is set to ”global script”, and `opcommand_grp` is not set

Property	Type	Description
opmessage	object	Object containing the data about the message sent by the update operation. The operation message object is described in detail above.
opmessage_grp	array	User groups to send messages to. The user groups must have the `usrgrpid` property defined. Property behavior: - *required* if `operationtype` is set to "send message" and `opmessage_usr` is not set
opmessage_usr	array	Users to send messages to. The users must have the `userid` property defined. Property behavior: - *required* if `operationtype` is set to "send message" and `opmessage_grp` is not set

Action filter

The action filter object defines a set of conditions that must be met to perform the configured action operations. It has the following properties.

Property	Type	Description
conditions	array	Set of filter conditions to use for filtering results. The conditions will be sorted in the order of their placement in the formula. Property behavior: - *required*
evaltype	integer	Filter condition evaluation method. Possible values: 0 - And/Or; 1 - And; 2 - Or; 3 - Custom expression. Property behavior: - *required*
eval_formula	string	Generated expression that will be used for evaluating filter conditions. The expression contains IDs that reference specific filter conditions by its `formulaid`. The value of `eval_formula` is equal to the value of `formula` for filters with a custom expression. Property behavior: - *read-only*
formula	string	User-defined expression to be used for evaluating conditions of filters with a custom expression. The expression must contain IDs that reference specific filter conditions by its `formulaid`. The IDs used in the expression must exactly match the ones defined in the filter conditions: no condition can remain unused or omitted. Property behavior: - *required* if `evaltype` is set to "custom expression"

Action filter condition

The action filter condition object defines a specific condition that must be checked before running the action operations.

Property	Type	Description
conditiontype	integer	Type of condition. Possible values if `eventsource` of Action object is set to ”event created by a trigger”: 0 - host group; 1 - host; 2 - trigger; 3 - event name; 4 - trigger severity; 6 - time period; 13 - host template; 16 - problem is suppressed; 25 - event tag; 26 - event tag value. Possible values if `eventsource` of Action object is set to ”event created by a discovery rule”: 7 - host IP; 8 - discovered service type; 9 - discovered service port; 10 - discovery status; 11 - uptime or downtime duration; 12 - received value; 18 - discovery rule; 19 - discovery check; 20 - proxy; 21 - discovery object; 29 - proxy group. Possible values if `eventsource` of Action object is set to ”event created by active agent autoregistration”: 20 - proxy; 22 - host name; 24 - host metadata; 29 - proxy group. Possible values if `eventsource` of Action object is set to ”internal event”: 0 - host group; 1 - host; 13 - host template; 23 - event type; 25 - event tag; 26 - event tag value. Possible values if `eventsource` of Action object is set to ”event created on service status update”: 25 - event tag; 26 - event tag value; 27 - service; 28 - service name. Property behavior: - *required*
value	string	Value to compare with. Property behavior: - *required*

Property	Type	Description
value2	string	Secondary value to compare with. Property behavior: - *required* if `eventsource` of Action object is set to ”event created by a trigger”, `conditiontype` is set to any possible value for trigger actions, and the type of condition (see below) is ”26” - *required* if `eventsource` of Action object is set to ”internal event”, `conditiontype` is set to any possible value for internal actions, and the type of condition (see below) is ”26” - *required* if `eventsource` of Action object is set to ”event created on service status update”, `conditiontype` is set to any possible value for service actions, and the type of condition (see below) is ”26”
formulaid	string	Arbitrary unique ID that is used to reference the condition from a custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward.
operator	integer	Condition operator. Possible values: 0 - *(default)* equals; 1 - does not equal; 2 - contains; 3 - does not contain; 4 - in; 5 - is greater than or equals; 6 - is less than or equals; 7 - not in; 8 - matches; 9 - does not match; 10 - Yes; 11 - No.

Note:
To better understand how to use filters with various types of expressions, see examples on the action.get and action.create method pages.

The following operators and values are supported for each condition type.

Condition	Condition name	Supported operators	Expected value
0	Host group	equals, does not equal	Host group ID.
1	Host	equals, does not equal	Host ID.
2	Trigger	equals, does not equal	Trigger ID.
3	Event name	contains, does not contain	Event name.
4	Trigger severity	equals, does not equal, is greater than or equals, is less than or equals	Trigger severity. Refer to the trigger `severity` property for a list of supported trigger severities.
5	Trigger value	equals	Trigger value. Refer to the trigger `value` property for a list of supported trigger values.
6	Time period	in, not in	Time when the event was triggered as a time period.
7	Host IP	equals, does not equal	One or several IP ranges to check, separated by commas. Refer to the network discovery configuration section for more information on supported formats of IP ranges.

Condition	Condition name	Supported operators	Expected value
8	Discovered service type	equals, does not equal	Type of discovered service. The type of service matches the type of the discovery check used to detect the service. Refer to the discovery check `type` property for a list of supported types.
9	Discovered service port	equals, does not equal	One or several port ranges, separated by commas.
10	Discovery status	equals	Status of a discovered object. Possible values: 0 - host or service up; 1 - host or service down; 2 - host or service discovered; 3 - host or service lost.
11	Uptime or downtime duration	is greater than or equals, is less than or equals	Time indicating how long has the discovered object been in the current status in seconds.
12	Received values	equals, does not equal, is greater than or equals, is less than or equals, contains, does not contain	Value returned when performing a Zabbix agent, SNMPv1, SNMPv2 or SNMPv3 discovery check.
13	Host template	equals, does not equal	Linked template ID.
16	Problem is suppressed	Yes, No	No value required: using the ”Yes” operator means that problem must be suppressed, ”No” - not suppressed.
18	Discovery rule	equals, does not equal	ID of the discovery rule.
19	Discovery check	equals, does not equal	ID of the discovery check.
20	Proxy	equals, does not equal	ID of the proxy.
21	Discovery object	equals	Type of object that triggered the discovery event. Possible values: 1 - discovered host; 2 - discovered service.
22	Host name	contains, does not contain, matches, does not match	Host name. Using a regular expression is supported for operators *matches* and *does not match* in autoregistration conditions.
23	Event type	equals	Specific internal event. Possible values: 0 - item in ”not supported” state; 1 - item in ”normal” state; 2 - LLD rule in ”not supported” state; 3 - LLD rule in ”normal” state; 4 - trigger in ”unknown” state; 5 - trigger in ”normal” state.
24	Host metadata	contains, does not contain, matches, does not match	Metadata of the auto-registered host. Using a regular expression is supported for operators *matches* and *does not match*.
25	Tag	equals, does not equal, contains, does not contain	Event tag.

Condition	Condition name	Supported operators	Expected value
26	Tag value	equals, does not equal, contains, does not contain	Event tag value.
27	Service	equals, does not equal	Service ID.
28	Service name	equals, does not equal	Service name.
29	Proxy group	equals, does not equal	ID of the proxy group.

action.create

Description

`object action.create(object/array actions)`

This method allows you to create new actions.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Actions to create.

Additionally to the standard action properties, the method accepts the following parameters.

Parameter	Type	Description
filter	object	Action filter object for the action.
operations	array	Action operations to create for the action.
recovery_operations	array	Action recovery operations to create for the action.
update_operations	array	Action update operations to create for the action.

Return values

`(object)` Returns an object containing the IDs of the created actions under the `actionids` property. The order of the returned IDs matches the order of the passed actions.

Examples

Create a trigger action

Create a trigger action that will begin once a trigger (with the word ”memory” in its name) from host ”10084” goes into a PROBLEM state. The action will have 4 configured operations. The first and immediate operation will send a message to all users in user group ”7” via media type ”1”. If the event is not resolved in 30 minutes, the second operation will run script ”5” (script with scope ”Action operation”) on all hosts in group ”2”. If the event is resolved, a recovery operation will notify all users who received any messages regarding the problem. If the event is updated, an acknowledge/update operation will notify (with a custom subject and message) all users who received any messages regarding the problem.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "action.create",
    "params": {
        "name": "Trigger action",
        "eventsource": 0,
        "esc_period": "30m",
        "filter": {
            "evaltype": 0,
            "conditions": [
```

```
                {
                    "conditiontype": 1,
                    "operator": 0,
                    "value": "10084"
                },
                {
                    "conditiontype": 3,
                    "operator": 2,
                    "value": "memory"
                }
            ]
        },
        "operations": [
            {
                "operationtype": 0,
                "esc_step_from": 1,
                "esc_step_to": 1,
                "opmessage_grp": [
                    {
                        "usrgrpid": "7"
                    }
                ],
                "opmessage": {
                    "default_msg": 1,
                    "mediatypeid": "1"
                }
            },
            {
                "operationtype": 1,
                "esc_step_from": 2,
                "esc_step_to": 2,
                "opconditions": [
                    {
                        "conditiontype": 14,
                        "operator": 0,
                        "value": "0"
                    }
                ],
                "opcommand_grp": [
                    {
                        "groupid": "2"
                    }
                ],
                "opcommand": {
                    "scriptid": "5"
                }
            }
        ],
        "recovery_operations": [
            {
                "operationtype": "11",
                "opmessage": {
                    "default_msg": 1
                }
            }
        ],
        "update_operations": [
            {
                "operationtype": "12",
                "opmessage": {
                    "default_msg": 0,
                    "message": "Custom update operation message body",
```

```
                    "subject": "Custom update operation message subject"
                }
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "actionids": [
            "17"
        ]
    },
    "id": 1
}
```

Create a discovery action

Create a discovery action that will link template ”10001” to discovered hosts.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "action.create",
    "params": {
        "name": "Discovery action",
        "eventsource": 1,
        "filter": {
            "evaltype": 0,
            "conditions": [
                {
                    "conditiontype": 21,
                    "operator": 0,
                    "value": "1"
                },
                {
                    "conditiontype": 10,
                    "operator": 0,
                    "value": "2"
                }
            ]
        },
        "operations": [
            {
                "operationtype": 6,
                "optemplate": [
                    {
                        "templateid": "10001"
                    }
                ]
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
```

```
        "actionids": [
            "18"
        ]
    },
    "id": 1
}
```

Using a custom expression filter

Create a trigger action that uses a custom expression - ”A and (B or C)” - for evaluating action conditions. Once a trigger with a severity higher or equal to ”Warning” from host ”10084” or host ”10106” goes into a PROBLEM state, the action will send a message to all users in user group ”7” via media type ”1”. The formula IDs ”A”, ”B” and ”C” have been chosen arbitrarily.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "action.create",
    "params": {
        "name": "Trigger action",
        "eventsource": 0,
        "esc_period": "15m",
        "filter": {
            "evaltype": 3,
            "formula": "A and (B or C)",
            "conditions": [
                {
                    "conditiontype": 4,
                    "operator": 5,
                    "value": "2",
                    "formulaid": "A"
                },
                {
                    "conditiontype": 1,
                    "operator": 0,
                    "value": "10084",
                    "formulaid": "B"
                },
                {
                    "conditiontype": 1,
                    "operator": 0,
                    "value": "10106",
                    "formulaid": "C"
                }
            ]
        },
        "operations": [
            {
                "operationtype": 0,
                "esc_step_from": 1,
                "esc_step_to": 1,
                "opmessage_grp": [
                    {
                        "usrgrpid": "7"
                    }
                ],
                "opmessage": {
                    "default_msg": 1,
                    "mediatypeid": "1"
                }
            }
        ]
    },
    "id": 1
```

```
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "actionids": [
            "18"
        ]
    },
    "id": 1
}
```

Create agent autoregistration rule

Create an autoregistration action that adds a host to host group ”2” when the host name contains ”SRV” or metadata contains ”AlmaLinux”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "action.create",
    "params": {
        "name": "Register Linux servers",
        "eventsource": "2",
        "filter": {
            "evaltype": "2",
            "conditions": [
                {
                    "conditiontype": "22",
                    "operator": "2",
                    "value": "SRV"
                },
                {
                    "conditiontype": "24",
                    "operator": "2",
                    "value": "AlmaLinux"
                }
            ]
        },
        "operations": [
            {
                "operationtype": "4",
                "opgroup": [
                    {
                        "groupid": "2"
                    }
                ]
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "actionids": [
            19
        ]
    },
    "id": 1
```

```
}
```

Create agent autoregistration rule with host tags

Create an autoregistration action that adds a host to host group "2" and adds two host tags.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "action.create",
    "params": {
        "name": "Register Linux servers with tags",
        "eventsource": "2",
        "operations": [
            {
                "operationtype": "4",
                "opgroup": [
                    {
                        "groupid": "2"
                    }
                ]
            },
            {
                "operationtype": "13",
                "optag": [
                    {
                        "tag": "location",
                        "value": "office"
                    },
                    {
                        "tag": "city",
                        "value": "Riga"
                    }
                ]
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "actionids": [
            20
        ]
    },
    "id": 1
}
```

See also

- Action filter
- Action operation
- Script

Source

CAction::create() in *ui/include/classes/api/services/CAction.php*.

action.delete

Description

`object action.delete(array actionIds)`

This method allows you to delete actions.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.
 The user will not be able to delete an action that contains an inaccessible proxy or proxy group in its filter conditions. See user group configuration (*Proxy access list*) for more information.

Parameters

`(array)` IDs of the actions to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted actions under the `actionids` property.

Examples

Delete multiple actions

Delete two actions.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "action.delete",
    "params": [
        "17",
        "18"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "actionids": [
            "17",
            "18"
        ]
    },
    "id": 1
}
```

Source

CAction::delete() in *ui/include/classes/api/services/CAction.php*.

action.get

Description

`integer/array action.get(object parameters)`

The method allows you to retrieve actions according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
actionids	ID/array	Return only actions with the given IDs.
groupids	ID/array	Return only actions that use the given host groups in action conditions.
hostids	ID/array	Return only actions that use the given hosts in action conditions.
triggerids	ID/array	Return only actions that use the given triggers in action conditions.
mediatypeids	ID/array	Return only actions that use the given media types to send messages.
usrgrpids	ID/array	Return only actions that are configured to send messages to the given user groups.
userids	ID/array	Return only actions that are configured to send messages to the given users.
scriptids	ID/array	Return only actions that are configured to run the given scripts.
selectFilter	query	Return a `filter` property with the action condition filter.
selectOperations	query	Return an `operations` property with action operations.
selectRecoveryOperations	query	Return a `recovery_operations` property with action recovery operations.
selectUpdateOperations	query	Return an `update_operations` property with action update operations.
sortfield	string/array	Sort the result by the given properties. Possible values: `actionid`, `name`, `status`.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve trigger actions

Retrieve all configured trigger actions together with action conditions and operations.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "action.get",
    "params": {
        "output": "extend",
        "selectOperations": "extend",
        "selectRecoveryOperations": "extend",
        "selectUpdateOperations": "extend",
        "selectFilter": "extend",
        "filter": {
            "eventsource": 0
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "actionid": "3",
            "name": "Report problems to Zabbix administrators",
            "eventsource": "0",
            "status": "1",
            "esc_period": "1h",
            "pause_suppressed": "1",
            "filter": {
                "evaltype": "0",
                "formula": "",
                "conditions": [],
                "eval_formula": ""
            },
            "operations": [
                {
                    "operationid": "3",
                    "actionid": "3",
                    "operationtype": "0",
                    "esc_period": "0",
                    "esc_step_from": "1",
                    "esc_step_to": "1",
                    "evaltype": "0",
                    "opconditions": [],
                    "opmessage": [
                        {
                            "default_msg": "1",
                            "subject": "",
                            "message": "",
                            "mediatypeid" => "0"
                        }
                    ],
                    "opmessage_grp": [
                        {
                            "usrgrpid": "7"
                        }
                    ]
                }
            ],
            "recovery_operations": [
                {
                    "operationid": "7",
                    "actionid": "3",
                    "operationtype": "11",
                    "evaltype": "0",
                    "opconditions": [],
                    "opmessage": {
                        "default_msg": "0",
                        "subject": "{TRIGGER.STATUS}: {TRIGGER.NAME}",
                        "message": "Trigger: {TRIGGER.NAME}\r\nTrigger state: {TRIGGER.STATUS}\r\nTrigger
                        "mediatypeid": "0"
                    }
                }
            ],
            "update_operations": [
                {
                    "operationid": "31",
                    "operationtype": "12",
                    "evaltype": "0",
                    "opmessage": {
```

```
                        "default_msg": "1",
                        "subject": "",
                        "message": "",
                        "mediatypeid": "0"
                    }
                },
                {
                    "operationid": "32",
                    "operationtype": "0",
                    "evaltype": "0",
                    "opmessage": {
                        "default_msg": "0",
                        "subject": "Updated: {TRIGGER.NAME}",
                        "message": "{USER.FULLNAME} updated problem at {EVENT.UPDATE.DATE} {EVENT.UPDATE.T
                        "mediatypeid": "1"
                    },
                    "opmessage_grp": [
                        {
                            "usrgrpid": "7"
                        }
                    ],
                    "opmessage_usr": []
                },
                {
                    "operationid": "33",
                    "operationtype": "1",
                    "evaltype": "0",
                    "opcommand": {
                        "scriptid": "3"
                    },
                    "opcommand_hst": [
                        {
                            "hostid": "10084"
                        }
                    ],
                    "opcommand_grp": []
                }
            ]
        }
    ],
    "id": 1
}
```


Retrieve discovery actions

Retrieve all configured discovery actions together with action conditions and operations. The filter uses the ”and” evaluation type, so the `formula` property is empty and `eval_formula` is generated automatically.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "action.get",
    "params": {
        "output": "extend",
        "selectOperations": "extend",
        "selectFilter": "extend",
        "filter": {
            "eventsource": 1
        }
    },
    "id": 1
}
```


Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "actionid": "2",
            "name": "Auto discovery. Linux servers.",
            "eventsource": "1",
            "status": "1",
            "esc_period": "0s",
            "pause_suppressed": "1",
            "filter": {
                "evaltype": "0",
                "formula": "",
                "conditions": [
                    {
                        "conditiontype": "10",
                        "operator": "0",
                        "value": "0",
                        "value2": "",
                        "formulaid": "B"
                    },
                    {
                        "conditiontype": "8",
                        "operator": "0",
                        "value": "9",
                        "value2": "",
                        "formulaid": "C"
                    },
                    {
                        "conditiontype": "12",
                        "operator": "2",
                        "value": "Linux",
                        "value2": "",
                        "formulaid": "A"
                    }
                ],
                "eval_formula": "A and B and C"
            },
            "operations": [
                {
                    "operationid": "1",
                    "actionid": "2",
                    "operationtype": "6",
                    "esc_period": "0s",
                    "esc_step_from": "1",
                    "esc_step_to": "1",
                    "evaltype": "0",
                    "opconditions": [],
                    "optemplate": [
                        {
                            "templateid": "10001"
                        }
                    ]
                },
                {
                    "operationid": "2",
                    "actionid": "2",
                    "operationtype": "4",
                    "esc_period": "0s",
                    "esc_step_from": "1",
                    "esc_step_to": "1",
                    "evaltype": "0",
```

```
                    "opconditions": [],
                    "opgroup": [
                        {
                            "groupid": "2"
                        }
                    ]
                }
            ]
        }
    ],
    "id": 1
}
```

See also

- Action filter
- Action operation

Source

CAction::get() in *ui/include/classes/api/services/CAction.php*.

action.update

Description

`object action.update(object/array actions)`

This method allows you to update existing actions.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Action properties to be updated.

The `actionid` property must be defined for each action, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard action properties, the method accepts the following parameters.

Parameter	Type	Description
filter	object	Action filter object to replace the current filter.
operations	array	Action operations to replace existing operations.
recovery_operations	array	Action recovery operations to replace existing recovery operations. Parameter behavior: - *supported* if `eventsource` of Action object is set to "event created by a trigger", "internal event", or "event created on service status update"
update_operations	array	Action update operations to replace existing update operations. Parameter behavior: - *supported* if `eventsource` of Action object is set to "event created by a trigger" or "event created on service status update"

Return values

`(object)` Returns an object containing the IDs of the updated actions under the `actionids` property.

Examples

Disable action

Disable an action, that is, set its status to "1".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "action.update",
    "params": {
        "actionid": "2",
        "status": "1"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "actionids": [
            "2"
        ]
    },
    "id": 1
}
```

See also

- Action filter
- Action operation

Source

CAction::update() in *ui/include/classes/api/services/CAction.php*.

Alert

This class is designed to work with alerts.

Object references:

- Alert

Available methods:

- alert.get - retrieve alerts

Alert object

The following objects are directly related to the `alert` API.

Alert

Note:
Alerts are created by Zabbix server and cannot be modified via the API.

The alert object contains information about whether certain action operations have been executed successfully. It has the following properties.

Property	Type	Description
alertid	ID	ID of the alert.
actionid	ID	ID of the action that generated the alert.
alerttype	integer	Alert type. Possible values: 0 - message; 1 - remote command.
clock	timestamp	Time when the alert was generated.

Property	Type	Description
error	string	Error text if there are problems sending a message or running a command.
esc_step	integer	Action escalation step during which the alert was generated.
eventid	ID	ID of the event that triggered the action.
mediatypeid	ID	ID of the media type that was used to send the message.
message	text	Message text. Property behavior: - *supported* if `alerttype` is set to ”message”
retries	integer	Number of times Zabbix tried to send the message.
sendto	string	Address, user name or other identifier of the recipient. Property behavior: - *supported* if `alerttype` is set to ”message”
status	integer	Status indicating whether the action operation has been executed successfully. Possible values if `alerttype` is set to ”message”: 0 - message not sent; 1 - message sent; 2 - failed after a number of retries; 3 - new alert is not yet processed by alert manager. Possible values if `alerttype` is set to ”remote command”: 0 - command not run; 1 - command run; 2 - tried to run the command on Zabbix agent, but it was unavailable.
subject	string	Message subject. Property behavior: - *supported* if `alerttype` is set to ”message”
userid	ID	ID of the user that the message was sent to.
p_eventid	ID	ID of problem event, which generated the alert.
acknowledgeid	ID	ID of acknowledgment, which generated the alert.

alert.get

Description

`integer/array alert.get(object parameters)`

The method allows you to retrieve alerts according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
alertids	ID/array	Return only alerts with the given IDs.
actionids	ID/array	Return only alerts generated by the given actions.
eventids	ID/array	Return only alerts generated by the given events.
groupids	ID/array	Return only alerts generated by objects from the given host groups.
hostids	ID/array	Return only alerts generated by objects from the given hosts.
mediatypeids	ID/array	Return only message alerts that used the given media types.
objectids	ID/array	Return only alerts generated by the given objects

Parameter	Type	Description
userids	ID/array	Return only message alerts that were sent to the given users.
eventobject	integer	Return only alerts generated by events related to objects of the given type. See event `object` for a list of supported object types. Default: 0 - trigger.
eventsource	integer	Return only alerts generated by events of the given type. See event `source` for a list of supported event types. Default: 0 - trigger events.
time_from	timestamp	Return only alerts that have been generated after the given time.
time_till	timestamp	Return only alerts that have been generated before the given time.
selectHosts	query	Return a `hosts` property with data of hosts that triggered the action operation.
selectMediatypes	query	Return a `mediatypes` property with an array of the media types that were used for the message alert. See `mediatype.get` for restrictions based on user type.
selectUsers	query	Return a `users` property with an array of the users that the message was addressed to. See `user.get` for restrictions based on user type.
sortfield	string/array	Sort the result by the given properties. Possible values: `alertid`, `clock`, `eventid`, `mediatypeid`, `sendto`, `status`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve alerts by action ID

Retrieve all alerts generated by action "3".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "alert.get",
    "params": {
        "output": "extend",
        "actionids": "3"
    },
```

```
    "id": 1
}
```

Response:


```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "alertid": "1",
            "actionid": "3",
            "eventid": "21243",
            "userid": "1",
            "clock": "1362128008",
            "mediatypeid": "1",
            "sendto": "support@company.com",
            "subject": "PROBLEM: Zabbix agent on Linux server is unreachable for 5 minutes: ",
            "message": "Trigger: Zabbix agent on Linux server is unreachable for 5 minutes: \nTrigger stat
            "status": "0",
            "retries": "3",
            "error": "",
            "esc_step": "1",
            "alerttype": "0",
            "p_eventid": "0",
            "acknowledgeid": "0"
        }
    ],
    "id": 1
}
```


See also

- Host
- Media type
- User

Source

CAlert::get() in *ui/include/classes/api/services/CAlert.php*.

API info

This class is designed to retrieve meta information about the API.

Available methods:

- apiinfo.version - retrieve the version of Zabbix API

apiinfo.version

Description

`string apiinfo.version(array)`

This method allows you to retrieve the version of the Zabbix API.

> **Attention:**
> This method is only available to unauthenticated users.

Parameters

`(array)` The method accepts an empty array.

Return values

`(string)` Returns the version of the Zabbix API.

Note:
Starting from Zabbix 2.0.4 the version of the API matches the version of Zabbix.

Examples

Retrieving the version of the API

Retrieve the version of the Zabbix API.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "apiinfo.version",
    "params": [],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": "8.2.0",
    "id": 1
}
```

Source

CAPIInfo::version() in *ui/include/classes/api/services/CAPIInfo.php*.

APM log

This class is designed to work with logs collected for application performance monitoring (APM).

Object references:

- APM log

Available methods:

- apmlog.get - retrieve log records

APM log object

The following objects are directly related to the `apmlog` API.

Note:
Each property corresponds to the column with the same name in the `otel_logs` table of the telemetry data source (e.g., `severity_text` property corresponds to the `SeverityText` column). For the list of columns, see Supported telemetry data columns.

APM log

The APM log object represents a log record (a text record of an event, such as a failed payment or a database timeout). A log record contains information about an event that occurred at a specific time, including its severity, message, and attributes. If a log record was produced during an operation that is part of a trace, it also contains the IDs of that trace and span.

The APM log object has the following properties:

Property	Type	Description
timestamp	integer	Unix time (in nanoseconds) when the event occurred. Property behavior: - *read-only*

Property	Type	Description
traceid	string	ID, if any, of the trace during which the log record was produced. Property behavior: - *read-only*
spanid	string	ID, if any, of the span during which the log record was produced. Property behavior: - *read-only*
trace_flags	integer	Trace flags of the trace context of the log record, as a bitmask. If the least significant bit is set (value 1), the trace is sampled. Property behavior: - *read-only*
severity_text	string	Severity level of the log record, as provided by the application (e.g., `INFO`, `WARN`, or `ERROR`). Property behavior: - *read-only*
severity_number	integer	Severity level of the log record, as a number. Possible values: 0 - UNSPECIFIED; 1-4 - TRACE; 5-8 - DEBUG; 9-12 - INFO; 13-16 - WARN; 17-20 - ERROR; 21-24 - FATAL. Property behavior: - *read-only*
service_name	string	Name of the service that produced the log record (e.g., `zabbix-frontend`), taken from the OpenTelemetry `service.name` resource attribute. Property behavior: - *read-only*
body	string	Message of the log record. Property behavior: - *read-only*
resource_schema_url	string	URL, if any, of the OpenTelemetry schema of the resource attributes. Property behavior: - *read-only*
resource_attributes	object	Key-value resource attributes of the log record (e.g., `host.name: eabf8c2a5ddc`). Property behavior: - *read-only*
scope_schema_url	string	URL, if any, of the OpenTelemetry schema of the instrumentation scope and its attributes. Property behavior: - *read-only*
scope_name	string	Name of the instrumentation scope that produced the log record. Property behavior: - *read-only*

Property	Type	Description
scope_version	string	Version, if any, of the instrumentation scope. Property behavior: - *read-only*
scope_attributes	object	Key-value attributes of the instrumentation scope. Property behavior: - *read-only*
log_attributes	object	Key-value log attributes of the log record (e.g., `log.file.path: /var/log/zabbix/zabbix_frontend.log`). Property behavior: - *read-only*
event_name	string	Name, if any, of the event type that the log record represents. Property behavior: - *read-only*

apmlog.get

Description

`integer/array apmlog.get(object parameters)`

The method allows you to retrieve log records according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
time_from	timestamp	Return only log records of events that occurred at or after the given time. Parameter behavior: - *required*
time_till	timestamp	Return only log records of events that occurred before the given time. Parameter behavior: - *required*
traceids	string/array	Return only log records produced during the given traces.
spanids	string/array	Return only log records produced during the given spans.
with_flags_on	integer	Return only log records that have all of the given trace flags set. To specify several flags, use the sum of their values.
with_flags_off	integer	Return only log records that have none of the given trace flags set. To specify several flags, use the sum of their values.

Parameter	Type	Description
resource_attributes	object/array	Return only log records with the given resource attributes. Format: `[{"key": "<key>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all log records. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
resource_attributes_evaltype	integer	Evaluation method for the conditions in `resource_attributes`. Possible values: 0 - *(default)* And/Or; 2 - Or.
scope_attributes	object/array	Return only log records with the given scope attributes. Format: `[{"key": "<key>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all log records. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
scope_attributes_evaltype	integer	Evaluation method for the conditions in `scope_attributes`. Possible values: 0 - *(default)* And/Or; 2 - Or.
log_attributes	object/array	Return only log records with the given log attributes. Format: `[{"key": "<key>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all log records. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
log_attributes_evaltype	integer	Evaluation method for the conditions in `log_attributes`. Possible values: 0 - *(default)* And/Or; 2 - Or.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Supported properties: `traceid`, `spanid`, `trace_flags`, `severity_text`, `severity_number`, `service_name`, `resource_schema_url`, `scope_schema_url`, `scope_name`, `scope_version`, `event_name`.

Parameter	Type	Description
search	object	Return results that match the given pattern (case-insensitive). Accepts an object, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a `LIKE "%...%"` search. Supported properties: `severity_text`, `service_name`, `body`, `resource_schema_url`, `scope_schema_url`, `scope_name`, `scope_version`, `event_name`.
sortfield	string/array	Sort the result by the given properties. Possible values: `timestamp`, `traceid`, `spanid`, `trace_flags`, `severity_text`, `severity_number`, `service_name`, `resource_schema_url`, `scope_schema_url`, `scope_name`, `scope_version`, `event_name`.
offset	integer	Number of records to skip before returning the result.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
limit	integer	
output	query	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving error logs of a trace

Retrieve log records with the ERROR severity (severity numbers 17-20) that are associated with the trace "8a4f1c7e2b9d4e6fa0c3b5d7e9f1a2c

Request:

```
{
    "jsonrpc": "2.0",
    "method": "apmlog.get",
    "params": {
        "output": "extend",
        "time_from": 1790848800,
        "time_till": 1790852400,
        "traceids": "8a4f1c7e2b9d4e6fa0c3b5d7e9f1a2c4",
        "filter": {
            "severity_number": [17, 18, 19, 20]
        },
        "sortfield": "timestamp",
        "sortorder": "ASC"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "timestamp": "1790849305091452376",
            "traceid": "8a4f1c7e2b9d4e6fa0c3b5d7e9f1a2c4",
```

```
            "spanid": "5d2e8f1a9c3b7046",
            "trace_flags": "1",
            "severity_text": "ERROR",
            "severity_number": "17",
            "service_name": "zabbix-frontend",
            "body": "mysqli_query(): (HY000/1205): Lock wait timeout exceeded; try restarting transaction"
            "resource_schema_url": "https://opentelemetry.io/schemas/1.26.0",
            "resource_attributes": {
                "service.name": "zabbix-frontend",
                "host.name": "6583ceeb5f94",
                "host.arch": "x86_64",
                "os.type": "linux"
            },
            "scope_schema_url": "",
            "scope_name": "zabbix-frontend",
            "scope_version": "",
            "scope_attributes": {},
            "log_attributes": {
                "log.file.path": "/var/log/zabbix/zabbix_frontend.log",
                "code.filepath": "/usr/share/zabbix/include/db.inc.php"
            },
            "event_name": ""
        }
    ],
    "id": 1
}
```


Searching log messages

Retrieve the time, service, and message of the 10 latest log records whose message contains ”timeout”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "apmlog.get",
    "params": {
        "output": ["timestamp", "service_name", "severity_text", "body"],
        "time_from": 1790848800,
        "time_till": 1790852400,
        "search": {
            "body": "timeout"
        },
        "sortfield": "timestamp",
        "sortorder": "DESC",
        "limit": 10
    },
    "id": 1
}
```

Response:


```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "timestamp": "1790849305091452376",
            "service_name": "zabbix-frontend",
            "severity_text": "ERROR",
            "body": "mysqli_query(): (HY000/1205): Lock wait timeout exceeded; try restarting transaction"
        },
        {
            "timestamp": "1790848961730284115",
            "service_name": "zabbix-frontend",
            "severity_text": "ERROR",
```

```
            "body": "mysqli_query(): (HY000/1205): Lock wait timeout exceeded; try restarting transaction"
        }
    ],
    "id": 1
}
```


See also

- APM trace
- APM span
- APM metric

Source

CApmLog::get() in *ui/include/classes/api/services/CApmLog.php*.

APM metric

This class is designed to work with metrics collected for application performance monitoring (APM).

Object references:

- APM metric
- Exemplar

Available methods:

- apmmetric.get - retrieve metric data points

APM metric object

The following objects are directly related to the `apmmetric` API.

Note:
Each property, except `type`, corresponds to a column in the metric tables (`otel_metrics_gauge`, `otel_metrics_sum`, `otel_metrics_histogram`, and `otel_metrics_exponential_histogram`) of the telemetry data source. Most properties correspond to the column with the same name (e.g., `metric_name` property corresponds to the `MetricName` column). Where the names differ, the data source columns are specified in the tables below. For the list of columns, see Supported telemetry data columns.

APM metric

The APM metric object represents a data point of a metric (the value, or aggregated values, of a measurement at a specific time).

The APM metric object has the following properties:

Property	Type	Description
type	integer	Type of the data point. Possible values: 0 - Gauge; 1 - Sum; 2 - Histogram; 3 - Exponential histogram. The type is determined by the table that holds the data point: `otel_metrics_gauge`, `otel_metrics_sum`, `otel_metrics_histogram`, or `otel_metrics_exponential_histogram`. Property behavior: - *read-only*

Property	Type	Description
resource_attributes	object	Key-value resource attributes of the data point (e.g., `host.name: eabf8c2a5ddc`). Property behavior: - *read-only*
resource_schema_url	string	URL, if any, of the OpenTelemetry schema of the resource attributes. Property behavior: - *read-only*
scope_name	string	Name of the instrumentation scope that produced the data point. Property behavior: - *read-only*
scope_version	string	Version, if any, of the instrumentation scope. Property behavior: - *read-only*
scope_attributes	object	Key-value attributes of the instrumentation scope. Property behavior: - *read-only*
scope_schema_url	string	URL, if any, of the OpenTelemetry schema of the instrumentation scope and its attributes. Property behavior: - *read-only*
service_name	string	Name of the service that produced the data point (e.g., `zabbix-frontend`), taken from the OpenTelemetry `service.name` resource attribute. Property behavior: - *read-only*
metric_name	string	Name of the metric that the data point belongs to (e.g., `http.server.request.duration`). Property behavior: - *read-only*
metric_description	string	Description, if any, of the metric. Property behavior: - *read-only*
metric_unit	string	Unit, if any, of the metric values (e.g., `s` or `By`; see UCUM). Property behavior: - *read-only*
attributes	object	Key-value data point attributes (e.g., `http.request.method: POST`). Property behavior: - *read-only*
start_time_unix	integer	Unix time (in nanoseconds) when the measurement period of the data point started. Property behavior: - *read-only*
time_unix	integer	Unix time (in nanoseconds) of the data point. For *Sum*, *Histogram*, and *Exponential histogram* data points, this is the end of the measurement period. Property behavior: - *read-only*

Property	Type	Description
value	float	Value of the data point. Applies to *Gauge* and *Sum* data points. Default: 0. Property behavior: - *read-only*
flags	integer	Bitmask of data point flags. If the least significant bit is set (value 1), no value was recorded for the data point. Property behavior: - *read-only*
exemplars	array	Array of exemplar objects, that is, individual measurements kept as examples, with the trace and span during which they were recorded. Property behavior: - *read-only*
aggregation_temporality	integer	How the values of the data point are aggregated over time. Applies to *Sum*, *Histogram*, and *Exponential histogram* data points. Possible values: 0 - *(default)* Unspecified; 1 - Delta (aggregated over the period since the previous data point); 2 - Cumulative (aggregated over the period since `start_time_unix`). Property behavior: - *read-only*
is_monotonic	boolean	Whether the sum can only increase or stay the same over time. Applies to *Sum* data points. Default: `false`. Property behavior: - *read-only*
count	integer	Number of measurements in the data point. Applies to *Histogram* and *Exponential histogram* data points. Default: 0. Property behavior: - *read-only*
sum	float	Sum of all measurements in the data point. Applies to *Histogram* and *Exponential histogram* data points. Default: 0. Property behavior: - *read-only*
bucket_counts	array	Array of integers with the number of measurements in each histogram bucket. The buckets are defined by `explicit_bounds`. Applies to *Histogram* data points. Default: `[]`. Property behavior: - *read-only*

Property	Type	Description
explicit_bounds	array	Array of floats with the boundaries between histogram buckets. Applies to *Histogram* data points. Default: `[]`. Property behavior: - *read-only*
min	float	Smallest of the measurements in the data point. Applies to *Histogram* and *Exponential histogram* data points. Default: 0. Property behavior: - *read-only*
max	float	Largest of the measurements in the data point. Applies to *Histogram* and *Exponential histogram* data points. Default: 0. Property behavior: - *read-only*
scale	integer	Scale of the exponential histogram, which determines the bucket boundaries and resolution. A higher scale gives narrower buckets and finer resolution; a lower scale gives wider buckets and coarser resolution. Applies to *Exponential histogram* data points. Default: 0. Property behavior: - *read-only*
zero_count	integer	Number of measurements in the zero bucket of the exponential histogram. Applies to *Exponential histogram* data points. Default: 0. Property behavior: - *read-only*
positive_bucket_counts	array	Array of integers with the number of positive measurements in each exponential histogram bucket. The bucket boundaries are determined by `scale`; the index of the first bucket is set in `positive_offset`. Applies to *Exponential histogram* data points. Default: `[]`. Property behavior: - *read-only*
positive_offset	integer	Index of the first bucket in `positive_bucket_counts`. Applies to *Exponential histogram* data points. Default: 0. Property behavior: - *read-only*

Property	Type	Description
negative_bucket_counts	array	Array of integers with the number of negative measurements in each exponential histogram bucket. The bucket boundaries are determined by `scale`; the index of the first bucket is set in `negative_offset`. Applies to *Exponential histogram* data points. Default: `[]`. Property behavior: - *read-only*
negative_offset	integer	Index of the first bucket in `negative_bucket_counts`. Applies to *Exponential histogram* data points. Default: 0. Property behavior: - *read-only*

Exemplar

The exemplar object has the following properties:

Property	Type	Description
traceid	string	ID, if any, of the trace during which the measurement was recorded. Data source column: `Exemplars.TraceId`.
spanid	string	ID, if any, of the span during which the measurement was recorded. Data source column: `Exemplars.SpanId`.
value	float	Value of the measurement. Data source column: `Exemplars.Value`.
time_unix	integer	Unix time (in nanoseconds) when the measurement was recorded. Data source column: `Exemplars.TimeUnix`.
filtered_attributes	object	Key-value attributes of the measurement that are not included in the `attributes` of the data point. Data source column: `Exemplars.FilteredAttributes`.

apmmetric.get

Description

`integer/array apmmetric.get(object parameters)`

The method allows you to retrieve metric data points according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
time_from	timestamp	Return only data points whose time is at or after the given time. Parameter behavior: - *required*

Parameter	Type	Description
time_till	timestamp	Return only data points whose time is before the given time. Parameter behavior: - *required*
types	integer/array	Return only data points of the given types. Refer to the APM metric object page for a list of supported metric types.
with_flags_on	integer	Return only data points that have all of the given flags set. To specify several flags, use the sum of their values.
with_flags_off	integer	Return only data points that have none of the given flags set. To specify several flags, use the sum of their values.
resource_attributes	object/array	Return only data points with the given resource attributes. Format: `[{"key": "<key>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all data points. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
resource_attributes_evaltype	integer	Evaluation method for the conditions in `resource_attributes`. Possible values: 0 - *(default)* And/Or; 2 - Or.
scope_attributes	object/array	Return only data points with the given scope attributes. Format: `[{"key": "<key>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all data points. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
scope_attributes_evaltype	integer	Evaluation method for the conditions in `scope_attributes`. Possible values: 0 - *(default)* And/Or; 2 - Or.
attributes	object/array	Return only data points with the given attributes. Format: `[{"key": "<key>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all data points. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
attributes_evaltype	integer	Evaluation method for the conditions in `attributes`. Possible values: 0 - *(default)* And/Or; 2 - Or.

Parameter	Type	Description
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Supported properties: `resource_schema_url`, `scope_schema_url`, `scope_name`, `scope_version`, `service_name`, `metric_name`, `metric_unit`, `flags`.
search	object	Return results that match the given pattern (case-insensitive). Accepts an object, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a `LIKE "%…%"` search. Supported properties: `resource_schema_url`, `scope_schema_url`, `scope_name`, `scope_version`, `service_name`, `metric_name`, `metric_description`, `metric_unit`.
sortfield	string/array	Sort the result by the given properties. Possible values: `type`, `resource_schema_url`, `scope_schema_url`, `scope_name`, `scope_version`, `service_name`, `metric_name`, `metric_unit`, `start_time_unix`, `time_unix`.
offset	integer	Number of records to skip before returning the result.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
limit	integer	
output	query	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving histogram data points

Retrieve the latest data point of the `http.server.request.duration` histogram metric of the `zabbix-frontend` service.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "apmmetric.get",
    "params": {
        "output": ["metric_name", "metric_unit", "time_unix", "count", "sum", "min", "max", "explicit_boun
        "time_from": 1790848800,
        "time_till": 1790852400,
        "types": 2,
        "filter": {
            "service_name": "zabbix-frontend",
            "metric_name": "http.server.request.duration"
        },
        "sortfield": "time_unix",
        "sortorder": "DESC",
        "limit": 1
    },
```

```
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "metric_name": "http.server.request.duration",
            "metric_unit": "s",
            "time_unix": "1790852390000000000",
            "count": "1842",
            "sum": "412.773",
            "min": "0.0061",
            "max": "1.098",
            "explicit_bounds": ["0.005", "0.05", "0.1", "0.25", "0.5", "1"],
            "bucket_counts": ["0", "211", "904", "512", "167", "41", "7"],
            "exemplars": [
                {
                    "traceid": "dee625f33cb0994e86c5cd4640e3b09c",
                    "spanid": "3f8b2c6d1e9a7054",
                    "value": "1.098",
                    "time_unix": "1790850513439147623",
                    "filtered_attributes": {}
                }
            ]
        }
    ],
    "id": 1
}
```

See also

- APM trace
- APM span
- APM log

Source

CApmMetric::get() in *ui/include/classes/api/services/CApmMetric.php*.

APM span

This class is designed to work with spans collected for application performance monitoring (APM).

Object references:

- APM span

Available methods:

- apmspan.get - retrieve spans

APM span object

The following objects are directly related to the `apmspan` API.

> **Note:**
> Each property corresponds to a column in the `otel_traces` table of the telemetry data source. For the list of columns, see Supported telemetry data columns.

APM span

The APM span object represents a span (a single operation within a trace, such as an HTTP request, a database query, or a function call).

Spans form a hierarchy, starting from the root span, which represents the first operation of the request. Every other span is a child span of the span during which its operation was started.

The APM span object has the following properties:

Property	Type	Description
timestamp	integer	Unix time (in nanoseconds) when the span started. Property behavior: - *read-only*
traceid	string	ID of the trace that the span belongs to. Property behavior: - *read-only*
spanid	string	ID of the span. Property behavior: - *read-only*
parent_spanid	string	ID of the parent span. Empty for the root span. Property behavior: - *read-only*
trace_state	string	Additional data, if any, that tracing tools attach to the trace (see W3C Trace Context). Property behavior: - *read-only*
span_name	string	Name of the operation that the span represents (e.g., `mysqli_query`). Property behavior: - *read-only*
span_kind	string	Role of the span. Possible values: `Unspecified` - span kind is not specified; `Internal` - operation performed within an application; `Server` - handling of an incoming request; `Client` - outgoing request; `Producer` - creation of an asynchronous operation (e.g., sending a message); `Consumer` - processing of an asynchronous operation. Property behavior: - *read-only*
service_name	string	Name of the service that produced the span (e.g., `zabbix-frontend`), taken from the OpenTelemetry `service.name` resource attribute. Property behavior: - *read-only*
resource_attributes	object	Key-value resource attributes of the span (e.g., `host.name: eabf8c2a5ddc`). Property behavior: - *read-only*
scope_name	string	Name of the instrumentation scope that produced the span. Property behavior: - *read-only*
scope_version	string	Version of the instrumentation scope. Property behavior: - *read-only*

Property	Type	Description
span_attributes	object	Key-value span attributes of the span (e.g., `db.system: mysql`). Property behavior: - *read-only*
duration	integer	Nanoseconds of the duration of the span. Property behavior: - *read-only*
status_code	string	Status of the operation that the span represents. Possible values: `Unset` - no explicit status has been assigned to the span; `Ok` - the operation was completed successfully; `Error` - the operation was completed with an error. Property behavior: - *read-only*
status_message	string	Additional information, if any, about the status of the span. Property behavior: - *read-only*
events	array	Array of event objects of the span, that is, notable points in time during the operation (e.g., exceptions). Property behavior: - *read-only*
links	array	Array of link objects of the span, which point to related spans in the same or another trace. Property behavior: - *read-only*

apmspan.get

Description

`integer/array apmspan.get(object parameters)`

The method allows you to retrieve spans according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
time_from	timestamp	Return only spans that started at or after the given time. Parameter behavior: - *required*
time_till	timestamp	Return only spans that started before the given time. Parameter behavior: - *required*
traceids	string/array	Return only spans that belong to the given traces.
spanids	string/array	Return only spans with the given IDs.

Parameter	Type	Description
parent_spanids	string/array	Return only spans with the given parent spans.
resource_attributes	object/array	Return only spans with the given resource attributes. Format: `[{"key": "<key>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all spans. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
resource_attributes_evaltype	integer	Evaluation method for the conditions in `resource_attributes`. Possible values: 0 - *(default)* And/Or; 2 - Or.
span_attributes	object/array	Return only spans with the given span attributes. Format: `[{"key": "<key>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all spans. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
span_attributes_evaltype	integer	Evaluation method for the conditions in `span_attributes`. Possible values: 0 - *(default)* And/Or; 2 - Or.
min_duration	integer	Return only spans whose duration is greater than or equal to the given value, in nanoseconds.
max_duration	integer	Return only spans whose duration is less than the given value, in nanoseconds.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Supported properties: `traceid`, `spanid`, `parent_spanid`, `trace_state`, `span_name`, `span_kind`, `service_name`, `scope_name`, `scope_version`, `status_code`, `status_message`.
search	object	Return results that match the given pattern (case-insensitive). Accepts an object, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a `LIKE "%...%"` search. Supported properties: `trace_state`, `span_name`, `service_name`, `scope_name`, `scope_version`, `status_message`.
sortfield	string/array	Sort the result by the given properties. Possible values: `timestamp`, `traceid`, `spanid`, `parent_spanid`, `trace_state`, `span_name`, `span_kind`, `service_name`, `scope_name`, `scope_version`, `duration`, `status_code`, `status_message`.

Parameter	Type	Description
offset	integer	Number of records to skip before returning the result.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
limit	integer	
output	query	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving the span hierarchy of a trace

Retrieve the first four spans of the trace "dee625f33cb0994e86c5cd4640e3b09c", sorted by start time. The `parent_spanid` property shows how the spans are related: the root span has no parent, and each other span refers to the span that started it.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "apmspan.get",
    "params": {
        "output": ["spanid", "parent_spanid", "service_name", "span_name", "duration", "status_code"],
        "time_from": 1790848800,
        "time_till": 1790852400,
        "traceids": "dee625f33cb0994e86c5cd4640e3b09c",
        "sortfield": "timestamp",
        "sortorder": "ASC",
        "limit": 4
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "spanid": "7c1a9e2f4b6d8a03",
            "parent_spanid": "",
            "service_name": "zabbix-nginx",
            "span_name": "POST /zabbix.php",
            "duration": "1104826531",
            "status_code": "Unset"
        },
        {
            "spanid": "3f8b2c6d1e9a7054",
            "parent_spanid": "7c1a9e2f4b6d8a03",
            "service_name": "zabbix-frontend",
            "span_name": "POST /zabbix.php",
            "duration": "1098230418",
            "status_code": "Unset"
        },
        {
            "spanid": "b3390ad41931ed08",
            "parent_spanid": "3f8b2c6d1e9a7054",
```

```
            "service_name": "zabbix-frontend",
            "span_name": "mysqli::real_connect",
            "duration": "36671360",
            "status_code": "Ok"
        },
        {
            "spanid": "e41f7a2c9d5b3086",
            "parent_spanid": "3f8b2c6d1e9a7054",
            "service_name": "zabbix-frontend",
            "span_name": "mysqli_query",
            "duration": "19704118",
            "status_code": "Ok"
        }
    ],
    "id": 1
}
```


See also

- APM trace
- APM log
- APM metric

Source

CApmSpan::get() in *ui/include/classes/api/services/CApmSpan.php*.

APM trace

This class is designed to work with traces collected for application performance monitoring (APM).

Object references:

- APM trace
- Event
- Link

Available methods:

- apmtrace.get - retrieve traces

APM trace object

The following objects are directly related to the `apmtrace` API.

Note:
Each property, except `span_count` and `error_count`, corresponds to a column in the `otel_traces` table of the telemetry data source. Most properties correspond to the column with the same name (e.g., `span_name` property corresponds to the `SpanName` column). Where the names differ, the data source columns are specified in the tables below. For the list of columns, see Supported telemetry data columns.

APM trace

The APM trace object represents a trace (the path that a request takes through your applications and services). A trace consists of one or more spans; each span represents a single operation.

The APM trace object is based on the root span of the trace (the span that starts the trace and has no parent span). All properties contain data of the root span, except `span_count` and `error_count`, which are calculated from all spans of the trace.

The APM trace object has the following properties:

Property	Type	Description
timestamp	integer	Unix time (in nanoseconds) when the root span started. Property behavior: - *read-only*
traceid	string	ID of the trace, shared by all spans that belong to the trace. Property behavior: - *read-only*
spanid	string	ID of the root span. Property behavior: - *read-only*
trace_state	string	Additional data, if any, that tracing tools attach to the trace (see W3C Trace Context). Property behavior: - *read-only*
span_name	string	Name of the operation that the root span represents (e.g., `POST /zabbix.php`). Property behavior: - *read-only*
span_kind	string	Role of the root span. Possible values: `Unspecified` - span kind is not specified; `Internal` - operation performed within an application; `Server` - handling of an incoming request; `Client` - outgoing request; `Producer` - creation of an asynchronous operation (e.g., sending a message); `Consumer` - processing of an asynchronous operation. Property behavior: - *read-only*
service_name	string	Name of the service that produced the root span (e.g., `zabbix-nginx`), taken from the OpenTelemetry `service.name` resource attribute. Property behavior: - *read-only*
resource_attributes	object	Key-value resource attributes of the root span (e.g., `host.name: eabf8c2a5ddc`). Property behavior: - *read-only*
scope_name	string	Name of the instrumentation scope that produced the root span. Property behavior: - *read-only*
scope_version	string	Version of the instrumentation scope. Property behavior: - *read-only*
span_attributes	object	Key-value span attributes of the root span (e.g., `http.method: POST`). Property behavior: - *read-only*

Property	Type	Description
duration	integer	Nanoseconds of the duration of the trace. Property behavior: - *read-only*
status_code	string	Status of the operation that the root span represents. Possible values: `Unset` - no explicit status has been assigned to the root span; `Ok` - the operation was completed successfully; `Error` - the operation was completed with an error. Property behavior: - *read-only*
status_message	string	Additional information, if any, about the status of the root span. Property behavior: - *read-only*
events	array	Array of event objects of the root span, that is, notable points in time during the operation (e.g., exceptions). Property behavior: - *read-only*
links	array	Array of link objects of the root span, which point to related spans in the same or another trace. Property behavior: - *read-only*
span_count	integer	Total number of spans in the trace. Property behavior: - *read-only*
error_count	integer	Number of spans in the trace with `status_code` set to `Error`. Property behavior: - *read-only*

Event

The event object has the following properties.

Property	Type	Description
timestamp	integer	Unix time (in nanoseconds) when the event occurred. Data source column: `Events.Timestamp`
name	string	Name of the event. Data source column: `Events.Name`
attributes	object	Key-value attributes with additional information, if any, about the event. Data source column: `Events.Attributes`

Link

The link object has the following properties.

Property	Type	Description
traceid	string	ID of the trace that contains the linked span. Data source column: `Links.TraceId`
spanid	string	ID of the linked span. Together with `traceid`, uniquely identifies the linked span. Data source column: `Links.SpanId`

Property	Type	Description
trace_state	string	Trace state, if any, of the linked span (see W3C Trace Context). Data source column: `Links.TraceState`
attributes	object	Key-value attributes with additional information about the link. Data source column: `Links.Attributes`

apmtrace.get

Description

`integer/array apmtrace.get(object parameters)`

The method allows you to retrieve traces according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
time_from	timestamp	Return only traces whose root span started at or after the given time. Parameter behavior: - *required*
time_till	timestamp	Return only traces whose root span started before the given time. Parameter behavior: - *required*
traceids	string/array	Return only traces with the given IDs.
resource_attributes	object/array	Return only traces whose root span has the given resource attributes. Format: `[{"key": "<key>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all traces. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
resource_attributes_evaltype	integer	Evaluation method for the conditions in `resource_attributes`. Possible values: 0 - *(default)* And/Or; 2 - Or.

Parameter	Type	Description
span_attributes	object/array	Return only traces whose root span has the given span attributes. Format: `[{"key": "<key>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all traces. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
span_attributes_evaltype	integer	Evaluation method for the conditions in `span_attributes`. Possible values: 0 - *(default)* And/Or; 2 - Or.
min_duration	integer	Return only traces whose root span `duration` is greater than or equal to the given value, in nanoseconds.
max_duration	integer	Return only traces whose root span `duration` is less than the given value, in nanoseconds.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Supported properties: `traceid`, `spanid`, `trace_state`, `span_name`, `span_kind`, `service_name`, `scope_name`, `scope_version`, `status_code`, `status_message`.
search	object	Return results that match the given pattern (case-insensitive). Accepts an object, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a `LIKE "%...%"` search. Supported properties: `trace_state`, `span_name`, `service_name`, `scope_name`, `scope_version`, `status_message`.
sortfield	string/array	Sort the result by the given properties. Possible values: `timestamp`, `traceid`, `spanid`, `trace_state`, `span_name`, `span_kind`, `service_name`, `scope_name`, `scope_version`, `duration`, `status_code`, `status_message`.
offset	integer	Number of records to skip before returning the result.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
limit	integer	
output	query	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving slow traces

Retrieve the two slowest traces of the `zabbix-nginx` service. Search traces from 10:00 to 11:00 on October 1, 2026 (UTC), and include only those that took 1 second or longer.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "apmtrace.get",
    "params": {
        "output": ["timestamp", "traceid", "span_name", "service_name", "duration", "status_code", "span_c
        "time_from": 1790848800,
        "time_till": 1790852400,
        "filter": {
            "service_name": "zabbix-nginx"
        },
        "min_duration": 1000000000,
        "sortfield": "duration",
        "sortorder": "DESC",
        "limit": 2
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "timestamp": "1790849274910226417",
            "traceid": "8a4f1c7e2b9d4e6fa0c3b5d7e9f1a2c4",
            "span_name": "POST /zabbix.php",
            "service_name": "zabbix-nginx",
            "duration": "30412907815",
            "status_code": "Error",
            "span_count": "49",
            "error_count": "2"
        },
        {
            "timestamp": "1790850512338104952",
            "traceid": "dee625f33cb0994e86c5cd4640e3b09c",
            "span_name": "POST /zabbix.php",
            "service_name": "zabbix-nginx",
            "duration": "1104826531",
            "status_code": "Unset",
            "span_count": "27",
            "error_count": "0"
        }
    ],
    "id": 1
}
```

See also

- APM span
- APM metric
- APM log

Source

CApmTrace::get() in *ui/include/classes/api/services/CApmTrace.php*.

Audit log

This class is designed to work with audit log.

Object references:

- Audit log

Available methods:

- auditlog.get - retrieve audit log records

Audit log object

The following objects are directly related to the `auditlog` API.

Audit log

The audit log object contains information about user actions. It has the following properties.

Property	Type	Description
auditid	ID	ID of audit log entry. Generated using CUID algorithm.
userid	ID	Audit log entry author userid.
username	string	Audit log entry author username.
clock	timestamp	Audit log entry creation timestamp.
ip	string	Audit log entry author IP address.
action	integer	Audit log entry action. Possible values: 0 - Add; 1 - Update; 2 - Delete; 4 - Logout; 7 - Execute; 8 - Login; 9 - Failed login; 10 - History clear; 11 - Config refresh; 12 - Push.

Property	Type	Description
resourcetype	integer	Audit log entry resource type. Possible values: 0 - User; 3 - Media type; 4 - Host; 5 - Action; 6 - Graph; 11 - User group; 13 - Trigger; 14 - Host group; 15 - Item; 16 - Image; 17 - Value map; 18 - Service; 19 - Map; 22 - Web scenario; 23 - Discovery rule; 25 - Script; 26 - Proxy; 27 - Maintenance; 28 - Regular expression; 29 - Macro; 30 - Template; 31 - Trigger prototype; 32 - Icon mapping; 33 - Dashboard; 34 - Event correlation; 35 - Graph prototype; 36 - Item prototype; 37 - Host prototype; 38 - Autoregistration; 39 - Module; 40 - Settings; 41 - Housekeeping; 42 - Authentication; 43 - Template dashboard; 44 - User role; 45 - API token; 46 - Scheduled report; 47 - High availability node; 48 - SLA; 49 - User directory; 50 - Template group; 51 - Connector; 52 - LLD rule; 53 - History; 54 - Multi-factor authentication; 55 - Proxy group; 56 - LLD rule prototype.
resourceid	ID	Audit log entry resource identifier.
resource_cuid	ID	Audit log entry resource unique identifier generated using the CUID algorithm.
resourcename	string	Audit log entry resource human readable name.
recordsetid	ID	Audit log entry recordset ID. The audit log records created during the same operation will have the same recordset ID. Generated using CUID algorithm.

Property	Type	Description
details	text	Audit log entry details. The details are stored as a JSON object, where each property name is a path to the property or nested object in which the change occurred, and where each value contains the data (in array format) about the change in this property or nested object. Possible value formats: ["add"] - Nested object has been added; ["add", "<value>"] - The property of the added object equals <value>; ["update"] - Nested object has been updated; ["update", "<new value>", "<old value>"] - The property of the updated object was changed from <old value> to <new value>; ["delete"] - Nested object has been deleted.

auditlog.get

Description

`integer/array auditlog.get(object parameters)`

The method allows you to retrieve audit log records according to the given parameters.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
auditids	ID/array	Return only audit log with the given IDs.
userids	ID/array	Return only audit log that were created by the given users.
time_from	timestamp	Returns only audit log entries that have been created after or at the given time.
time_till	timestamp	Returns only audit log entries that have been created before or at the given time.
sortfield	string/array	Sort the result by the given properties. Possible values: `auditid`, `userid`, `clock`.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve audit log

Retrieve two latest audit log records.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "auditlog.get",
    "params": {
        "output": "extend",
        "sortfield": "clock",
        "sortorder": "DESC",
        "limit": 2
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "auditid": "cksstgfam0001yhdcc41y20q2",
            "userid": "1",
            "username": "Admin",
            "clock": "1629975715",
            "ip": "127.0.0.1",
            "action": "1",
            "resourcetype": "0",
            "resourceid": "0",
            "resource_cuid": "0",
            "resourcename": "Jim",
            "recordsetid": "cksstgfal0000yhdcso67ondl",
            "details": "{\"user.name\":[\"update\",\"Jim\",\"\"],\"user.medias[37]\":[\"add\"],\"user.medi
        },
        {
            "auditid": "ckssofl0p0001yhdcqxclsg8r",
            "userid": "1",
            "username": "Admin",
            "clock": "1629967278",
            "ip": "127.0.0.1",
            "action": "0",
            "resourcetype": "0",
            "resourceid": "20",
            "resource_cuid": "0",
            "resourcename": "John",
            "recordsetid": "ckssofl0p0000yhdcpxyo1jgo",
            "details": "{\"user.username\":[\"add\",\"John\"], \"user.userid:\":[\"add\",\"20\"],\"user.us
        }
    ],
    "id": 1
}
```

See also

- Audit log object

Source

CAuditLog::get() in *ui/include/classes/api/services/CAuditLog.php*.

Authentication

This class is designed to work with authentication settings.

Object references:

- Authentication

Available methods:

- authentication.get - retrieve authentication settings
- authentication.update - update authentication settings

Authentication object

The following objects are directly related to the `authentication` API.

Authentication

The authentication object has the following properties.

Property	Type	Description
authentication_type	integer	Default authentication. Possible values: 0 - *(default)* Internal; 1 - LDAP.
http_auth_enabled	integer	HTTP authentication. Possible values: 0 - *(default)* Disabled; 1 - Enabled. Property behavior: - *supported* if `$ZBX_FEATURE_FLAGS['http_auth_enabled']` is set to `true` in the frontend configuration file.
http_login_form	integer	Default login form. Possible values: 0 - *(default)* Zabbix login form; 1 - HTTP login form. Property behavior: - *supported* if `$ZBX_FEATURE_FLAGS['http_auth_enabled']` is set to `true` in the frontend configuration file.
http_strip_domains	string	Domain name to remove. Property behavior: - *supported* if `$ZBX_FEATURE_FLAGS['http_auth_enabled']` is set to `true` in the frontend configuration file.
http_case_sensitive	integer	HTTP case-sensitive login. Possible values: 0 - Off; 1 - *(default)* On. Property behavior: - *supported* if `$ZBX_FEATURE_FLAGS['http_auth_enabled']` is set to `true` in the frontend configuration file.
ldap_auth_enabled	integer	LDAP authentication. Possible values: 0 - *(default)* Disabled; 1 - Enabled.
ldap_case_sensitive	integer	LDAP case-sensitive login. Possible values: 0 - Off; 1 - *(default)* On.

Property	Type	Description
ldap_userdirectoryid	ID	ID of the default user directory for LDAP authentication. Used for user groups with `gui_access` set to LDAP or System default. Property behavior: - *required* if `ldap_auth_enabled` is set to ”Enabled”
saml_auth_enabled	integer	SAML authentication. Possible values: 0 - *(default)* Disabled; 1 - Enabled.
saml_case_sensitive	integer	SAML case-sensitive login. Possible values: 0 - Off; 1 - *(default)* On.
passwd_min_length	integer	Password minimal length requirement. Possible values range from 1 to 70. Default: 8.
passwd_check_rules	integer	Password checking rules. Possible bitmap values: 0 - Check password length; 1 - Check if password uses uppercase and lowercase Latin letters; 2 - Check if password uses digits; 4 - Check if password uses special characters; 8 - *(default)* Check if password is not in the list of commonly used passwords and does not contain derivations of word ”Zabbix” or user’s name, last name, or username. This is a bitmask field; any sum of possible bitmap values is acceptable (for example, 15 for checking all rules).
ldap_jit_status	integer	Status of LDAP provisioning. Possible values: 0 - Disabled for configured LDAP IdPs; 1 - Enabled for configured LDAP IdPs.
saml_jit_status	integer	Status of SAML provisioning. Possible values: 0 - Disabled for configured SAML IdPs; 1 - Enabled for configured SAML IdPs.
jit_provision_interval	string	Time interval between JIT provision requests for logged-in user. Accepts a number of seconds (e.g., 3600), or a value with a time suffix (e.g., `3600s`, `60m`, `1h`, `1d`, `1M`, `1y`), including support for months and years. Minimum value: 1h. Default: 1h. Available only for LDAP provisioning.
disabled_usrgrpid	ID	ID of the user group to assign the deprovisioned user to. The user group must be disabled and cannot be enabled or deleted when configured. Property behavior: - *required* if `ldap_jit_status` is set to ”Enabled for configured LDAP IdPs”, or `saml_jit_status` is set to ”Enabled for configured SAML IdPs”

Property	Type	Description
mfa_status	integer	Multi-factor authentication. Possible values: 0 - Disabled (for all configured MFA methods); 1 - Enabled (for all configured MFA methods).
mfaid	ID	Default MFA method for user groups with MFA enabled. Property behavior: - *required* if `mfa_status` is set to "Enabled"

authentication.get

Description

`object authentication.get(object parameters)`

The method allows you to retrieve authentication object according to the given parameters.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports only one parameter.

Parameter	Type	Description
output	query	This parameter is described in the reference commentary.

Return values

`(object)` Returns authentication object.

Examples

Request:

```
{
    "jsonrpc": "2.0",
    "method": "authentication.get",
    "params": {
        "output": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "authentication_type": "0",
        "http_auth_enabled": "0",
        "http_login_form": "0",
        "http_strip_domains": "",
        "http_case_sensitive": "1",
        "ldap_auth_enabled": "0",
        "ldap_case_sensitive": "1",
        "ldap_userdirectoryid": "0",
        "saml_auth_enabled": "0",
        "saml_case_sensitive": "0",
```

```
        "passwd_min_length": "8",
        "passwd_check_rules": "15",
        "jit_provision_interval": "1h",
        "saml_jit_status": "0",
        "ldap_jit_status": "0",
        "disabled_usrgrpid": "9",
        "mfa_status": "0",
        "mfaid": "0"
    },
    "id": 1
}
```


Source

CAuthentication::get() in *ui/include/classes/api/services/CAuthentication.php*.

authentication.update

Description

`object authentication.update(object authentication)`

This method allows you to update existing authentication settings.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Authentication properties to be updated.

Return values

`(array)` Returns an array with the names of updated parameters.

Examples

Request:


```
{
    "jsonrpc": "2.0",
    "method": "authentication.update",
    "params": {
        "http_auth_enabled": 1,
        "http_case_sensitive": 0,
        "http_login_form": 1
    },
    "id": 1
}
```


Response:


```
{
    "jsonrpc": "2.0",
    "result": [
        "http_auth_enabled",
        "http_case_sensitive",
        "http_login_form"
    ],
    "id": 1
}
```


Source

CAuthentication::update() in *ui/include/classes/api/services/CAuthentication.php*.

Autoregistration

This class is designed to work with autoregistration.

Object references:

- Autoregistration

Available methods:

- autoregistration.get - retrieve autoregistration
- autoregistration.update - update autoregistration

Autoregistration object

The following objects are directly related to the `autoregistration` API.

Autoregistration

The autoregistration object has the following properties.

Property	Type	Description
tls_accept	integer	Type of allowed incoming connections for autoregistration. Possible values: 1 - allow unencrypted connections; 2 - allow TLS with PSK; 3 - allow both unencrypted and TLS with PSK connections.
tls_psk_identity	string	PSK identity; must be paired with only one PSK (across autoregistration, hosts, and proxies). Do not include sensitive information in the PSK identity, as it is sent unencrypted over the network to inform the receiver which PSK to use. Property behavior: - *write-only*
tls_psk	string	Pre-shared key (PSK); must be at least 32 hex digits. Property behavior: - *write-only*

autoregistration.get

Description

`object autoregistration.get(object parameters)`

The method allows you to retrieve autoregistration object according to the given parameters.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports only one parameter.

Parameter	Type	Description
output	query	This parameter is described in the reference commentary.

Return values

(object) Returns autoregistration object.

Examples

Request:

```
{
   "jsonrpc": "2.0",
    "method": "autoregistration.get",
    "params": {
        "output": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "tls_accept": "3"
    },
    "id": 1
}
```

Source

CAutoregistration::get() in *ui/include/classes/api/services/CAutoregistration.php*.

autoregistration.update

Description

`object autoregistration.update(object autoregistration)`

This method allows you to update existing autoregistration.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

(object) Autoregistration properties to be updated.

Return values

(boolean) Returns boolean true as result on successful update.

Examples

Request:

```
{
    "jsonrpc": "2.0",
    "method": "autoregistration.update",
    "params": {
        "tls_accept": "3",
        "tls_psk_identity": "PSK 001",
        "tls_psk": "11111595725ac58dd977beef14b97461a7c1045b9a1c923453302c5473193478"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": true,
```

```
    "id": 1
}
```

Source

CAutoregistration::update() in *ui/include/classes/api/services/CAutoregistration.php*.

Complex event processing

This class is designed to work with complex event processing.

Object references:

- CEP rule
- CEP rule filter
  - CEP rule filter condition
- CEP rule time window
- CEP rule operation
- CEP rule operation filter
- CEP rule operation filter condition

Available methods:

- ceprule.create - create new CEP rules
- ceprule.delete - delete CEP rules
- ceprule.get - retrieve CEP rules
- ceprule.update - update CEP rules
- ceprule.resettimewindows - reset time windows for CEP rule

CEP rule object

The following objects are directly related to the `ceprule` API.

CEP rule

The CEP rule object has the following properties.

Property	Type	Description
cep_ruleid	ID	ID of the CEP rule. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the CEP rule. Property behavior: - *required* for create operations
window_type	integer	Time window type: 0 - (default) None 1 - Simple (multiple); 2 - Cause and symptom grouping; 3 - Simple (once); 4 - Event pattern match.
stop	integer	Possible values: 0 - (default) Continue with the next CEP rule according to sort order; 1 - Stop after executing this rule.
sortorder	integer	Execution order of the CEP rule.
description	string	Description of the CEP rule.
status	integer	Whether the CEP rule is enabled or disabled. Possible values: 0 - *(default)* enabled; 1 - disabled.

Property	Type	Description
error	string	Error text if the CEP rule processing failed on rule level. Property behavior: - *read-only*

CEP rule filter

The CEP rule filter object defines a set of conditions that must be met for events to be processed by CEP rule. It has the following properties.

Property	Type	Description
conditions	array	Set of filter conditions to use for filtering results. The conditions will be sorted in the order of their placement in the formula. Property behavior: - *required*
evaltype	integer	Filter condition evaluation method. Possible values: 0 - And/Or; 1 - And; 2 - Or; 3 - Custom expression. Property behavior: - *required*
eval_formula	string	Generated expression that will be used for evaluating filter conditions. The expression contains IDs that reference specific filter conditions by its `formulaid`. The value of `eval_formula` is equal to the value of `formula` for filters with a custom expression. Property behavior: - *read-only*
formula	string	User-defined expression to be used for evaluating conditions of filters with a custom expression. The expression must contain IDs that reference specific filter conditions by its `formulaid`. The IDs used in the expression must exactly match the ones defined in the filter conditions: no condition can remain unused or omitted. Property behavior: - *required* if `evaltype` is set to ”custom expression”

CEP rule filter condition

The CEP rule filter condition object defines a specific condition that must be checked before processing events by the CEP rule.

Property	Type	Description
type	integer	Type of condition. Possible values: 3 - event name; 4 - severity; 6 - time period; 25 - tag name; 26 - tag value; 36 - host visible name; 37 - host group name. Property behavior: - *required*
operator	integer	Condition operator. Possible values: 0 - equals (default); 1 - does not equal; 2 - contains; 3 - does not contain; 4 - in; 5 - is greater than or equals; 6 - is less than or equals; 7 - not in. Property behavior: - *required*.
event_name	string	Event name for a specific operation type. Property behavior: - *required* if `type` is set to ”event name”
severity	string	Event severity for a specific operation type. Property behavior: - *required* if `type` is set to ”severity”
time_period	string	Time period for a specific operation type. Property behavior: - *required* if `type` is set to ”time period”
tag	string	Tag name for a specific operation type. Property behavior: - *required* if `type` is set to ”tag name” or ”tag value”
tag_value	string	Tag value for a specific operation type. Property behavior: - *required* if `type` is set to ”tag value”
event_severity	string	Event severity for a specific operation type. Property behavior: - *required* if `type` is set to ”severity”
host_name	string	Host name for a specific operation type. Property behavior: - *required* if `type` is set to ”host visible name”
host_group_name	string	Host group name for a specific operation type. Property behavior: - *required* if `type` is set to ”host group name”

Property	Type	Description
formulaid	string	Arbitrary unique ID that is used to reference the condition from a custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward. Property behavior: - *required* if CEP rule filter `evaltype` is set to ”custom expression”

CEP rule time window

The CEP rule time window object defines the time window of the CEP rule. It has the following properties.

Property	Type	Description
duration	string	Duration of the CEP rule window (365 days maximum). Accepts seconds and time units with suffix. Default: `10m`
capacity	integer	Capacity of the CEP rule window. Supports user macros. Default: 0 (unlimited)
group_by_host_group	integer	Group by flag for host group. Possible values: 0 - no (default); 1 - yes.
group_by_host	integer	Group by flag for host. Possible values: 0 - no (default); 1 - yes.
group_by_tags	integer	Group by flag for tag. Possible values: 0 - no (default); 1 - yes.
tags	string	Tag names to use for grouping. Property behavior: - *required* if `group_by_tag` is set to ”1”
event_count_tag	string	Tag name for displaying the count of affected events. Property behavior: - *supported* if CEP rule `window_type` is set to ”2”; not allowed otherwise
script	string	Script to execute. Property behavior: - *supported* if CEP rule `window_type` is set to ”4”; not allowed otherwise

CEP rule operation

The CEP rule operation object defines an operation that will be performed when a CEP rule is executed. It has the following properties.

Property	Type	Description
sortorder	integer	Execution order (ascending) of the operation. Property behavior: - *required*
execute_when	integer	Type of condition for operation. Possible values: 0 - event occurred; 1 - event added; 2 - event evicted; 3 - window closed; 4 - event pattern matched; 5 - event closed. Property behavior: - *required*
scope	integer	Scope of the operation. Possible values: 0 - current event (default); 1 - other events; 2 - all events. Possible values when `execute_when` is ”Window closed” or ”Pattern matched”: 2 - all events. Possible values when `execute_when` is ”Event ocurred”: 0 - current event.
filter	object	Filter object.
scope_filter	object	Scope filter object. scope_filter is editable only if the `scope` field contains ”Other events” or ”All events” value. Otherwise scope_filter accepts only object with default values.

Property	Type	Description
type	integer	CEP operation type. Possible values: 1 - set event name; 2 - close event; 3 - discard event; 4 - set event severity; 5 - increase event severity; 6 - decrease event severity; 7 - suppress event; 8 - unsuppress event; 9 - copy first event as new; 10 - copy last event as new; 11 - add tag; 12 - set tag; 13 - set tag value; 14 - increase tag value; 15 - decrease tag value; 16 - rename tag; 17 - remove tag; 18 - close window. Property behavior: - *required* - *supported*: 1, 2, 3, 4, 5, 6, 7, 8, 11, 12, 13, 14, 15, 16, 17 - supported if ”Event occurred” (see `execute_when`) 1, 2, 4, 5, 6, 7, 8, 11, 12, 13, 14, 15, 16, 17, 18 - supported if ”Event added” (see `execute_when`) 1, 2, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18 - supported if ”Event evicted” (see `execute_when`) 1, 2, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18 - supported if ”Event closed” (see `execute_when`) 1, 2, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17 - supported if ”Window closed” (see `execute_when`) 1, 2, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18 - supported if ”Event pattern matched” (see `execute_when`)
event_name	string	Event name for a specific operation type. Property behavior: - *required* if `type` is set to ”Set event name”
tag	string	Tag name for specific operation types. Property behavior: - *required* if `type` is set to ”Add tag”, ”Set tag”, ”Set tag value”, ”Increase tag value”, ”Decrease tag value”, ”Rename tag”, ”Remove tag”
new_tag	string	New tag name for a specific operation type. Property behavior: - *required* if `type` is set to ”Rename tag”
tag_value	string	Tag value for specific operation types. Property behavior: - *required* if `type` is set to ”Add tag”, ”Set tag”, ”Set tag value”, ”Increase tag value”, ”Decrease tag value”, ”Rename tag”, ”Remove tag”
severity	integer	Event severity for a specific operation type. Property behavior: - *required* if `type` is set to ”Set event severity”

Property	Type	Description
suppress_duration	string	Duration of the event suppression. Accepts seconds and time unit with suffix. Property behavior: - *required* if `type` is set to ”Suppress event”

CEP rule operation filter

The CEP rule operation filter object defines a set of conditions that must be met to perform the configured CEP operations. It has the following properties.

Property	Type	Description
conditions	array	Set of operation filter conditions to use for filtering results. The conditions will be sorted in the order of their placement in the formula. Property behavior: - *required*
evaltype	integer	Filter condition evaluation method. Possible values: 0 - And/Or; 1 - And; 2 - Or; 3 - Custom expression. Property behavior: - *required*
eval_formula	string	Generated expression that will be used for evaluating filter conditions. The expression contains IDs that reference specific filter conditions by its `formulaid`. The value of `eval_formula` is equal to the value of `formula` for filters with a custom expression. Property behavior: - *read-only*
formula	string	Arbitrary unique ID that is used to reference the condition from a custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward. Property behavior: - *required*

CEP rule operation filter condition

The CEP rule operation filter condition object defines a specific condition that must be checked before running the CEP rule operations.

Property	Type	Description
type	integer	Type of operation condition. Possible values: 25 - tag name; 26 - tag value; 30 - problem opened; 31 - first; 32 - last; 33 - symptom; 34 - cloned; 35 - suppressed. Property behavior: - *required* - *supported*: 25, 26, 30, 33, 34, 35 - supported if ”Event occurred” (see `execute_when`)
operator	integer	Condition operator. Possible values: 0 - equals (default); 1 - does not equal; 2 - contains; 3 - does not contain; 5 - is greater than or equals; 6 - is less than or equals; 10 - true; 11 - false. Property behavior: - *required* - *supported*: 0, 1, 2, 3 - supported values if `type` set to ”Tag name” 0, 1, 2, 3, 5, 6 - supported values if `type` set to ”Tag value” 10, 11 - supported values if `type` set to ”Problem opened”, ”First”, ”Last”, ”Symptom”, ”Cloned”, ”Suppressed”.
tag	string	Tag name of operation condition. Property behavior: - *required* if `type` is set to ”Tag name” or ”Tag value”
tag_value	string	Value to compare with. Property behavior: - *required* if `type` is set to ”Tag value”

CEP rule operation scope filter

The CEP rule operation scope filter object defines a set of conditions that must be met to apply the configured CEP operations to events. It has the following properties.

Property	Type	Description
conditions	array	Set of operation scope filter conditions to use for filtering results. The conditions will be sorted in the order of their placement in the formula. Property behavior: - *required*

Property	Type	Description
evaltype	integer	Filter condition evaluation method. Possible values: 0 - And/Or; 1 - And; 2 - Or; 3 - Custom expression. Property behavior: - *required*
eval_formula	string	Generated expression that will be used for evaluating filter conditions. The expression contains IDs that reference specific filter conditions by its `formulaid`. The value of `eval_formula` is equal to the value of `formula` for filters with a custom expression. Property behavior: - *read-only*
formula	string	Arbitrary unique ID that is used to reference the condition from a custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward. Property behavior: - *required*

CEP rule operation scope filter condition

The CEP rule operation scope filter condition object defines a specific condition that must be checked before applying the CEP rule operations to events.

Property	Type	Description
type	integer	Type of operation condition. Possible values: 25 - tag name; 26 - tag value; 30 - problem opened; 31 - first; 32 - last; 33 - symptom; 34 - cloned; 35 - suppressed. Property behavior: - *required* - *supported*: 25, 26, 30, 33, 34, 35 - supported if ”Event occurred” (see `execute_when`)

Property	Type	Description
operator	integer	Condition operator. Possible values: 0 - equals (default); 1 - does not equal; 2 - contains; 3 - does not contain; 5 - is greater than or equals; 6 - is less than or equals; 10 - true; 11 - false. Property behavior: - *required* - *supported*: 0, 1, 2, 3 - supported values if `type` set to ”Tag name” 0, 1, 2, 3, 5, 6 - supported values if `type` set to ”Tag value” 10, 11 - supported values if `type` set to ”Problem opened”, ”First”, ”Last”, ”Symptom”, ”Cloned”, ”Suppressed”.
tag	string	Tag name of operation condition. Property behavior: - *required* if `type` is set to ”Tag name” or ”Tag value”
tag_value	string	Value to compare with. Property behavior: - *required* if `type` is set to ”Tag value”

ceprule.create

Description

`object ceprule.create(object/array CEP rules)`

This method allows you to create new CEP rules.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` CEP rules to create.

Additionally to the standard CEP properties, the method accepts the following parameters.

Parameter	Type	Description
operations	array	CEP rule operations to create for the CEP rule. Parameter behavior: - *required*
filter	object	CEP rule filter object for the CEP rule. Parameter behavior: - *required*

Return values

`(object)` Returns an object containing the IDs of the created CEP rules under the `cep_ruleids` property. The order of the returned IDs matches the order of the passed CEP rules.

Examples

Create a new CEP rule

Create a CEP rule using evaluation method `AND/OR` with one condition and one window. By default, the rule will be enabled.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "ceprule.create",
    "params": {
      "name": "Group events by cause",
      "stop": "1",
      "sortorder": "0",
      "description": "",
      "status": "0",
      "window_type": "2",
      "filter": {
        "evaltype": "0",
        "formula": "",
        "conditions": [
          {
            "type": "3",
            "operator": "2",
            "event_name": "Printer",
            "tag": "",
            "tag_value": "",
            "severity": "0",
            "host_name": "",
            "host_group_name": "",
            "time_period": "",
            "formulaid": ""
          }
        ]
      },
      "window": {
        "duration": "1d",
        "capacity": "0",
        "group_by_host_group": "0",
        "group_by_host": "1",
        "group_by_tags": "1",
        "tags": [
          "cause"
        ],
        "event_count_tag": "",
        "script": ""
      },
      "operations": []
    },
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "cep_ruleids": [
      "1"
    ]
  },
  "id": 1
}
```

See also

- CEP rule

Source

CCepRule::create() in *ui/include/classes/api/services/CCepRule.php*.

ceprule.delete

Description

`object ceprule.delete(array cep_ruleids)`

This method allows you to delete CEP rules.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the CEP rules to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted CEP rules under the `cep_ruleids` property.

Example

Delete multiple CEP rules

Delete two CEP rules.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "ceprule.delete",
    "params": [
        "3",
        "4"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "cep_ruleids": [
            "3",
            "4"
        ]
    },
    "id": 1
}
```

Source

CCepRule::delete() in *ui/include/classes/api/services/CCepRule.php*.

ceprule.get

Description

`integer/array ceprule.get(object parameters)`

The method allows you to retrieve CEP rules according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
cep_ruleids	ID/array	Return only CEP rules with the given IDs.
selectFilter	query	Return a `filter` property with the CEP rule conditions.
selectOperations	query	Return an `operations` property with the CEP rule operations.
selectWindow	query	Return a `[window]`(/manual/api/reference/ceprule/object#cep-rule-time-window) property with the CEP rule time window.
sortfield	string/array	Sort the result by the given properties. Possible values: `cep_ruleid`, `name`, `stop`, `sortorder`, `status`.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve CEP rules

Retrieve enabled CEP rules, including conditions, windows, and operations.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "ceprule.get",
    "params": {
        "output": "extend",
        "filter": {"status": 0},
        "selectFilter": "extend",
        "selectWindow": "extend",
        "selectOperations": "extend"
    },
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "cep_ruleid": "2",
```

```
      "name": "Group events by cause",
      "stop": "0",
      "sortorder": "1",
      "description": "",
      "status": "0",
      "window_type": "2",
      "error": "",
      "filter": {
        "evaltype": "0",
        "eval_formula": "A",
        "formula": "",
        "conditions": [
          {
            "type": "3",
            "operator": "2",
            "event_name": "Printer",
            "tag": "",
            "tag_value": "",
            "severity": "0",
            "host_name": "",
            "host_group_name": "",
            "time_period": "",
            "formulaid": "A"
          }
        ]
      },
      "window": {
        "duration": "1d",
        "capacity": "0",
        "group_by_host_group": "0",
        "group_by_host": "1",
        "group_by_tags": "1",
        "tags": [
          "cause"
        ],
        "event_count_tag": "symptoms",
        "script": ""
      },
      "operations": []
    },
    {
      "cep_ruleid": "10",
      "name": "Suppress consecutive events",
      "stop": "1",
      "sortorder": "2",
      "description": "",
      "status": "0",
      "window_type": "3",
      "error": "",
      "filter": {
        "evaltype": "0",
        "eval_formula": "A",
        "formula": "",
        "conditions": [
          {
            "type": "3",
            "operator": "2",
            "event_name": "Printer",
            "tag": "",
            "tag_value": "",
            "severity": "0",
            "host_name": "",
```

```
          "host_group_name": "",
          "time_period": "",
          "formulaid": "A"
        }
      ]
    },
    "window": {
      "duration": "3m",
      "capacity": "0",
      "group_by_host_group": "0",
      "group_by_host": "1",
      "group_by_tags": "1",
      "tags": [
        "cause"
      ],
      "event_count_tag": "",
      "script": ""
    },
    "operations": [
      {
        "sortorder": "1",
        "execute_when": "1",
        "type": "7",
        "event_name": "",
        "severity": "0",
        "suppress_duration": "6m",
        "tag": "",
        "new_tag": "",
        "tag_value": "",
        "scope": "0",
        "filter": {
          "evaltype": "0",
          "eval_formula": "A",
          "formula": "",
          "conditions": [
            {
              "type": "33",
              "operator": "10",
              "tag": "",
              "tag_value": "",
              "formulaid": "A"
            }
          ]
        },
        "scope_filter": {
          "evaltype": "0",
          "eval_formula": "",
          "formula": "",
          "conditions": []
        }
      },
      {
        "sortorder": "2",
        "execute_when": "2",
        "type": "8",
        "event_name": "",
        "severity": "0",
        "suppress_duration": "0",
        "tag": "",
        "new_tag": "",
        "tag_value": "",
        "scope": "0",
```

```
          "filter": {
            "evaltype": "0",
            "eval_formula": "A",
            "formula": "",
            "conditions": [
              {
                "type": "35",
                "operator": "10",
                "tag": "",
                "tag_value": "",
                "formulaid": "A"
              }
            ]
          },
          "scope_filter": {
            "evaltype": "0",
            "eval_formula": "",
            "formula": "",
            "conditions": []
          }
        }
      ]
    }
  ],
  "id": 1
}
```

See also

- CEP rule

Source

CCepRule::get() in *ui/include/classes/api/services/CCepRule.php*.

ceprule.resettimewindows

Description

`object ceprule.resettimewindows(object/array cep_ruleid)`

This method allows you to reset time windows for an existing CEP rule.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` CEP rule object must have a `cep_ruleid` property defined.

Return values

`(object)` Returns an object containing the ID of the CEP rule whose time windows have been reset under the `cep_ruleid` property.

Examples

Reset time windows for CEP rule

Request:

```
{
    "jsonrpc": "2.0",
    "method": "ceprule.resettimewindows",
    "params": {
        "cep_ruleid": "1"
    },
```

```
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
      "cep_ruleid": "1"
    },
    "id": 1
}
```

See also

- CEP rule

Source

CCepRule::resettimewindows() in *ui/include/classes/api/services/CCepRule.php*.

ceprule.update

Description

`object ceprule.update(object/array CEP rules)`

This method allows you to update existing CEP rules.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` CEP rule properties to be updated.

The `cep_ruleid` property must be defined for each CEP rule, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard CEP rule properties, the method accepts the following parameters.

Parameter	Type	Description
filter	object	CEP rule filter object to replace the current filter.
window	array	CEP rule window to be updated. Property behavior: - *required* if `window_type` is changed to ”Cause and symptom grouping” and ”group_by_host_group”, ”group_by_host”, ”group_by_tags” are all set to NO; - *required* if `window_type` is changed to ”Event pattern match”; - *supported* if `window_type` is ”Simple (multiple)” or ”Simple (once)”
operations	array	CEP rule operations to replace existing operations. Property behavior: - *required* if `window_type` is changed to ”None” (0), ”Simple (multiple)” (1), ”Simple (once)” (3) or ”Event pattern match” (4).

Return values

`(object)` Returns an object containing the IDs of the updated CEP rules under the `cep_ruleids` property.

Examples

Disable CEP rule

Request:

```
{
    "jsonrpc": "2.0",
    "method": "ceprule.update",
    "params": {
        "cep_ruleid": "1",
        "status": "1"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "cep_ruleids": [
            "1"
        ]
    },
    "id": 1
}
```

Replace condition

Request:

```
{
    "jsonrpc": "2.0",
    "method": "ceprule.update",
    "params": {
        "cep_ruleid": "3",
        "filter": {
        "evaltype": "0",
        "formula": "",
            "conditions": [
                {
                    "type": "3",
                    "operator": "2",
                    "event_name": "Inaccessible port",
                    "severity": "0"
                }
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "cep_ruleids": [
            "1"
        ]
    },
    "id": 1
}
```

See also

- CEP rule

Source

CCepRule::update() in *ui/include/classes/api/services/CCepRule.php*.

Configuration

This class is designed to export and import Zabbix configuration data.

Available methods:

- configuration.export - export configuration data
- configuration.import - import configuration data
- configuration.importcompare - compare import file with current system elements

configuration.export

Description

`string configuration.export(object parameters)`

This method allows you to export configuration data as a serialized string.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the objects to be exported and the format to use.

Parameter	Type	Description
format	string	Format in which the data must be exported. Possible values: `yaml` - YAML; `xml` - XML; `json` - JSON; `raw` - unprocessed PHP array. Parameter behavior: - *required*
prettyprint	boolean	Make the output more human readable by adding indentation. Possible values: `true` - add indentation; `false` - *(default)* do not add indentation.
options	object	Objects to be exported. The `options` object has the following parameters: `host_groups` - `(array)` IDs of host groups to export; `hosts` - `(array)` IDs of hosts to export; `images` - `(array)` IDs of images to export; `maps` - `(array)` IDs of maps to export; `mediaTypes` - `(array)` IDs of media types to export; `template_groups` - `(array)` IDs of template groups to export; `templates` - `(array)` IDs of templates to export; `dashboards` - `(array)` IDs of dashboards to export; `global_regexes` - `(array)` IDs of global regular expressions to export. *Admin* and *User* type users may export only those objects for which they have *read-only* or *read-write* permission, as well as images, but not media types. Parameter behavior: - *required*.

Return values

`(string)` Returns a serialized string containing the requested configuration data.

Examples

Exporting a template

Export the configuration of template "10571" as an XML string.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "configuration.export",
    "params": {
        "options": {
            "templates": [
                "10571"
            ]
        },
        "format": "xml"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": "<?xml version=\"1.0\" encoding=\"UTF-8\"?>\n<zabbix_export><version>8.2</version><template_
    "id": 1
}
```

Source

CConfiguration::export() in *ui/include/classes/api/services/CConfiguration.php*.

configuration.import

Description

`boolean configuration.import(object parameters)`

This method allows you to import configuration data from a serialized string.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters containing the data to import and rules how the data should be handled.

Parameter	Type	Description
format	string	Format of the serialized string. Possible values: `yaml` - YAML; `xml` - XML; `json` - JSON. Parameter behavior: - *required*
source	string	Serialized string containing the configuration data. Parameter behavior: - *required*

Parameter	Type	Description
rules	object	Rules on how new and existing objects should be imported. *Admin* type users may import only those objects for which they have *read-write* permission, as well as maps. For example, a host and its entities (items, triggers, graphs, etc.) may be imported only if the user's user group has permission to the host group that the imported host will belong to. Images and media types cannot be imported by *Admin* type users. The `rules` parameter is described in detail in the table below. Parameter behavior: - *required*
returnMissingObjects	boolean	Rules if response object should contain information about missing objects; default: `false`.

Attention:
If no rules are given, the configuration will not be updated.

The `rules` object supports the following parameters.

Parameter	Type	Description
dashboards	object	Rules on how to import dashboards. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new dashboards will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`and a particular importable dashboard is present in the system, it will be fully updated. Pages and widgets will be fully rewritten by the new ones; default: `false`.
discoveryRules	object	Rules on how to import LLD rules. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new LLD rules will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing LLD rules will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, LLD rules not present in the imported data will be deleted from the database; default: `false`.
global_regexes	object	Rules on how to import global regular expressions. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new global regular expressions will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing global regular expressions will be updated, with their subexpressions fully overwritten; default: `false`.
graphs	object	Rules on how to import graphs. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new graphs will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing graphs will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, graphs not present in the imported data will be deleted from the database; default: `false`.

Parameter	Type	Description
host_groups	object	Rules on how to import host groups. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new host groups will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing host groups will be updated; default: `false`.
hosts	object	Rules on how to import hosts. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new hosts will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing hosts will be updated; default: `false`.
httptests	object	Rules on how to import web scenarios. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new web scenarios will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing web scenarios will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, web scenarios not present in the imported data will be deleted from the database; default: `false`.
images	object	Rules on how to import images. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new images will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing images will be updated; default: `false`.
items	object	Rules on how to import items. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new items will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing items will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, items not present in the imported data will be deleted from the database; default: `false`.
maps	object	Rules on how to import maps. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new maps will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing maps will be updated; default: `false`.
mediaTypes	object	Rules on how to import media types. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new media types will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing media types will be updated; default: `false`.

Parameter	Type	Description
templateDashboards	object	Rules on how to import template dashboards. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new template dashboards will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing template dashboards will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, template dashboards not present in the imported data will be deleted from the database; default: `false`.
template_groups	object	Rules on how to import template groups. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new template groups will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing template groups will be updated; default: `false`.
templateLinkage	object	Rules on how to import template links. Supported parameters: `createMissing` - `(boolean)` if set to `true`, templates that are not linked to the host or template being imported, but are present in the imported data, will be linked; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, templates that are linked to the host or template being imported, but are not present in the imported data, will be unlinked without removing entities (items, triggers, etc.) inherited from the unlinked templates; default: `false`.
templates	object	Rules on how to import templates. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new templates will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing templates will be updated; default: `false`.
triggers	object	Rules on how to import triggers. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new triggers will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing triggers will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, triggers not present in the imported data will be deleted from the database; default: `false`.
valueMaps	object	Rules on how to import host or template value maps. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new value maps will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing value maps will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, value maps not present in the imported data will be deleted from the database; default: `false`.

Return values

`(boolean)` Returns `true` if importing has been successful.

Examples

Importing a template

Import the template configuration contained in the XML string. If any items or triggers in the XML string are missing, they will be deleted from the database, and everything else will be left unchanged.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "configuration.import",
    "params": {
        "format": "xml",
        "rules": {
            "templates": {
                "createMissing": true,
                "updateExisting": true
            },
            "items": {
                "createMissing": true,
                "updateExisting": true,
                "deleteMissing": true
            },
            "triggers": {
                "createMissing": true,
                "updateExisting": true,
                "deleteMissing": true
            },
            "valueMaps": {
                "createMissing": true,
                "updateExisting": false
            }
        },
        "source": "<?xml version=\"1.0\" encoding=\"UTF-8\"?>\n<zabbix_export><version>8.2</version><templ
    },
    "id": 1
}
```


Response:


```
{
    "jsonrpc": "2.0",
    "result": true,
    "id": 1
}
```


Source

CConfiguration::import() in *ui/include/classes/api/services/CConfiguration.php*.

configuration.importcompare

Description

`array configuration.importcompare(object parameters)`

This method allows you to compare import file with current system elements and shows what will be changed if this import file will be imported.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters containing the possible data to import and rules how the data should be handled.

Parameter	Type	Description
format	string	Format of the serialized string. Possible values: `yaml` - YAML; `xml` - XML; `json` - JSON. Parameter behavior: - *required*
source	string	Serialized string containing the configuration data. Parameter behavior: - *required*
rules	object	Rules on how new and existing objects should be compared. *Admin* and *User* type users may compare new objects only with existing ones for which they have *read-only* or *read-write* permission. The `rules` parameter is described in detail in the table below. Parameter behavior: - *required*
returnMissingObjects	boolean	Rules if response object should contain information about missing objects; default: `false`.

Attention:
If no rules are given, there will be nothing to update and result will be empty.

Note:
Comparison will be done only for host groups and templates. Triggers and graphs will be compared only for imported templates, any other will be considered as "new".

The `rules` object supports the following parameters.

Parameter	Type	Description
dashboards	object	Rules on how to import dashboards. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new dashboards will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`and a particular importable dashboard is present in the system, it will be fully updated. Pages and widgets will be fully rewritten by the new ones; default: `false`.
discoveryRules	object	Rules on how to import LLD rules. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new LLD rules will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing LLD rules will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, LLD rules not present in the imported data will be deleted from the database; default: `false`.

Parameter	Type	Description
global_regexes	object	Rules on how to import global regular expressions. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new global regular expressions will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing global regular expressions will be updated, with their subexpressions fully overwritten; default: `false`.
graphs	object	Rules on how to import graphs. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new graphs will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing graphs will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, graphs not present in the imported data will be deleted from the database; default: `false`.
host_groups	object	Rules on how to import host groups. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new host groups will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing host groups will be updated; default: `false`.
httptests	object	Rules on how to import web scenarios. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new web scenarios will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing web scenarios will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, web scenarios not present in the imported data will be deleted from the database; default: `false`.
images	object	Rules on how to import images. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new images will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing images will be updated; default: `false`. This parameter will make no difference to the output. It is allowed only for consistency with `configuration.import`.
items	object	Rules on how to import items. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new items will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing items will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, items not present in the imported data will be deleted from the database; default: `false`.

Parameter	Type	Description
maps	object	Rules on how to import maps. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new maps will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing maps will be updated; default: `false`. This parameter will make no difference to the output. It is allowed only for consistency with `configuration.import`.
mediaTypes	object	Rules on how to import media types. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new media types will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing media types will be updated; default: `false`. This parameter will make no difference to the output. It is allowed only for consistency with `configuration.import`.
templateDashboards	object	Rules on how to import template dashboards. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new template dashboards will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing template dashboards will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, template dashboards not present in the imported data will be deleted from the database; default: `false`.
template_groups	object	Rules on how to import template groups. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new template groups will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing template groups will be updated; default: `false`.
templateLinkage	object	Rules on how to import template links. Supported parameters: `createMissing` - `(boolean)` if set to `true`, templates that are not linked to the host or template being imported, but are present in the imported data, will be linked; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, templates that are linked to the host or template being imported, but are not present in the imported data, will be unlinked without removing entities (items, triggers, etc.) inherited from the unlinked templates; default: `false`.
templates	object	Rules on how to import templates. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new templates will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing templates will be updated; default: `false`.

Parameter	Type	Description
triggers	object	Rules on how to import triggers. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new triggers will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing triggers will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, triggers not present in the imported data will be deleted from the database; default: `false`.
valueMaps	object	Rules on how to import host or template value maps. Supported parameters: `createMissing` - `(boolean)` if set to `true`, new value maps will be created; default: `false`; `updateExisting` - `(boolean)` if set to `true`, existing value maps will be updated; default: `false`; `deleteMissing` - `(boolean)` if set to `true`, value maps not present in the imported data will be deleted from the database; default: `false`.

Return values

`(array)` Returns an array with changes in configuration, that will be made.

Examples

Comparing the import of a template

Compare the template contained in the XML string to the current system elements, and show what will be changed if this template will be imported.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "configuration.importcompare",
    "params": {
        "format": "xml",
        "rules": {
            "discoveryRules": {
                "createMissing": true,
                "updateExisting": true,
                "deleteMissing": true
            },
            "graphs": {
                "createMissing": true,
                "updateExisting": true,
                "deleteMissing": true
            },
            "host_groups": {
                "createMissing": true,
                "updateExisting": true
            },
            "template_groups": {
                "createMissing": true,
                "updateExisting": true
            },
            "httptests": {
                "createMissing": true,
                "updateExisting": true,
                "deleteMissing": true
            },
            "items": {
                "createMissing": true,
                "updateExisting": true,
```

```
                "deleteMissing": true
            },
            "templateLinkage": {
                "createMissing": true,
                "deleteMissing": true
            },
            "templates": {
                "createMissing": true,
                "updateExisting": true
            },
            "templateDashboards": {
                "createMissing": true,
                "updateExisting": true,
                "deleteMissing": true
            },
            "triggers": {
                "createMissing": true,
                "updateExisting": true,
                "deleteMissing": true
            },
            "valueMaps": {
                "createMissing": true,
                "updateExisting": true,
                "deleteMissing": true
            }
        },
        "source": "<?xml version=\"1.0\" encoding=\"UTF-8\"?>\n<zabbix_export><version>8.2</version><templ
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templates": {
            "updated": [
                {
                    "before": {
                        "uuid": "5aef0444a82a4d8cb7a95dc4c0c85330",
                        "template": "New template",
                        "name": "New template",
                        "groups": [
                            {
                                "name": "Templates"
                            }
                        ]
                    },
                    "after": {
                        "uuid": "5aef0444a82a4d8cb7a95dc4c0c85330",
                        "template": "New template",
                        "name": "New template",
                        "groups": [
                            {
                                "name": "Templates"
                            }
                        ]
                    },
                    "items": {
                        "added": [
                            {
                                "after": {
```

```
                                "uuid": "648006da5971424ead0c47ddbbf1ea2e",
                                "name": "CPU utilization",
                                "key": "system.cpu.util",
                                "value_type": "FLOAT",
                                "units": "%"
                            },
                            "triggers": {
                                "added": [
                                    {
                                        "after": {
                                            "uuid": "736225012c534ec480c2a66a91322ce0",
                                            "expression": "avg(/New template/system.cpu.util,3m)>70",
                                            "name": "CPU utilization too high on 'New host' for 3 minu
                                            "priority": "WARNING"
                                        }
                                    }
                                ]
                            }
                        }
                    ],
                    "removed": [
                        {
                            "before": {
                                "uuid": "6805d4c39a624a8bab2cc8ab63df1ab3",
                                "name": "CPU load",
                                "key": "system.cpu.load",
                                "value_type": "FLOAT"
                            },
                            "triggers": {
                                "removed": [
                                    {
                                        "before": {
                                            "uuid": "ab4c2526c2bc42e48a633082255ebcb3",
                                            "expression": "avg(/New template/system.cpu.load,3m)>2",
                                            "name": "CPU load too high on 'New host' for 3 minutes",
                                            "priority": "WARNING"
                                        }
                                    }
                                ]
                            }
                        }
                    ],
                    "updated": [
                        {
                            "before": {
                                "uuid": "7f1e6f1e48aa4a128e5b6a958a5d11c3",
                                "name": "Zabbix agent ping",
                                "key": "agent.ping"
                            },
                            "after": {
                                "uuid": "7f1e6f1e48aa4a128e5b6a958a5d11c3",
                                "name": "Zabbix agent ping",
                                "key": "agent.ping",
                                "delay": "3m"
                            }
                        }
                    ]
                }
            }
        ]
    }
},
```

```
    "id": 1
}
```

Source

CConfiguration::importcompare() in *ui/include/classes/api/services/CConfiguration.php*.

Connector

This class is designed to work with connectors.

Object references:

- Connector
- Tag filter

Available methods:

- connector.create - create new connectors
- connector.delete - delete connectors
- connector.get - retrieve connectors
- connector.update - update connectors

Connector object

The following objects are directly related to the `connector` API.

Connector

The connector object has the following properties.

Property	Type	Description
connectorid	ID	ID of the connector. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the connector. Property behavior: - *required* for create operations
url	string	Endpoint URL, that is, URL of the receiver. User macros are supported. Property behavior: - *required* for create operations
protocol	integer	Communication protocol. Possible values: 0 - *(default)* Zabbix Streaming Protocol v1.0.
data_type	integer	Data type. Possible values: 0 - *(default)* Item values; 1 - Events.

Property	Type	Description
item_value_type	integer	A sum of item value types to be sent. Possible values: 1 - Numeric (float); 2 - Character; 4 - Log; 8 - Numeric (unsigned); 16 - Text; 32 - Binary; 64 - JSON. Default: 31 - All item types (except binary and JSON). Property behavior: - *supported* if `data_type` is set to "Item values".
max_records	integer	Maximum number of events or items that can be sent within one message. Possible values: 0-2147483647 (max value of 32-bit signed integer). Default: 0 - Unlimited.
max_senders	integer	Number of sender processes to run for this connector. Possible values: 1-100. Default: 1.
max_attempts	integer	Number of attempts. Possible values: 1-5. Default: 1.
attempt_interval	string	The interval between retry attempts. Accepts seconds. Possible values: 0s-10s. Default: 5s. Property behavior: - *supported* if `max_attempts` is greater than 1.
timeout	string	Timeout. Time suffixes are supported (e.g., `30s`, `1m`). User macros are supported. Possible values: 1s-60s. Default: 5s.
http_proxy	string	HTTP(S) proxy connection string given as *[protocol]://[username[:password]@]proxy.example.com[:port]*. User macros are supported.
authtype	integer	HTTP authentication method. Possible values: 0 - *(default)* None; 1 - Basic; 2 - NTLM; 3 - Kerberos; 4 - Digest; 5 - Bearer.

Property	Type	Description
username	string	User name. User macros are supported. Property behavior: - *supported* if `authtype` is set to ”Basic”, ”NTLM”, ”Kerberos”, or ”Digest”
password	string	Password. User macros are supported. Property behavior: - *supported* if `authtype` is set to ”Basic”, ”NTLM”, ”Kerberos”, or ”Digest”
token	string	Bearer token. User macros are supported. Property behavior: - *required* if `authtype` is set to ”Bearer”
verify_peer	integer	Whether to validate that the host’s certificate is authentic. Possible values: 0 - Do not validate; 1 - *(default)* Validate.
verify_host	integer	Whether to validate that the host name for the connection matches the one in the host’s certificate. Possible values: 0 - Do not validate; 1 - *(default)* Validate.
ssl_cert_file	string	Public SSL Key file path. User macros are supported.
ssl_key_file	string	Private SSL Key file path. User macros are supported.
ssl_key_password	string	Password for SSL Key file. User macros are supported.
description	text	Description of the connector.
status	integer	Whether the connector is enabled. Possible values: 0 - Disabled; 1 - *(default)* Enabled.
tags_evaltype	integer	Tag evaluation method. Possible values: 0 - *(default)* And/Or; 2 - Or.

Tag filter

The tag filter allows you to export only matching item values or events. If not set, everything will be exported. The tag filter object has the following properties.

Property	Type	Description
tag	string	Tag name. Property behavior: - *required*

Property	Type	Description
operator	integer	Condition operator. Possible values: 0 - *(default)* Equals; 1 - Does not equal; 2 - Contains; 3 - Does not contain; 12 - Exists; 13 - Does not exist.
value	string	Tag value.

connector.create

Description

`object connector.create(object/array connectors)`

This method allows you to create new connector objects.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Connector objects to create.

Additionally to the standard connector properties, the method accepts the following parameters.

Parameter	Type	Description
tags	array	Connector tag filter.

Return values

`(object)` Returns an object containing the IDs of the created connectors under the `connectorids` property. The order of the returned IDs matches the order of the passed connectors.

Examples

Creating a connector

Create a connector to export trigger events with a tag filter. HTTP authentication will be performed using Bearer token.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "connector.create",
    "params": [
        {
            "name": "Export of events",
            "data_type": 1,
            "url": "{$DATA_EXPORT_URL}",
            "authtype": 5,
            "token": "{$DATA_EXPORT_BEARER_TOKEN}",
            "tags": [
                {
                    "tag": "service",
                    "operator": 0,
                    "value": "mysqld"
                }
            ]
        }
```

```
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "connectorid": [
            "3"
        ]
    },
    "id": 1
}
```

Source

CConnector::create() in *ui/include/classes/api/services/CConnector.php*.

connector.delete

Description

`object connector.delete(array connectorids)`

This method allows you to delete connector entries.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the connectors to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted connectors under the `connectorids` property.

Examples

Deleting multiple connectors

Delete two connector entries.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "connector.delete",
    "params": [
        3,
        5
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "connectorids": [
            "3",
            "5"
        ]
    },
```

```
    "id": 1
}
```

Source

CConnector::delete() in *ui/include/classes/api/services/CConnector.php*.

connector.get

Description

`integer/array connector.get(object parameters)`

The method allows you to retrieve connector objects according to the given parameters.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
connectorids	ID/array	Return only connectors with the given IDs.
selectTags	query	Return a `tags` property with connector tag filter. Supports `count`.
sortfield	string/array	Sort the result by the given properties. Possible values: `connectorid`, `name`, `data_type`, `status`.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving all connectors

Retrieve all data about all connectors and their properties.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "connector.get",
    "params": {
        "output": "extend",
        "selectTags": ["tag", "operator", "value"],
        "preservekeys": true
```

```
    },
    "id": 1
}
```

Response:


```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "connectorid": "1",
            "name": "Export of item values",
            "protocol": "0",
            "data_type": "0",
            "url": "{$DATA_EXPORT_VALUES_URL}",
            "item_value_type": "31",
            "authtype": "4",
            "username": "{$DATA_EXPORT_VALUES_USERNAME}",
            "password": "{$DATA_EXPORT_VALUES_PASSWORD}",
            "token": "",
            "max_records": "0",
            "max_senders": "4",
            "max_attempts": "2",
            "attempt_interval": "10s",
            "timeout": "10s",
            "http_proxy": "{$DATA_EXPORT_VALUES_PROXY}",
            "verify_peer": "1",
            "verify_host": "1",
            "ssl_cert_file": "{$DATA_EXPORT_VALUES_SSL_CERT_FILE}",
            "ssl_key_file": "{$DATA_EXPORT_VALUES_SSL_KEY_FILE}",
            "ssl_key_password": "",
            "description": "",
            "status": "1",
            "tags_evaltype": "0",
            "tags": [
                {
                    "tag": "component",
                    "operator": "0",
                    "value": "memory"
                }
            ]
        },
        {
            "connectorid": "2",
            "name": "Export of events",
            "protocol": "0",
            "data_type": "1",
            "url": "{$DATA_EXPORT_EVENTS_URL}",
            "item_value_type": "31",
            "authtype": "5",
            "username": "",
            "password": "",
            "token": "{$DATA_EXPORT_EVENTS_BEARER_TOKEN}",
            "max_records": "0",
            "max_senders": "2",
            "max_attempts": "1",
            "attempt_interval": "5s",
            "timeout": "5s",
            "http_proxy": "",
            "verify_peer": "1",
            "verify_host": "1",
            "ssl_cert_file": "",
            "ssl_key_file": "",
```

```
            "ssl_key_password": "",
            "description": "",
            "status": "1",
            "tags_evaltype": "0",
            "tags": [
                {
                    "tag": "scope",
                    "operator": "0",
                    "value": "performance"
                }
            ]
        }
    ],
    "id": 1
}
```

Source

CConnector:get() in *ui/include/classes/api/services/CConnector.php*.

connector.update

Description

`object connector.update(object/array connectors)`

This method allows you to update existing connectors.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Connector properties to be updated.

The `connectorid` property must be defined for each connector, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard connector properties, the method accepts the following parameters.

Parameter	Type	Description
tags	array	Connector tag filter to replace the current tag filter.

Return values

`(object)` Returns an object containing the IDs of the updated connectors under the `connectorids` property.

Examples

Changing HTTP authentication type

Change HTTP authentication type to Bearer for connector with ID "3".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "connector.update",
    "params": {
        "connectorid": 3,
        "authtype": 5,
        "token": "{$DATA_EXPORT_BEARER_TOKEN}"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "connectorids": [
            "3"
        ]
    },
    "id": 1
}
```

Updating tag filter

Change tag filter for connector with ID ”5”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "connector.update",
    "params": [
        {
            "connectorid": 5,
            "tags_evaltype": 2,
            "tags": [
                {
                    "tag": "service",
                    "operator": 0,
                    "value": "mysqld"
                },
                {
                    "tag": "error",
                    "operator": 12,
                    "value": ""
                }
            ]
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "connectorids": [
            "5"
        ]
    },
    "id": 1
}
```

Source

CConnector::update() in *ui/include/classes/api/services/CConnector.php*.

Correlation

This class is designed to work with correlations.

:::noteimportant This feature is deprecated and may be removed in future releases. See Complex event processing for updated event correlation features. :::note

Object references:

- Correlation

- Correlation operation
- Correlation filter
  - Correlation filter condition

Available methods:

- correlation.create - create new correlations
- correlation.delete - delete correlations
- correlation.get - retrieve correlations
- correlation.update - update correlations

Correlation object

The following objects are directly related to the `correlation` API.

Correlation

The correlation object has the following properties.

Property	Type	Description
correlationid	ID	ID of the correlation. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the correlation. Property behavior: - *required* for create operations
description	string	Description of the correlation.
status	integer	Whether the correlation is enabled or disabled. Possible values: 0 - *(default)* enabled; 1 - disabled.

Correlation operation

The correlation operation object defines an operation that will be performed when a correlation is executed. It has the following properties.

Property	Type	Description
type	integer	Type of operation. Possible values: 0 - close old events; 1 - close new event. Property behavior: - *required*

Correlation filter

The correlation filter object defines a set of conditions that must be met to perform the configured correlation operations. It has the following properties.

Property	Type	Description
conditions	array	Set of filter conditions to use for filtering results. The conditions will be sorted in the order of their placement in the formula. Property behavior: - *required*
evaltype	integer	Filter condition evaluation method. Possible values: 0 - And/Or; 1 - And; 2 - Or; 3 - Custom expression. Property behavior: - *required*
eval_formula	string	Generated expression that will be used for evaluating filter conditions. The expression contains IDs that reference specific filter conditions by its `formulaid`. The value of `eval_formula` is equal to the value of `formula` for filters with a custom expression. Property behavior: - *read-only*
formula	string	User-defined expression to be used for evaluating conditions of filters with a custom expression. The expression must contain IDs that reference specific filter conditions by its `formulaid`. The IDs used in the expression must exactly match the ones defined in the filter conditions: no condition can remain unused or omitted. Property behavior: - *required* if `evaltype` is set to ”custom expression”

Correlation filter condition

The correlation filter condition object defines a specific condition that must be checked before running the correlation operations.

Property	Type	Description
type	integer	Type of condition. Possible values: 0 - old event tag; 1 - new event tag; 2 - new event host group; 3 - event tag pair; 4 - old event tag value; 5 - new event tag value. Property behavior: - *required*
tag	string	Event tag (old or new). Property behavior: - *required* if `type` is set to ”old event tag”, ”new event tag”, ”old event tag value”, or ”new event tag value”
groupid	ID	ID of the host group. Property behavior: - *required* if `type` is set to ”new event host group”

Property	Type	Description
oldtag	string	Old event tag. Property behavior: - *required* if `type` is set to ”event tag pair”
newtag	string	Old event tag. Property behavior: - *required* if `type` is set to ”event tag pair”
value	string	Event tag (old or new) value. Property behavior: - *required* if `type` is set to ”old event tag value” or ”new event tag value”
formulaid	string	Arbitrary unique ID that is used to reference the condition from a custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward.
operator	integer	Condition operator. Property behavior: - *required* if `type` is set to ”new event host group”, ”old event tag value”, or ”new event tag value”

Note:
To better understand how to use filters with various types of expressions, see examples on the correlation.get and correlation.create method pages.

The following operators and values are supported for each condition type.

Condition	Condition name	Supported operators	Expected value
2	Host group	=, <>	Host group ID.
4	Old event tag value	=, <>, like, not like	string
5	New event tag value	=, <>, like, not like	string

correlation.create

Description

`object correlation.create(object/array correlations)`

This method allows you to create new correlations.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Correlations to create.

Additionally to the standard correlation properties, the method accepts the following parameters.

Parameter	Type	Description
operations	array	Correlation operations to create for the correlation. Parameter behavior: - *required*

Parameter	Type	Description
filter	object	Correlation filter object for the correlation. Parameter behavior: *- required*

Return values

`(object)` Returns an object containing the IDs of the created correlations under the `correlationids` property. The order of the returned IDs matches the order of the passed correlations.

Examples

Create a new event tag correlation

Create a correlation using evaluation method `AND/OR` with one condition and one operation. By default, the correlation will be enabled.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "correlation.create",
    "params": {
        "name": "new event tag correlation",
        "filter": {
            "evaltype": 0,
            "conditions": [
                {
                    "type": 1,
                    "tag": "ok"
                }
            ]
        },
        "operations": [
            {
                "type": 0
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "correlationids": [
            "1"
        ]
    },
    "id": 1
}
```

Using a custom expression filter

Create a correlation that will use a custom filter condition. The formula IDs ”A” or ”B” have been chosen arbitrarily. Condition type will be ”Host group” with operator ”<>”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "correlation.create",
    "params": {
        "name": "new host group correlation",
```

```
        "description": "a custom description",
        "status": 0,
        "filter": {
            "evaltype": 3,
            "formula": "A or B",
            "conditions": [
                {
                    "type": 2,
                    "operator": 1,
                    "formulaid": "A"
                },
                {
                    "type": 2,
                    "operator": 1,
                    "formulaid": "B"
                }
            ]
        },
        "operations": [
            {
                "type": 1
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "correlationids": [
            "2"
        ]
    },
    "id": 1
}
```

See also

- Correlation filter
- Correlation operation

Source

CCorrelation::create() in *ui/include/classes/api/services/CCorrelation.php*.

correlation.delete

Description

`object correlation.delete(array correlationids)`

This method allows you to delete correlations.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the correlations to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted correlations under the `correlationids` property.

Example

Delete multiple correlations

Delete two correlations.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "correlation.delete",
    "params": [
        "1",
        "2"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "correlationids": [
            "1",
            "2"
        ]
    },
    "id": 1
}
```

Source

CCorrelation::delete() in *ui/include/classes/api/services/CCorrelation.php*.

correlation.get

Description

`integer/array correlation.get(object parameters)`

The method allows you to retrieve correlations according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
correlationids	ID/array	Return only correlations with the given IDs.
selectFilter	query	Return a `filter` property with the correlation conditions.
selectOperations	query	Return an `operations` property with the correlation operations.
sortfield	string/array	Sort the result by the given properties. Possible values: `correlationid`, `name`, `status`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	

Parameter	Type	Description
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve correlations

Retrieve all configured correlations together with correlation conditions and operations. The filter uses the "and/or" evaluation type, so the `formula` property is empty and `eval_formula` is generated automatically.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "correlation.get",
    "params": {
        "output": "extend",
        "selectOperations": "extend",
        "selectFilter": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "correlationid": "1",
            "name": "Correlation 1",
            "description": "",
            "status": "0",
            "filter": {
                "evaltype": "0",
                "formula": "",
                "conditions": [
                    {
                        "type": "3",
                        "oldtag": "error",
                        "newtag": "ok",
                        "formulaid": "A"
                    }
                ],
                "eval_formula": "A"
            },
            "operations": [
                {
                    "type": "0"
                }
            ]
        }
    ],
    "id": 1
}
```

See also

- Correlation filter
- Correlation operation

Source

CCorrelation::get() in *ui/include/classes/api/services/CCorrelation.php*.

correlation.update

Description

`object correlation.update(object/array correlations)`

This method allows you to update existing correlations.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Correlation properties to be updated.

The `correlationid` property must be defined for each correlation, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard correlation properties, the method accepts the following parameters.

Parameter	Type	Description
filter	object	Correlation filter object to replace the current filter.
operations	array	Correlation operations to replace existing operations.

Return values

`(object)` Returns an object containing the IDs of the updated correlations under the `correlationids` property.

Examples

Disable correlation

Request:

```
{
    "jsonrpc": "2.0",
    "method": "correlation.update",
    "params": {
        "correlationid": "1",
        "status": "1"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "correlationids": [
            "1"
        ]
    },
    "id": 1
}
```

Replace conditions, but keep the evaluation method

Request:

```
{
    "jsonrpc": "2.0",
    "method": "correlation.update",
    "params": {
        "correlationid": "1",
        "filter": {
            "conditions": [
                {
                    "type": 3,
                    "oldtag": "error",
                    "newtag": "ok"
                }
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "correlationids": [
            "1"
        ]
    },
    "id": 1
}
```

See also

- Correlation filter
- Correlation operation

Source

CCorrelation::update() in *ui/include/classes/api/services/CCorrelation.php*.

Dashboard

This class is designed to work with dashboards.

Object references:

- Dashboard
- Dashboard page
    - Dashboard widget
        * Dashboard widget field
- Dashboard user group
- Dashboard user

Available methods:

- dashboard.create - create new dashboards
- dashboard.delete - delete dashboards
- dashboard.get - retrieve dashboards
- dashboard.update - update dashboards

Dashboard object

The following objects are directly related to the `dashboard` API.

Dashboard

The dashboard object has the following properties.

Property	Type	Description
dashboardid	ID	ID of the dashboard. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the dashboard. Property behavior: - *required* for create operations
userid	ID	ID of the user that is the owner of the dashboard.
private	integer	Type of dashboard sharing. Possible values: 0 - public dashboard; 1 - *(default)* private dashboard.
display_period	integer	Default slideshow interval (in seconds). Possible values: 10, 30, 60, 120, 600, 1800, 3600. Default: 30.
auto_start	integer	Auto start slideshow. Possible values: 0 - do not auto start slideshow; 1 - *(default)* auto start slideshow.

Dashboard page

The dashboard page object has the following properties.

Property	Type	Description
dashboard_pageid	ID	ID of the dashboard page. Property behavior: - *read-only*
name	string	Dashboard page name. Default: empty string.
display_period	integer	Dashboard slideshow interval (in seconds). Possible values: 0, 10, 30, 60, 120, 600, 1800, 3600. Default: 0 (will use the default slideshow interval).
widgets	array	Array of the dashboard widget objects.

Dashboard widget

The dashboard widget object has the following properties.

Property	Type	Description
widgetid	ID	ID of the dashboard widget. Property behavior: - *read-only*

Property	Type	Description
type	string	Type of the dashboard widget. Possible values: actionlog - Action log; clock - Clock; discovery - Discovery status; favgraphs - Favorite graphs; favmaps - Favorite maps; gauge - Gauge; geomap - Geomap; graph - Graph (classic); graphprototype - Graph prototype; honeycomb - Honeycomb; hostavail - Host availability; hostcard - Host card; hostnavigator - Host navigator; itemcard - Item card; itemhistory - Item history; itemnavigator - Item navigator; item - Item value; map - Map; navtree - Map Navigation Tree; piechart - Pie chart; problemhosts - Problem hosts; problems - Problems; problemsbysv - Problems by severity; scatterplot - Scatter plot; slareport - SLA report; svggraph - Graph; systeminfo - System information; tophosts - Top hosts; topitems - Top items; toptriggers - Top triggers; trigover - Trigger overview; url - URL; web - Web monitoring. Property behavior: - *required*
name	string	Custom widget name.
x	integer	A horizontal position from the left side of the dashboard. Possible values range from 0 to 71.
y	integer	A vertical position from the top of the dashboard. Possible values range from 0 to 63.
width	integer	The widget width. Possible values range from 1 to 72.
height	integer	The widget height. Possible values range from 1 to 64.
view_mode	integer	The widget view mode. Possible values: 0 - *(default)* default widget view; 1 - with hidden header;
fields	array	Array of the dashboard widget field objects. Property behavior: - see individual widgets in Dashboard widget fields

Dashboard widget field

The dashboard widget field object has the following properties.

Property	Type	Description
type	integer	Type of the widget field. Possible values: 0 - Integer; 1 - String; 2 - Host group; 3 - Host; 4 - Item; 5 - Item prototype; 6 - Graph; 7 - Graph prototype; 8 - Map; 9 - Service; 10 - SLA; 11 - User; 12 - Action; 13 - Media type. Property behavior: - *required*
name	string	Widget field name. Possible values: see Dashboard widget fields. Property behavior: - *required*
value	mixed	Widget field value depending on the type. Possible values: see Dashboard widget fields. Property behavior: - *required*

Dashboard user group

List of dashboard permissions based on user groups. It has the following properties.

Property	Type	Description
usrgrpid	ID	ID of the user group. Property behavior: - *required*
permission	integer	Type of permission level. Possible values: 2 - read only; 3 - read-write. Property behavior: - *required*

Dashboard user

List of dashboard permissions based on users. It has the following properties.

Property	Type	Description
userid	ID	ID of the user. Property behavior: - *required*
permission	integer	Type of permission level. Possible values: 2 - read only; 3 - read-write. Property behavior: - *required*

dashboard.create

Description

`object dashboard.create(object/array dashboards)`

This method allows you to create new dashboards.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Dashboards to create.

Additionally to the standard dashboard properties, the method accepts the following parameters.

Parameter	Type	Description
pages	array	Dashboard pages to be created for the dashboard. Dashboard pages will be ordered in the same order as specified. Parameter behavior: - *required*
users	array	Dashboard user shares to be created on the dashboard.
userGroups	array	Dashboard user group shares to be created on the dashboard.

Return values

`(object)` Returns an object containing the IDs of the created dashboards under the `dashboardids` property. The order of the returned IDs matches the order of the passed dashboards.

Examples

Creating a dashboard

Create a dashboard named ”My dashboard” with one Problems widget with tags and using two types of sharing (user group and user) on a single dashboard page.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
```

```
                "widgets": [
                    {
                        "type": "problems",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 1,
                                "name": "tags.0.tag",
                                "value": "service"
                            },
                            {
                                "type": 0,
                                "name": "tags.0.operator",
                                "value": 1
                            },
                            {
                                "type": 1,
                                "name": "tags.0.value",
                                "value": "zabbix_server"
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": "7",
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": "4",
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "2"
        ]
    },
    "id": 1
}
```

See also

- Dashboard page
- Dashboard widget
- Dashboard widget field
- Dashboard user
- Dashboard user group

Source

CDashboard::create() in *ui/include/classes/api/services/CDashboard.php*.

dashboard.delete

Description

`object dashboard.delete(array dashboardids)`

This method allows you to delete dashboards.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the dashboards to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted dashboards under the `dashboardids` property.

Examples

Deleting multiple dashboards

Delete two dashboards.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.delete",
    "params": [
        "2",
        "3"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "2",
            "3"
        ]
    },
    "id": 1
}
```

Source

CDashboard::delete() in *ui/include/classes/api/services/CDashboard.php*.

dashboard.get

Description

`integer/array dashboard.get(object parameters)`

The method allows you to retrieve dashboards according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
dashboardids	ID/array	Return only dashboards with the given IDs.
selectPages	query	Return a `pages` property with dashboard pages, correctly ordered.
selectUsers	query	Return a `users` property with users that the dashboard is shared with.
selectUserGroups	query	Return a `userGroups` property with user groups that the dashboard is shared with.
sortfield	string/array	Sort the result by the given properties. Possible values: `dashboardid`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving a dashboard by ID

Retrieve all data about dashboards "1" and "2".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.get",
    "params": {
        "output": "extend",
        "selectPages": "extend",
        "selectUsers": "extend",
        "selectUserGroups": "extend",
        "dashboardids": [
            "1",
            "2"
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "dashboardid": "1",
            "name": "Dashboard",
            "userid": "1",
            "private": "0",
            "display_period": "30",
            "auto_start": "1",
            "users": [],
            "userGroups": [],
            "pages": [
                {
                    "dashboard_pageid": "1",
                    "name": "",
                    "display_period": "0",
                    "widgets": [
                        {
                            "widgetid": "9",
                            "type": "systeminfo",
                            "name": "",
                            "x": "12",
                            "y": "8",
                            "width": "12",
                            "height": "5",
                            "view_mode": "0",
                            "fields": []
                        },
                        {
                            "widgetid": "8",
                            "type": "problemsbysv",
                            "name": "",
                            "x": "12",
                            "y": "4",
                            "width": "12",
                            "height": "4",
                            "view_mode": "0",
                            "fields": []
                        },
                        {
                            "widgetid": "7",
                            "type": "problemhosts",
                            "name": "",
                            "x": "12",
                            "y": "0",
                            "width": "12",
                            "height": "4",
                            "view_mode": "0",
                            "fields": []
                        },
                        {
                            "widgetid": "6",
                            "type": "discovery",
                            "name": "",
                            "x": "6",
                            "y": "9",
                            "width": "18",
                            "height": "4",
                            "view_mode": "0",
                            "fields": []
                        },
```

```
                    {
                        "widgetid": "5",
                        "type": "web",
                        "name": "",
                        "x": "0",
                        "y": "9",
                        "width": "18",
                        "height": "4",
                        "view_mode": "0",
                        "fields": []
                    },
                    {
                        "widgetid": "4",
                        "type": "problems",
                        "name": "",
                        "x": "0",
                        "y": "3",
                        "width": "12",
                        "height": "6",
                        "view_mode": "0",
                        "fields": []
                    },
                    {
                        "widgetid": "3",
                        "type": "favmaps",
                        "name": "",
                        "x": "8",
                        "y": "0",
                        "width": "12",
                        "height": "3",
                        "view_mode": "0",
                        "fields": []
                    },
                    {
                        "widgetid": "1",
                        "type": "favgraphs",
                        "name": "",
                        "x": "0",
                        "y": "0",
                        "width": "12",
                        "height": "3",
                        "view_mode": "0",
                        "fields": []
                    }
                ]
            },
            {
                "dashboard_pageid": "2",
                "name": "",
                "display_period": "0",
                "widgets": []
            },
            {
                "dashboard_pageid": "3",
                "name": "Custom page name",
                "display_period": "60",
                "widgets": []
            }
        ]
    },
    {
        "dashboardid": "2",
```

```
            "name": "My dashboard",
            "userid": "1",
            "private": "1",
            "display_period": "60",
            "auto_start": "1",
            "users": [
                {
                    "userid": "4",
                    "permission": "3"
                }
            ],
            "userGroups": [
                {
                    "usrgrpid": "7",
                    "permission": "2"
                }
            ],
            "pages": [
                {
                    "dashboard_pageid": "4",
                    "name": "",
                    "display_period": "0",
                    "widgets": [
                        {
                            "widgetid": "10",
                            "type": "problems",
                            "name": "",
                            "x": "0",
                            "y": "0",
                            "width": "12",
                            "height": "5",
                            "view_mode": "0",
                            "fields": [
                                {
                                    "type": "2",
                                    "name": "groupids",
                                    "value": "4"
                                }
                            ]
                        }
                    ]
                }
            ]
        }
    ],
    "id": 1
}
```


See also

- Dashboard page
- Dashboard widget
- Dashboard widget field
- Dashboard user
- Dashboard user group

Source

CDashboard::get() in *ui/include/classes/api/services/CDashboard.php*.

dashboard.update

Description

`object dashboard.update(object/array dashboards)`

This method allows you to update existing dashboards.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Dashboard properties to be updated.

The `dashboardid` property must be defined for each dashboard, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard dashboard properties, the method accepts the following parameters.

Parameter	Type	Description
pages	array	Dashboard pages to replace the existing dashboard pages. Dashboard pages are updated by the `dashboard_pageid` property. New dashboard pages will be created for objects without `dashboard_pageid` property and the existing dashboard pages will be deleted if not reused. Dashboard pages will be ordered in the same order as specified. Only the specified properties of the dashboard pages will be updated.
users	array	Dashboard user shares to replace the existing elements.
userGroups	array	Dashboard user group shares to replace the existing elements.

Return values

`(object)` Returns an object containing the IDs of the updated dashboards under the `dashboardids` property.

Examples

Renaming a dashboard

Rename a dashboard to "SQL server status".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.update",
    "params": {
        "dashboardid": "2",
        "name": "SQL server status"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "2"
        ]
    },
    "id": 1
}
```

Updating dashboard pages

Rename the first dashboard page, replace widgets on the second dashboard page and add a new page as the third one. Delete all other dashboard pages.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.update",
    "params": {
        "dashboardid": "2",
        "pages": [
            {
                "dashboard_pageid": 1,
                "name": "Renamed Page"
            },
            {
                "dashboard_pageid": 2,
                "widgets": [
                    {
                        "type": "clock",
                        "x": 0,
                        "y": 0,
                        "width": 12,
                        "height": 3
                    }
                ]
            },
            {
                "display_period": 60
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "2"
        ]
    },
    "id": 1
}
```

Change dashboard owner

Available only for admins and super admins.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.update",
    "params": {
        "dashboardid": "2",
        "userid": "1"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "2"
        ]
```

```
    },
    "id": 1
}
```

See also

- Dashboard page
- Dashboard widget
- Dashboard widget field
- Dashboard user
- Dashboard user group

Source

CDashboard::update() in *ui/include/classes/api/services/CDashboard.php*.

Dashboard widget fields

This page contains navigation links for dashboard widget parameters and possible property values for the respective dashboard widget field objects.

To see the parameters and property values for each widget, please use the sidebar to navigate to the corresponding widget page.

1 Action log

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Action log* widget in `dashboard.create` and `dashboard.update` methods.

> **Attention:**
> Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Action log* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Action log* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Recipients	11	userids.0	User ID. Note: To configure multiple users, create a dashboard widget field object for each user with an incremented number in the property name.
Actions	12	actionids.0	Action ID. Note: To configure multiple actions, create a dashboard widget field object for each action with an incremented number in the property name.
Media types	13	mediatypeids.0	Media type ID. Note: To configure multiple media types, create a dashboard widget field object for each media type with an incremented number in the property name.

Parameter	type	name	value
Status	0	statuses.0	0 - In progress; 1 - Sent/Executed; 2 - Failed. Note: To configure multiple values, create a dashboard widget field object for each value with an incremented number in the property name.
Search string	1	message	Any string value.
Time period	1	time_period._referenceDASHBOARD	`DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source. Default: `DASHBOARD._timeperiod` Alternatively, you can set the time period only in the *From* and *To* parameters.
From	1	time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.to` is set
To	1	time_period.to	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.from` is set
Sort entries by	0	sort_triggers	3 - Time (ascending); 4 - *(default)* Time (descending); 5 - Type (ascending); 6 - Type (descending); 7 - Status (ascending); 8 - Status (descending); 11 - Recipient (ascending); 12 - Recipient (descending).
Show lines	0	show_lines	Possible values range from 1 to 100. Default: 25.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Action log* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring an *Action log* widget

Configure an *Action log* widget that displays 10 entries of action operation details, sorted by time (in ascending order). In addition, display details only for those action operations that attempted to send an email to user "1", but were unsuccessful.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
```

```
                "widgets": [
                    {
                        "type": "actionlog",
                        "name": "Action log",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 0,
                                "name": "show_lines",
                                "value": 10
                            },
                            {
                                "type": 0,
                                "name": "sort_triggers",
                                "value": 3
                            },
                            {
                                "type": 11,
                                "name": "userids.0",
                                "value": 1
                            },
                            {
                                "type": 13,
                                "name": "mediatypeids.0",
                                "value": 1
                            },
                            {
                                "type": 0,
                                "name": "statuses.0",
                                "value": 2
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
```

```
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

2 Clock

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Clock* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Clock* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Clock* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - *(default)* 15 minutes.
Time type	0	time_type	0 - *(default)* Local time; 1 - Server time; 2 - Host time.
Clock type	0	clock_type	0 - *(default)* Analog; 1 - Digital.

The following parameters are supported if *Time type* is set to ”Host time”.

Parameter	type	name	value
Item	4	itemid.0	Item ID. Parameter behavior: - *required*

The following parameters are supported if *Clock type* is set to ”Digital”.

Parameter	type	name	value
Show	0	show.0	1 - Date; 2 - *(default)* Time; 3 - Time zone. Note: To configure multiple values, create a dashboard widget field object for each value with an incremented number in the property name.

Advanced configuration

The following advanced configuration parameters are supported if *Clock type* is set to ”Digital”.

Parameter	type	name	value
Background color	1	bg_color	Hexadecimal color code (e.g., `FF0000`). Default: "" (empty).

Date

The following advanced configuration parameters are supported if *Clock type* is set to ”Digital”, and *Show* is set to ”Date”.

Parameter	type	name	value
Bold	0	date_bold	0 - *(default)* Disabled; 1 - Enabled.
Color	1	date_color	Hexadecimal color code (e.g., `FF0000`). Default: "" (empty).

Time

The following advanced configuration parameters are supported if *Clock type* is set to ”Digital”, and *Show* is set to ”Time”.

Parameter	type	name	value
Bold	0	time_bold	0 - *(default)* Disabled; 1 - Enabled.
Color	1	time_color	Hexadecimal color code (e.g., `FF0000`). Default: "" (empty).
Seconds	0	time_sec	0 - Disabled; 1 - *(default)* Enabled.
Format	0	time_format	0 - *(default)* 24-hour; 1 - 12-hour.

Time zone

The following advanced configuration parameters are supported if *Clock type* is set to ”Digital”, and *Show* is set to ”Time zone”.

Parameter	type	name	value
Bold	0	tzone_bold	0 - *(default)* Disabled; 1 - Enabled.
Color	1	tzone_color	Hexadecimal color code (e.g., `FF0000`). Default: "" (empty).
Time zone	1	tzone_timezone	Valid timezone string (e.g., `Europe/Riga`, `system`, `UTC`, etc.). For the full list of supported time zones please refer to PHP documentation. Default: `local`. Parameter behavior: - *supported* if *Time type* is set to ”Local time” or ”Server time”
Format	0	tzone_format	0 - *(default)* Short; 1 - Full. Parameter behavior: - *supported* if *Time type* is set to ”Local time” or ”Server time”

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Clock* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Clock* widget

Configure a *Clock* widget that displays local date, time and time zone in a customized digital clock.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "clock",
                        "name": "Clock",
                        "x": 0,
                        "y": 0,
                        "width": 12,
                        "height": 3,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 0,
                                "name": "clock_type",
                                "value": 1
                            },
                            {
                                "type": 0,
                                "name": "show.0",
                                "value": 1
                            },
                            {
                                "type": 0,
                                "name": "show.1",
                                "value": 2
                            },
                            {
                                "type": 0,
                                "name": "show.2",
                                "value": 3
                            },
                            {
                                "type": 1,
                                "name": "date_color",
                                "value": "E1E1E1"
                            },
                            {
                                "type": 0,
                                "name": "time_bold",
                                "value": 1
                            },
                            {
                                "type": 1,
                                "name": "tzone_color",
                                "value": "E1E1E1"
                            },
                            {
                                "type": 1,
```

```
                                "name": "tzone_timezone",
                                "value": "Europe/Riga"
                            },
                            {
                                "type": 0,
                                "name": "tzone_format",
                                "value": 1
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

3 Discovery status

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Discovery status* widget in `dashboard.create` and `dashboard.update` methods.

Parameters

The following parameters are supported for the *Discovery status* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Discovery status* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring *Discovery status* widget

Configure a *Discovery status* widget with the refresh interval set to 15 minutes.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "discovery",
                        "name": "Discovery status",
                        "x": 0,
                        "y": 0,
                        "width": 18,
                        "height": 3,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 0,
                                "name": "rf_rate",
                                "value": 900
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- dashboard.create
- dashboard.update

4 Favorite graphs

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Favorite graphs* widget in dashboard.create and dashboard.update methods.

Parameters

The following parameters are supported for the *Favorite graphs* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - *(default)* 15 minutes.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Favorite graphs* widget. For more information on configuring a dashboard, see dashboard.create.

Configuring a *Favorite graphs* widget

Configure a *Favorite graphs* widget with the refresh interval set to 10 minutes.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "favgraphs",
                        "name": "Favorite graphs",
                        "x": 0,
                        "y": 0,
                        "width": 12,
```

```
                        "height": 3,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 0,
                                "name": "rf_rate",
                                "value": 600
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

5 Favorite maps

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Favorite maps* widget in `dashboard.create` and `dashboard.update` methods.

Parameters

The following parameters are supported for the *Favorite maps* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - *(default)* 15 minutes.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Favorite maps* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Favorite maps* widget

Configure a *Favorite maps* widget with the refresh interval set to 10 minutes.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "favmaps",
                        "name": "Favorite maps",
                        "x": 0,
                        "y": 0,
                        "width": 12,
                        "height": 3,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 0,
                                "name": "rf_rate",
                                "value": 600
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

6 Gauge

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Gauge* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Gauge* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Gauge* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Item	4	itemid.0	Item ID. Parameter behavior: - *required* if *Item (Widget)* is not set
Item (Widget)	1	itemid._reference	Instead of Item ID: `ABCDE._itemid` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for items. Parameter behavior: - *required* if *Item* is not set
Min	1	min	Any numeric value. Suffixes (e.g., "1d", "2w", "4K", "8G") are supported. Default: "0".
Max	1	max	Any numeric value. Suffixes (e.g., "1d", "2w", "4K", "8G") are supported. Default: "100".
Value arc	1	value_arc_color	Hexadecimal color code (e.g., `FF0000`). Default: "" (empty).

Parameter	type	name	value
Arc background	1	empty_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).
Background	1	bg_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).
Show	0	show.0	1 - Description; 2 - Value; 3 - Needle; 4 - Scale; 5 - Value arc. Note: To configure multiple values, create a dashboard widget field object for each value with an incremented number in the property name. Default: 1, 2, 4, 5. Values "Needle" and "Scale" are not supported if both: - no dashboard widget field object for *Show* with the value "Value arc" is set; - *Show arc* advanced configuration parameter is set to "Disabled". Advanced configuration parameters for *Show* options are not supported if no dashboard widget field objects with the respective values are set.
Override host	1	override_hostid._reference	ABCDE._hostid - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for hosts; DASHBOARD._hostid - set the dashboard host selector as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.

Advanced configuration

The following advanced configuration parameters are supported for the *Gauge* widget.

Note:
The number in the *Thresholds* property name (e.g., thresholds.0.color) references the threshold place in a list, sorted in ascending order. However, if thresholds are configured in a different order, the values will be sorted in ascending order after updating widget configuration in Zabbix frontend (e.g., "thresholds.0.threshold":"5" → "thresholds.0.threshold":"1"; "thresholds.1.threshold":"1" → "thresholds.1.threshold": "5").

Parameter	type	name	value
Angle	0	angle	Possible values: 180 *(default)* or 270.
Description			
Description	1	description	Any string value, including macros. Supported macros: {HOST.*}, {ITEM.*}, {INVENTORY.*}, user macros. Default: {ITEM.NAME}.
Size	0	desc_size	Possible values range from 1 to 100. Default: 15.
Vertical position	0	desc_v_pos	0 - Top; 1 - *(default)* Bottom.
Bold	0	desc_bold	0 - *(default)* Disabled; 1 - Enabled.

Parameter	type	name	value
Color	1	desc_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).
Value			
Decimal places	0	decimal_places	Possible values range from 1 to 10. Default: 2.
Size	0	value_size	Possible values range from 1 to 100. Default: 25.
Bold	0	value_bold	0 - *(default)* Disabled; 1 - Enabled.
Color	1	value_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).
Units			
Units (checkbox)	0	units_show	0 - Disabled; 1 - *(default)* Enabled.
Units (value)	1	units	Any string value. Parameter behavior: - *supported* if *Units* (checkbox) is set to ”Enabled”
Size	0	units_size	Possible values range from 1 to 100. Default: 25. Parameter behavior: - *supported* if *Units* (checkbox) is set to ”Enabled”
Bold	0	units_bold	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Units* (checkbox) is set to ”Enabled”
Position	0	units_pos	0 - Before value; 1 - Above value; 2 - *(default)* After value; 3 - Below value. Parameter behavior: - *supported* if *Units* (checkbox) is set to ”Enabled” This parameter is ignored if set to one of the following time-related units: unixtime, uptime, s.
Color	1	units_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).
Value arc			
Arc size	0	value_arc_size	Possible values range from 1 to 100. Default: 20.
Needle			
Color	1	needle_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty). Parameter behavior: - *supported* if a dashboard widget field object for *Show* with the value ”Value arc” is set, or *Show arc* is set to ”Enabled”
Scale			

Parameter	type	name	value
Show units	0	scale_show_units	0 - Disabled; 1 - *(default)* Enabled. Parameter behavior: - *supported* if *Units* (checkbox) is set to ”Enabled” and either a dashboard widget field object for *Show* with the value ”Value arc” is set, or *Show arc* is set to ”Enabled”
Size	0	scale_size	Possible values range from 1 to 100. Default: 15. Parameter behavior: - *supported* if a dashboard widget field object for *Show* with the value ”Value arc” is set, or *Show arc* is set to ”Enabled”
Decimal places	0	scale_decimal_places	Possible values range from 1 to 10. Default: 0. Parameter behavior: - *supported* if a dashboard widget field object for *Show* with the value ”Value arc” is set, or *Show arc* is set to ”Enabled”
Thresholds			
Color	1	thresholds.0.color	Hexadecimal color code (e.g., FF0000).
Threshold	1	thresholds.0.threshold	Any numeric value. Suffixes (e.g., ”1d”, ”2w”, ”4K”, ”8G”) are supported.
Show labels	0	th_show_labels	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Thresholds* are set and either a dashboard widget field object for *Show* with the value ”Value arc” is set or *Show arc* is set to ”Enabled”
Show arc	0	th_show_arc	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Thresholds* are set
Arc size	0	th_arc_size	Possible values range from 1 to 100. Default: 5. Parameter behavior: - *supported* if *Show arc* is set to ”Enabled”

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Gauge* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Gauge* widget

Configure a *Gauge* widget that displays the item value for the item ”44474” (Interface enp0s3: Bits sent). In addition, visually fine-tune the widget with multiple advanced options, including thresholds.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
```

```
        "pages": [
            {
                "widgets": [
                    {
                        "type": "gauge",
                        "name": "Gauge",
                        "x": 0,
                        "y": 0,
                        "width": 18,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 4,
                                "name": "itemid.0",
                                "value": 44474
                            },
                            {
                                "type": 1,
                                "name": "min",
                                "value": "100000"
                            },
                            {
                                "type": 1,
                                "name": "max",
                                "value": "1000000"
                            },
                            {
                                "type": 0,
                                "name": "show.0",
                                "value": 1
                            },
                            {
                                "type": 0,
                                "name": "show.1",
                                "value": 2
                            },
                            {
                                "type": 0,
                                "name": "show.2",
                                "value": 3
                            },
                            {
                                "type": 0,
                                "name": "show.4",
                                "value": 4
                            },
                            {
                                "type": 0,
                                "name": "show.5",
                                "value": 5
                            },
                            {
                                "type": 0,
                                "name": "angle",
                                "value": 270
                            },
                            {
                                "type": 0,
                                "name": "desc_size",
                                "value": 10
                            },
```

```
                {
                    "type": 0,
                    "name": "desc_bold",
                    "value": 1
                },
                {
                    "type": 0,
                    "name": "decimal_places",
                    "value": 0
                },
                {
                    "type": 0,
                    "name": "value_bold",
                    "value": 1
                },
                {
                    "type": 0,
                    "name": "units_size",
                    "value": 15
                },
                {
                    "type": 0,
                    "name": "units_pos",
                    "value": 3
                },
                {
                    "type": 1,
                    "name": "needle_color",
                    "value": "3C3C3C"
                },
                {
                    "type": 1,
                    "name": "thresholds.0.color",
                    "value": "FF465C"
                },
                {
                    "type": 1,
                    "name": "thresholds.0.threshold",
                    "value": "700000"
                },
                {
                    "type": 1,
                    "name": "thresholds.1.color",
                    "value": "FFD54F"
                },
                {
                    "type": 1,
                    "name": "thresholds.1.threshold",
                    "value": "500000"
                },
                {
                    "type": 1,
                    "name": "thresholds.2.color",
                    "value": "0EC9AC"
                },
                {
                    "type": 1,
                    "name": "thresholds.2.threshold",
                    "value": "100000"
                },
                {
                    "type": 0,
```

```
                                "name": "th_show_labels",
                                "value": 1
                            },
                            {
                                "type": 0,
                                "name": "th_show_arc",
                                "value": 1
                            },
                            {
                                "type": 0,
                                "name": "th_arc_size",
                                "value": 15
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

7 Geomap

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Geomap* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Geomap* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Geomap* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.
Hosts	3	hostids.0	Host ID. Note: To configure multiple hosts, create a dashboard widget field object for each host with an incremented number in the property name. For multiple hosts, the parameter *Host groups* must either be not configured at all or configured with at least one host group that the configured hosts belong to. This parameter is not supported if configuring the widget on a template dashboard.
Hosts (Widget/Dashboard)	1	hostids._reference	Instead of Host ID: `DASHBOARD.hostids` - set the dashboard host selector as the data source for hosts; `ABCDE._hostids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Tags			
Evaluation type	0	evaltype	0 - *(default)* And/Or; 2 - Or. This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Tag name	1	tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Tags* This parameter is not supported if configuring the widget on a template dashboard.
Operator	0	tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Tags* This parameter is not supported if configuring the widget on a template dashboard.
Tag value	1	tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Tags* This parameter is not supported if configuring the widget on a template dashboard.
Initial view	1	default_view	Comma separated *latitude*, *longitude*, *zoom level* (*optional*, possible values range from 0 to 30). Example: `40.6892494,-74.0466891,10`.
Clustering			
Clustering mode	0	clustering_mode	0 - *(default)* Auto; 1 - Zoom level.
Clustering zoom level	0	clustering_zoom_level	Possible values range from 0 to 30. Default: 0. Parameter behavior: - *supported* if *Clustering mode* is set to "Zoom level"
Reference	1	reference	Any string value consisting of 5 characters (e.g., `ABCDE` or `JBPNL`). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Geomap* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Geomap* widget

Configure a *Geomap* widget that displays hosts from host groups "2" and "22" based on the following tag configuration: tag with

the name ”component” contains value ”node”, and tag with the name ”location” equals value ”New York”. In addition, set the map initial view to coordinates ”40.6892494” (latitude), ”-74.0466891” (longitude) with the zoom level ”10”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "geomap",
                        "name": "Geomap",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 2,
                                "name": "groupids.0",
                                "value": 22
                            },
                            {
                                "type": 2,
                                "name": "groupids.1",
                                "value": 2
                            },
                            {
                                "type": 1,
                                "name": "default_view",
                                "value": "40.6892494,-74.0466891,10"
                            },
                            {
                                "type": 0,
                                "name": "evaltype",
                                "value": 2
                            },
                            {
                                "type": 1,
                                "name": "tags.0.tag",
                                "value": "component"
                            },
                            {
                                "type": 0,
                                "name": "tags.0.operator",
                                "value": 0
                            },
                            {
                                "type": 1,
                                "name": "tags.0.value",
                                "value": "node"
                            },
                            {
                                "type": 1,
                                "name": "tags.1.tag",
                                "value": "location"
                            },
```

```
                            {
                                "type": 0,
                                "name": "tags.1.operator",
                                "value": 1
                            },
                            {
                                "type": 1,
                                "name": "tags.1.value",
                                "value": "New York"
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

8 Graph

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Graph* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Graph* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Graph* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Reference	1	reference	Any string value consisting of 5 characters (e.g., `ABCDE` or `JBPNL`). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Data set

The following parameters are supported for configuring a *Data set*.

Note:
The first number in the property name (e.g., ds.0.hosts.0, ds.0.items.0) represents the particular data set, while the second number, if present, represents the configured host or item.

Parameter	type	name	value
Data set type	0	ds.0.dataset_type	0 - Item list; 1 - *(default)* Item patterns.
Items	4	ds.0.itemids.0	Item ID. When configuring the widget on a template dashboard, only items configured on the template should be set. Note: To configure multiple items, create a dashboard widget field object for each item with an incremented number in the property name. Parameter behavior: - *required* if *Data set type* is set to "Item list" and *Items (Widget)* is not set
Items (Widget)	1	ds.0.itemids.0._reference	Instead of Item ID: `ABCDE._itemid` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for items. Note: To configure multiple widgets, create a dashboard widget field object for each widget with an incremented number in the property name. Parameter behavior: - *required* if *Data set type* is set to "Item list" and *Items* is not set
Color	1	ds.0.color.0	Hexadecimal color code (e.g., `FF0000`). Parameter behavior: - *required* if *Data set type* is set to "Item list"
Host patterns	1	ds.0.hosts.0	Host name or pattern (e.g., "Zabbix*"). Parameter behavior: - *required* if *Data set type* is set to "Item patterns" This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Item patterns	1	ds.0.items.0	Item name or pattern (e.g., ”*: Number of processed *values per second”). When configuring the widget on a template dashboard, only the patterns for items configured on the template should be set. Parameter behavior: - *required* if *Data set type* is set to ”Item patterns”
Color	1	ds.0.color	Hexadecimal color code (e.g., FF0000). Parameter behavior: - *supported* if *Data set type* is set to ”Item patterns” and *Color palette* is not set
Color palette	0	ds.0.color_palette	Color palette index. Possible values range from 0 to 11. Default: 0. Parameter behavior: - *supported* if *Data set type* is set to ”Item patterns” and *Color* is not set
Draw	0	ds.0.type	0 - *(default)* Line; 1 - Points; 2 - Staircase; 3 - Bar.
Stacked	0	ds.0.stacked	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Draw* is set to ”Line”, ”Staircase”, or ”Bar”
Width	0	ds.0.width	Possible values range from 1 to 10. Default: 1. Parameter behavior: - *supported* if *Draw* is set to ”Line” or ”Staircase”
Point size	0	ds.0.pointsize	Possible values range from 1 to 10. Default: 3. Parameter behavior: - *supported* if *Draw* is set to ”Points”
Transparency	0	ds.0.transparency	Possible values range from 1 to 10. Default: 5.
Fill	0	ds.0.fill	Possible values range from 1 to 10. Default: 3. Parameter behavior: - *supported* if *Draw* is set to ”Line” or ”Staircase”
Missing data	0	ds.0.missingdatafunc	0 - *(default)* None; 1 - Connected; 2 - Treat as 0; 3 - Last known. Parameter behavior: - *supported* if *Draw* is set to ”Line” or ”Staircase”

Parameter	type	name	value
Override host	1	ds.0.override_hostid.ABCDE._hostid	ABCDE._hostid - set a compatible widget (with its "Reference" parameter set to "ABCDE") as the data source for hosts; DASHBOARD._hostid - set the dashboard host selector as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Y-axis	0	ds.0.axisy	0 - *(default)* Left; 1 - Right.
Invert values	0	ds.0.invert_values	0 - (default) Disabled; 1 - Enabled.
Time shift	1	ds.0.timeshift	Valid time string (e.g., 3600, 1h, etc.). You may use time suffixes. Negative values are also allowed. Default: "" (empty).
Aggregation function	0	ds.0.aggregate_function	0 - *(default)* not used; 1 - min; 2 - max; 3 - avg; 4 - count; 5 - sum; 6 - first; 7 - last.
Aggregation interval	1	ds.0.aggregate_interval	Valid time string (e.g., 3600, 1h, etc.). You may use time suffixes. Default: 1h.
Aggregate	0	ds.0.aggregate_grouping	0 - *(default)* Each item; 1 - Data set. Parameter behavior: - *supported* if *Aggregation function* is set to "min", "max", "avg", "count", "sum", "first", or "last"
Approximation	0	ds.0.approximation	1 - min; 2 - *(default)* avg; 4 - max; 7 - all.
Data set label	1	ds.0.data_set_label	Any string value. Default: "" (empty).
Item tags			
Evaluation type	0	ds.0.item_tags_evaltype	0 - *(default)* And/Or; 2 - Or.
Tag name	1	ds.0.item_tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*

Parameter	type	name	value
Operator	0	ds.0.item_tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*
Tag value	1	ds.0.item_tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*

Display options

The following parameters are supported for configuring *Display options*.

Parameter	type	name	value
History data selection	0	source	0 - *(default)* Auto; 1 - History; 2 - Trends.
Simple triggers	0	simple_triggers	0 - *(default)* Disabled; 1 - Enabled.
Working time	0	working_time	0 - *(default)* Disabled; 1 - Enabled.
Host names in labels	0	show_hostnames	0 - *(default)* Auto; 1 - Show; 2 - Hide. This parameter is not supported if configuring the widget on a template dashboard.
Percentile line (left)			
Status	0	percentile_left	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Y-axis* (in *Data set* configuration) is set to ”Left”
Value	0	percentile_left_value	Possible values range from 1 to 100. Parameter behavior: - *supported* if *Y-axis* (in *Data set* configuration) is set to ”Left”
Percentile line (right)			

Parameter	type	name	value
Status	0	percentile_right	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Y-axis* (in *Data set* configuration) is set to ”Right”
Value	0	percentile_right_value	Possible values range from 1 to 100. Parameter behavior: - *supported* if *Y-axis* (in *Data set* configuration) is set to ”Right”

Time period

The following parameters are supported for configuring *Time period*.

Parameter	type	name	value
Time period	1	time_period._reference	`DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its *Reference* parameter set to ”`ABCDE`”) as the data source. Default: `DASHBOARD._timeperiod` Alternatively, you can set the time period only in the *From* and *To* parameters.
From	1	time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.to` is set
To	1	time_period.to	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.from` is set

Axes

The following parameters are supported for configuring *Axes*.

Parameter	type	name	value
Left Y	0	lefty	0 - Disabled; 1 - *(default)* Enabled. Parameter behavior: - *supported* if *Y-axis* (in *Data set* configuration) is set to ”Left”
Right Y	0	righty	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Y-axis* (in *Data set* configuration) is set to ”Right”
Scale	0	lefty_scale righty_scale	0 - *(default)* Linear; 1 - Logarithmic.
Min	1	lefty_min righty_min	Any numeric value. Default: "" (empty).

Parameter	type	name	value
Max	1	lefty_max	Any numeric value. Default: "" (empty).
		righty_max	
Units (type)	0	lefty_units	0 - *(default)* Auto; 1 - Static.
		righty_units	
Units (value)	1	lefty_static_units	Any string value. Default: "" (empty).
		righty_static_units	
X-Axis	0	xaxis	0 - Disabled; 1 - *(default)* Enabled.

Legend

The following parameters are supported for configuring *Legend*.

Parameter	type	name	value
Show legend	0	legend	0 - Disabled; 1 - *(default)* Enabled.
Display min/avg/max	0	legend_statistic	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Show legend* is set to ”Enabled”
Show aggregation function	0	legend_aggregation	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Show legend* is set to ”Enabled”
Rows	0	legend_lines_mode	0 - *(default)* Fixed; 1 - Variable. Parameter behavior: - *supported* if *Show legend* is set to ”Enabled”
Number of rows/ Maximum number of rows	0	legend_lines	Possible values range from 1 to 10. Default: 1. Parameter behavior: - *supported* if *Show legend* is set to ”Enabled”
Number of columns	0	legend_columns	Possible values range from 1 to 4. Default: 4. Parameter behavior: - *supported* if *Show legend* is set to ”Enabled”, and *Display min/avg/max* is set to ”Disabled”

Problems

The following parameters are supported for configuring *Problems*.

Parameter	type	name	value
Show problems	0	show_problems	0 - *(default)* Disabled; 1 - Enabled.

Parameter	type	name	value
Selected items only	0	graph_item_problems	0 - Disabled; 1 - *(default)* Enabled.
Problem hosts	1	problemhosts.0	Host name. Note: The number in the property name references the configured host. To configure multiple hosts, create a dashboard widget field object for each host with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Severity	0	severities.0	0 - Not classified; 1 - Information; 2 - Warning; 3 - Average; 4 - High; 5 - Disaster. Default: empty (all enabled). Note: To configure multiple values, create a dashboard widget field object for each value with an incremented number in the property name.
Problem	1	problem_name	Problem event name (case insensitive, full name or part of it).
Problem tags			
Evaluation type	0	evaltype	0 - *(default)* And/Or; 2 - Or.
Tag name	1	tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Operator	0	tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Tag value	1	tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*

Overrides

The following parameters are supported for configuring *Overrides*.

Note:
The first number in the property name (e.g., or.0.hosts.0, or.0.items.0) represents the particular data set, while the second number, if present, represents the configured host or item.

Parameter	type	name	value
Host patterns	1	or.0.hosts.0	Host name or pattern (e.g., `Zabbix*`). This parameter is not supported if configuring the widget on a template dashboard. Parameter behavior: - *required* if configuring *Overrides*
Item patterns	1	or.0.items.0	Item name or pattern (e.g., `*: Number of processed *values per second`). When configuring the widget on a template dashboard, only the patterns for items configured on the template should be set. Parameter behavior: - *required* if configuring *Overrides*
Base color	1	or.0.color	Hexadecimal color code (e.g., `FF0000`).
Width	0	or.0.width	Possible values range from 1 to 10.
Draw	0	or.0.type	0 - Line; 1 - Points; 2 - Staircase; 3 - Bar.
Transparency	0	or.0.transparency	Possible values range from 1 to 10.
Fill	0	or.0.fill	Possible values range from 1 to 10.
Point size	0	or.0.pointsize	Possible values range from 1 to 10.
Missing data	0	or.0.missingdatafunc	0 - None; 1 - Connected; 2 - Treat as 0; 3 - Last known.
Y-axis	0	or.0.axisy	0 - Left; 1 - Right.
Invert values	0	or.0.invert_values	0 - (default) Disabled; 1 - Enabled.
Time shift	1	or.0.timeshift	Valid time string (e.g., 3600, 1h, etc.). You may use time suffixes. Negative values are allowed.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Graph* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Graph* widget

Configure a *Graph* widget in the following way:

- 2 data sets for a total of 9 items on 1 host.
- The first data set is of type "Item list" and consists of 3 items that are represented by lines with a different color, but the same width, transparency, and fill.
- The second data set is of type "Item patterns", consists of 6 items, has a configured aggregation, and is represented by a line with a custom color, width, transparency, and fill.
- The second data set also has a custom data set label.
- Data in the graph are displayed for a time period of the last 3 hours.
- Problems in the graph are displayed only for the configured items.
- Graph has two Y axes of which the right Y axis displays values only for the second data set.
- Graph legend displays configured items in 4 rows, as well as minimum, maximum and average values of the data sets.

Request:

```
{
    "jsonrpc": "2.0",
```

```
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "svggraph",
                        "name": "Graph",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 0,
                                "name": "ds.0.dataset_type",
                                "value": 0
                            },
                            {
                                "type": 4,
                                "name": "ds.0.itemids.0",
                                "value": 23264
                            },
                            {
                                "type": 1,
                                "name": "ds.0.color.0",
                                "value": "FF0000"
                            },
                            {
                                "type": 4,
                                "name": "ds.0.itemids.1",
                                "value": 23269
                            },
                            {
                                "type": 1,
                                "name": "ds.0.color.1",
                                "value": "BF00FF"
                            },
                            {
                                "type": 4,
                                "name": "ds.0.itemids.2",
                                "value": 23257
                            },
                            {
                                "type": 1,
                                "name": "ds.0.color.2",
                                "value": "0040FF"
                            },
                            {
                                "type": 0,
                                "name": "ds.0.width",
                                "value": 3
                            },
                            {
                                "type": 0,
                                "name": "ds.0.transparency",
                                "value": 3
                            },
```

```
                    {
                        "type": 0,
                        "name": "ds.0.fill",
                        "value": 1
                    },
                    {
                        "type": 1,
                        "name": "ds.1.hosts.0",
                        "value": "Zabbix server"
                    },
                    {
                        "type": 1,
                        "name": "ds.1.items.0",
                        "value": "*: Number of processed *values per second"
                    },
                    {
                        "type": 1,
                        "name": "ds.1.color",
                        "value": "000000"
                    },
                    {
                        "type": 0,
                        "name": "ds.1.transparency",
                        "value": 0
                    },
                    {
                        "type": 0,
                        "name": "ds.1.fill",
                        "value": 0
                    },
                    {
                        "type": 0,
                        "name": "ds.1.axisy",
                        "value": 1
                    },
                    {
                        "type": 0,
                        "name": "ds.1.aggregate_function",
                        "value": 3
                    },
                    {
                        "type": 1,
                        "name": "ds.1.aggregate_interval",
                        "value": "1m"
                    },
                    {
                        "type": 0,
                        "name": "ds.1.aggregate_grouping",
                        "value": 1
                    },
                    {
                        "type": 1,
                        "name": "ds.1.data_set_label",
                        "value": "Number of processed values per second"
                    },
                    {
                        "type": 0,
                        "name": "graph_time",
                        "value": 1
                    },
                    {
                        "type": 1,
```

```
                                "name": "time_period.from",
                                "value": "now-3h"
                            },
                            {
                                "type": 1,
                                "name": "time_period.to",
                                "value": "now"
                            },
                            {
                                "type": 0,
                                "name": "legend_statistic",
                                "value": 1
                            },
                            {
                                "type": 0,
                                "name": "legend_lines",
                                "value": 4
                            },
                            {
                                "type": 0,
                                "name": "show_problems",
                                "value": 1
                            },
                            {
                                "type": 1,
                                "name": "reference",
                                "value": "YZABC"
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field

- `dashboard.create`
- `dashboard.update`

9 Graph (classic)

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Graph (classic)* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Graph (classic)* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Graph (classic)* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Source	0	source_type	0 - *(default)* Graph; 1 - Simple graph.
Graph	6	graphid.0	Graph ID. Parameter behavior: - *required* if *Source* is set to ”Graph”
Graph (Widget)	1	graphid._reference	Instead of Graph ID: `ABCDE._graphid` - set a compatible widget (with its *Reference* parameter set to ”ABCDE”) as the data source for graphs.
Item	4	itemid.0	Parameter behavior: - *required* if *Source* is set to ”Simple graph” and *Graph* is not set Item ID. Parameter behavior: - *required* if *Source* is set to ”Simple graph” and *Item (Widget)* is not set
Item (Widget)	1	itemid._reference	Instead of Item ID: `ABCDE._itemid` - set a compatible widget (with its *Reference* parameter set to ”ABCDE”) as the data source for items.
Time period	1	time_period._reference	Parameter behavior: - *required* if *Source* is set to ”Simple graph” and *Item* is not set `DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its *Reference* parameter set to ”ABCDE”) as the data source. Default: `DASHBOARD._timeperiod` Alternatively, you can set the time period only in the *From* and *To* parameters.

Parameter	type	name	value
From	1	time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.to` is set
To	1	time_period.to	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.from` is set
Show legend	0	show_legend	0 - Disabled; 1 - *(default)* Enabled.
Override host	1	override_hostid._reference	`ABCDE._hostid` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for hosts; `DASHBOARD._hostid` - set the dashboard host selector as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Reference	1	reference	Any string value consisting of 5 characters (e.g., `ABCDE` or `JBPNL`). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Graph (classic)* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Graph (classic)* widget

Configure a *Graph (classic)* widget that displays a simple graph for the item "42269".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "graph",
                        "name": "Graph (classic)",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 0,
                                "name": "source_type",
                                "value": 1
```

```
                            },
                            {
                                "type": 4,
                                "name": "itemid.0",
                                "value": 42269
                            },
                            {
                                "type": 1,
                                "name": "reference",
                                "value": "RSTUV"
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

10 Graph prototype

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Graph prototype* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Graph prototype* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Graph prototype* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Source	0	source_type	2 - *(default)* Graph prototype; 3 - Simple graph prototype.
Graph prototype	7	graphid.0	Graph prototype ID. Parameter behavior: - *required* if *Source* is set to ”Graph prototype”
Item prototype	5	itemid.0	Item prototype ID. Parameter behavior: - *required* if *Source* is set to ”Simple graph prototype”
Time period	1	time_period._reference	`DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its *Reference* parameter set to ”ABCDE”) as the data source. Default: `DASHBOARD._timeperiod` Alternatively, you can set the time period only in the *From* and *To* parameters.
From	1	time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.to` is set
To	1	time_period.to	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.from` is set
Show legend	0	show_legend	0 - Disabled; 1 - *(default)* Enabled.
Override host	1	override_hostid._reference	`ABCDE._hostid` - set a compatible widget (with its *Reference* parameter set to ”ABCDE”) as the data source for hosts; `DASHBOARD._hostid` - set the dashboard host selector as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Columns	0	columns	Possible values range from 1 to 24. Default: 2.
Rows	0	rows	Possible values range from 1 to 16. Default: 1.

Parameter	type	name	value
Reference	1	reference	Any string value consisting of 5 characters (e.g., ABCDE or JBPNL). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Graph prototype* widget. For more information on configuring a dashboard, see dashboard.create.

Configuring a *Graph prototype* widget

Configure a *Graph prototype* widget that displays a grid of 3 graphs (3 columns, 1 row) created from an item prototype (ID: "42316") by low-level discovery.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "graphprototype",
                        "name": "Graph prototype",
                        "x": 0,
                        "y": 0,
                        "width": 48,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 0,
                                "name": "source_type",
                                "value": 3
                            },
                            {
                                "type": 5,
                                "name": "itemid.0",
                                "value": 42316
                            },
                            {
                                "type": 0,
                                "name": "columns",
                                "value": 3
                            },
                            {
                                "type": 1,
                                "name": "reference",
                                "value": "OPQWX"
                            }
                        ]
                    }
                ]
            }
        ],
```

```
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

11 Honeycomb

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Honeycomb* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Honeycomb* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Honeycomb* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.

Parameter	type	name	value
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.
Hosts	3	hostids.0	Host ID. Note: To configure multiple hosts, create a dashboard widget field object for each host with an incremented number in the property name. For multiple hosts, the parameter *Host groups* must either be not configured at all or configured with at least one host group that the configured hosts belong to. This parameter is not supported if configuring the widget on a template dashboard.
Hosts (Widget/Dashboard)	1	hostids._reference	Instead of Host ID: `DASHBOARD.hostids` - set the dashboard host selector as the data source for hosts; `ABCDE._hostids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Host tags			
Evaluation type	0	evaltype_host	0 - *(default)* And/Or; 2 - Or. This parameter is not supported if configuring the widget on a template dashboard.
Tag name	1	host_tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Operator	0	host_tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Tag value	1	host_tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Item patterns	1	items.0	Item name or pattern. Note: To configure multiple item patterns, create a dashboard widget field object for each item pattern with an incremented number in the property name. Parameter behavior: - *required*
Item tags			
Evaluation type	0	evaltype_item	0 - *(default)* And/Or; 2 - Or.
Tag name	1	item_tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*
Operator	0	item_tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*

Parameter	type	name	value
Tag value	1	item_tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*
Show hosts in maintenance	0	maintenance	0 - *(default)* Disabled; 1 - Enabled.
Show	0	show.0	1 - Primary label; 2 - Secondary label. Note: To configure multiple values, create a dashboard widget field object for each value with an incremented number in the property name. Default: 1, 2.
Reference	1	reference	Any string value consisting of 5 characters (e.g., `ABCDE` or `JBPNL`). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Advanced configuration

The following advanced configuration parameters are supported for the *Honeycomb* widget.

Note:
The number in the *Thresholds* property name (e.g., thresholds.0.color) references the threshold place in a list, sorted in ascending order. However, if thresholds are configured in a different order, the values will be sorted in ascending order after updating widget configuration in Zabbix frontend (e.g., `"thresholds.0.threshold":"5"` → `"thresholds.0.threshold":"1"`; `"thresholds.1.threshold":"1"` → `"thresholds.1.threshold": "5"`).

Parameter	type	name	value
Primary label			
Type	0	primary_label_type	0 - *(default)* Text; 1 - Value.
Text	1	primary_label	Any string value, including macros. Supported macros: `{HOST.*}`, `{ITEM.*}`, `{INVENTORY.*}`, user macros. Default: `{HOST.NAME}` Parameter behavior: - *supported* if *Type* is set to "Text"
Decimal places	0	primary_label_decimal_places	Possible values range from 0 to 6. Default: 2. Parameter behavior: - *supported* if *Type* is set to "Value"
Size (type)	0	primary_label_size_type	0 - *(default)* Auto; 1 - Custom.

Parameter	type	name	value
Size	0	primary_label_size	Possible values range from 1 to 100. Default: 20. Parameter behavior: - *supported* if *Size* (type) is set to ”Custom”
Bold	0	primary_label_bold	0 - *(default)* Disabled; 1 - Enabled.
Color	1	primary_label_color	Hexadecimal color code (e.g., FF0000). Default: based on theme of Settings object and User object: 1F2C33 for ”blue-theme”, ”blue-classic-theme”, or ”hc-light”; EEEEEE for ”dark-theme”, ”dark-classic-theme”, or ”hc-dark”; E5F1FF for ”dark-blue-theme”.
Units (checkbox)	0	primary_label_units_show	0 - Disabled; 1 - *(default)* Enabled. Parameter behavior: - *supported* if *Type* is set to ”Value”
Units (value)	1	primary_label_units	Any string value. "" (empty) Parameter behavior: - *supported* if *Type* is set to ”Value” and *Units* (checkbox) is set to ”Enabled”
Position	0	primary_label_units_pos	0 - Before value; 1 - *(default)* After value. Parameter behavior: - *supported* if *Type* is set to ”Value” and *Units* (checkbox) is set to ”Enabled” This parameter is ignored if set to one of the following time-related units: unixtime, uptime, s.
Secondary label			
Type	0	secondary_label_type	0 - Text; 1 - *(default)* Value.
Text	1	secondary_label	Any string value, including macros. Supported macros: {HOST.*}, {ITEM.*}, {INVENTORY.*}, user macros. Default: {{ITEM.LASTVALUE}.fmtnum(2)} Parameter behavior: - *supported* if *Type* is set to ”Text”
Decimal places	0	secondary_label_decimal_places	Possible values range from 0 to 6. Default: 2. Parameter behavior: - *supported* if *Type* is set to ”Value”
Size (type)	0	secondary_label_size_type	0 - *(default)* Auto; 1 - Custom.
Size	0	secondary_label_size	Possible values range from 1 to 100. Default: 30. Parameter behavior: - *supported* if *Size* (type) is set to ”Custom”
Bold	0	secondary_label_bold	0 - Disabled; 1 - *(default)* Enabled.

Parameter	type	name	value
Color	1	secondary_label_color	Hexadecimal color code (e.g., FF0000). Default: based on theme of Settings object and User object: 1F2C33 for ”blue-theme”, ”blue-classic-theme”, or ”hc-light”; EEEEEE for ”dark-theme”, ”dark-classic-theme”, or ”hc-dark”; E5F1FF for ”dark-blue-theme”.
Units (checkbox)	0	secondary_label_units_show	0 - Disabled; 1 - *(default)* Enabled. Parameter behavior: - *supported* if *Type* is set to ”Value”
Units (value)	1	secondary_label_units	Any string value. "" (empty) Parameter behavior: - *supported* if *Type* is set to ”Value” and *Units* (checkbox) is set to ”Enabled”
Position	0	secondary_label_position	0 - Before value; 1 - *(default)* After value. Parameter behavior: - *supported* if *Type* is set to ”Value” and *Units* (checkbox) is set to ”Enabled” This parameter is ignored if set to one of the following time-related units: unixtime, uptime, s.
Background color			
Background color	1	bg_color	Hexadecimal color code (e.g., FF0000). Default: based on theme of Settings object and User object: D9E7ED for ”blue-theme” and ”blue-classic-theme”; 3D5059 for ”dark-theme” and ”dark-classic-theme”; 004FA8 for ”dark-blue-theme”; AAD7E9 for ”hc-light”; 335463 for ”hc-dark”.
Thresholds			
Color interpolation	0	interpolation	0 - Disabled; 1 - *(default)* Enabled.
Color	1	thresholds.0.color	Hexadecimal color code (e.g., FF0000).
Threshold	1	thresholds.0.threshold	Any numeric value. Suffixes (e.g., ”1d”, ”2w”, ”4K”, ”8G”) are supported.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Honeycomb* widget. For more information on configuring a dashboard, see dashboard.create.

Configuring a *Honeycomb* widget

Configure a *Honeycomb* widget that displays the utilization of Zabbix server processes. In addition, change the primary label of honeycomb cells and visually fine-tune the widget with thresholds.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": "30",
        "auto_start": "1",
```

```
        "pages": [
            {
                "widgets": [
                    {
                        "type": "honeycomb",
                        "name": "Honeycomb",
                        "x": "0",
                        "y": "0",
                        "width": "24",
                        "height": "5",
                        "view_mode": "0",
                        "fields": [
                            {
                                "type": 2,
                                "name": "groupids.0",
                                "value": 4
                            },
                            {
                                "type": 3,
                                "name": "hostids.0",
                                "value": 10084
                            },
                            {
                                "type": 1,
                                "name": "items.0",
                                "value": "Zabbix server: Utilization*"
                            },
                            {
                                "type": 1,
                                "name": "primary_label",
                                "value": "{ITEM.NAME}"
                            },
                            {
                                "type": 1,
                                "name": "thresholds.0.color",
                                "value": "0EC9AC"
                            },
                            {
                                "type": 1,
                                "name": "thresholds.0.threshold",
                                "value": "0"
                            },
                            {
                                "type": 1,
                                "name": "thresholds.1.color",
                                "value": "FFD54F"
                            },
                            {
                                "type": 1,
                                "name": "thresholds.1.threshold",
                                "value": "70"
                            },
                            {
                                "type": 1,
                                "name": "thresholds.2.color",
                                "value": "FF465C"
                            },
                            {
                                "type": 1,
                                "name": "thresholds.2.threshold",
                                "value": "90"
                            },
```

```
                            {
                                "type": 1,
                                "name": "reference",
                                "value": "KSTMQ"
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

12 Host availability

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Host availability* widget in `dashboard.create` and `dashboard.update` methods.

Parameters

The following parameters are supported for the *Host availability* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - *(default)* 15 minutes.

Parameter	type	name	value
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.
Interface type	0	interface_type.0	0 - None; 1 - Zabbix agent (passive checks); 2 - SNMP; 3 - IPMI; 4 - JMX; 5 - Zabbix agent (active checks). Default: 1, 2, 3, 4, 5 (all enabled). Note: To configure multiple values, create a dashboard widget field object for each value with an incremented number in the property name.
Layout	0	layout	0 - *(default)* Horizontal; 1 - Vertical.
Show hosts in maintenance	0	maintenance	0 - *(default)* Disabled; 1 - Enabled.
Show only totals	0	only_totals	0 - *(default)* Disabled; 1 - Enabled.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Host availability* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Host availability* widget

Configure a *Host availability* widget that displays availability information (in a vertical layout) for hosts in host group "4" with "Zabbix agent" and "SNMP" interfaces configured.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
```

```
                        "type": "hostavail",
                        "name": "Host availability",
                        "x": 0,
                        "y": 0,
                        "width": 18,
                        "height": 3,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 2,
                                "name": "groupids.0",
                                "value": 4
                            },
                            {
                                "type": 0,
                                "name": "interface_type",
                                "value": 1
                            },
                            {
                                "type": 0,
                                "name": "interface_type",
                                "value": 2
                            },
                            {
                                "type": 0,
                                "name": "layout",
                                "value": 1
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field

- `dashboard.create`
- `dashboard.update`

13 Host card

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Host card* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Host card* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Host card* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Host	3	hostid.0	Host ID. Parameter behavior: - *required* if *Host (Widget/Dashboard)* is not set This parameter is not supported if configuring the widget on a template dashboard.
Host (Widget/Dashboard)	1	hostid._reference	Instead of Host ID: `DASHBOARD.hostid` - set the dashboard host selector as the data source for host; `ABCDE._hostid` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for host. Parameter behavior: - *required* if *Host* is not set This parameter is not supported if configuring the widget on a template dashboard.
Show suppressed problems	0	show_suppressed	0 - *(default)* Disabled; 1 - Enabled.

Parameter	type	name	value
Show	0	sections.0	0 - Host groups; 1 - Description; 2 - Monitoring; 3 - Availability; 4 - Monitored by; 5 - Templates; 6 - Inventory; 7 - Tags. Note: The number in the property name references section order in the section list. To configure multiple sections, create a dashboard widget field object for each section with an incremented number in the property name.

The following parameters are supported if *Show* is set to ”Inventory”.

Parameter	type	name	value
Inventory fields	0	inventory.0	Inventory ID. Note: To configure multiple inventory fields, create a dashboard widget field object for each inventory field with an incremented number in the property name.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Host card* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Host card* widget

Configure a *Host card* widget that displays these sections: ”Monitoring”, ”Availability”, ”Monitored by”, ”Inventory”, and ”Tags”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "hostcard",
                        "name": "Host card",
                        "x": 0,
                        "y": 0,
                        "width": 14,
                        "height": 7,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 3,
                                "name": "hostid.0",
                                "value": 10084
                            },
                            {
                                "type": 0,
                                "name": "show_suppressed",
                                "value": 1
```

```
                            },
                            {
                                "type": 0,
                                "name": "sections.0",
                                "value": 2
                            },
                            {
                                "type": 0,
                                "name": "sections.1",
                                "value": 3
                            },
                            {
                                "type": 0,
                                "name": "sections.2",
                                "value": 4
                            },
                            {
                                "type": 0,
                                "name": "sections.3",
                                "value": 6
                            },
                            {
                                "type": 0,
                                "name": "sections.4",
                                "value": 7
                            },
                            {
                                "type": 0,
                                "name": "inventory.0",
                                "value": 25
                            },
                            {
                                "type": 0,
                                "name": "inventory.1",
                                "value": 26
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
```

```
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

14 Host navigator

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Host navigator* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Host navigator* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Host navigator* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Host patterns	1	hosts.0	Host name or pattern. Note: To configure multiple host patterns, create a dashboard widget field object for each host pattern with an incremented number in the property name. For multiple host patterns, the parameter *Host groups* must either be not configured at all or configured with at least one host group that the hosts matching the configured host patterns belong to. This parameter is not supported if configuring the widget on a template dashboard.
Host status	0	status	-1 - *(default)* Any; 0 - Enabled; 1 - Disabled. This parameter is not supported if configuring the widget on a template dashboard.
Host tags			
Evaluation type	0	host_tags_evaltype	0 - *(default)* And/Or; 2 - Or. This parameter is not supported if configuring the widget on a template dashboard.
Tag name	1	host_tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Operator	0	host_tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Tag value	1	host_tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Severity	0	severities.0	0 - Not classified; 1 - Information; 2 - Warning; 3 - Average; 4 - High; 5 - Disaster. Default: empty (all enabled). Note: To configure multiple values, create a dashboard widget field object for each value with an incremented number in the property name.
Show hosts in maintenance	0	maintenance	0 - *(default)* Disabled; 1 - Enabled.
Show problems	0	show_problems	0 - All; 1 - *(default)* Unsuppressed; 2 - None.
Group by			
Attribute	0	group_by.0.attribute	0 - Host group; 1 - Tag value; 2 - Severity. Note: The number in the property name references attribute order in the grouping attribute list. Parameter behavior: - *required* if configuring *Group by*
Value	1	group_by.0.tag_name	Any string value. Note: The number in the property name references the grouping attribute set in the *Attribute* parameter. Parameter behavior: - *required* if configuring *Group by* and *Attribute* is set to ”Tag value”
Host limit	0	show_lines	Possible values range from 1 to 9999. Default: 100. This parameter is not supported if configuring the widget on a template dashboard.
Reference	1	reference	Any string value consisting of 5 characters (e.g., `ABCDE` or `JBPNL`). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Host navigator* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Host navigator* widget

Configure a *Host navigator* widget that displays hosts grouped by their host group and, then, by the ”city” tag value.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": "30",
        "auto_start": "1",
        "pages": [
            {
                "widgets": [
                    {
                        "type": "hostnavigator",
                        "name": "Host navigator",
                        "x": "0",
                        "y": "0",
                        "width": "12",
                        "height": "5",
                        "view_mode": "0",
                        "fields": [
                            {
                                "type": 2,
                                "name": "groupids.0",
                                "value": 2
                            },
                            {
                                "type": 2,
                                "name": "groupids.1",
                                "value": 4
                            },
                            {
                                "type": 0,
                                "name": "group_by.0.attribute",
                                "value": 0
                            },
                            {
                                "type": 0,
                                "name": "group_by.1.attribute",
                                "value": 1
                            },
                            {
                                "type": 1,
                                "name": "group_by.1.tag_name",
                                "value": "city"
                            },
                            {
                                "type": 1,
                                "name": "reference",
                                "value": "SWKLB"
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
```

```
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

15 Item card

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Item card* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Item card* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Item card* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Item	4	itemid.0	Item ID. Parameter behavior: - *required* if *Item (Widget)* is not set
Item (Widget)	1	itemid._reference	Instead of Item ID: `ABCDE._itemid` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for item. Parameter behavior: - *required* if *Item* is not set

Parameter	type	name	value
Show	0	sections.0	0 - Description; 1 - Error text; 2 - Metrics; 3 - Latest data; 4 - Type of information; 5 - Triggers; 6 - Host interface; 7 - Type; 8 - Host inventory; 9 - Tags. Note: The number in the property name references section order in the section list. To configure multiple sections, create a dashboard widget field object for each section with an incremented number in the property name.
Override host	1	override_hostid._reference	`ABCDE._hostid` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for hosts; `DASHBOARD._hostid` - set the dashboard host selector as the data source for hosts. Default: `""` (empty) This parameter is not supported if configuring the widget on a template dashboard.

Sparkline

The following parameters are supported if *Show* is set to "Latest data".

Parameter	type	name	value
Width	0	sparkline.width	Possible values range from 0 to 10. Default: 1.
Fill	0	sparkline.fill	Possible values range from 0 to 10. Default: 3.
Color	1	sparkline.color	Hexadecimal color code (e.g., `FF0000`). Default: `42A5F5`.
Time period	1	sparkline.time_period.reference	`DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its `reference` parameter equal to `ABCDE`) as the data source. Default: `""` (empty) Alternatively, you can set the time period only in the *From* and *To* parameters.
From	1	sparkline.time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Default: now-1h. Parameter behavior: - *supported* if *Time period* is not set - *required* if `sparkline.time_period.to` is set

Parameter	type	name	value
To	1	sparkline.time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Default: now. Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.to` is set
History data	0	sparkline.history	0 - *(default)* Auto; 1 - History; 2 - Trends.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Item card* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring an *Item card* widget

Configure an *Item card* widget that displays these sections: ”Description”, ”Latest data”, ”Triggers”, and ”Tags”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "itemcard",
                        "name": "Item card",
                        "x": 0,
                        "y": 0,
                        "width": 14,
                        "height": 7,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 4,
                                "name": "itemid.0",
                                "value": 42257
                            },
                            {
                                "type": 0,
                                "name": "sections.0",
                                "value": 0
                            },
                            {
                                "type": 0,
                                "name": "sections.1",
                                "value": 3
                            },
                            {
                                "type": 0,
                                "name": "sections.2",
                                "value": 5
                            },
                            {
```

```
                                "type": 0,
                                "name": "sections.3",
                                "value": 9
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

16 Item history

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Item history* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Item history* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Item history* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Layout	0	layout	0 - *(default)* Horizontal; 1 - Vertical.
Columns (see below)			
Show lines	0	show_lines	Possible values range from 1 to 100. Default: 25.
Override host	1	override_hostid._reference	`ABCDE._hostid` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for hosts; `DASHBOARD._hostid` - set the dashboard host selector as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Advanced configuration (see below)			
Reference	1	reference	Any string value consisting of 5 characters (e.g., `ABCDE` or `JBPNL`). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Columns

Columns have common parameters and additional parameters depending on the configuration of the *Item* parameter.

Note:
For all parameters related to columns, the number in the property name (e.g., columns.0.name) references a column for which the parameter is configured.

The following parameters are supported for all columns.

Parameter	type	name	value
Name	1	columns.0.name	Any string value. Parameter behavior: - *required*
Item	4	columns.0.itemid	Item ID. When configuring the widget on a template dashboard, only items configured on the template should be set. Parameter behavior: - *required*
Base color	1	columns.0.base_color	Hexadecimal color code (e.g., `FF0000`). Default: "" (empty).

The following column parameters are supported if the configured *Item* is a numeric type item.

Parameter	type	name	value
Display	0	columns.0.display	1 - *(default)* As is; 2 - Bar; 3 - Indicators.
Min	1	columns.0.min	Any numeric value. Parameter behavior: - *supported* if *Display* is set to ”Bar” or ”Indicators”
Max	1	columns.0.max	Any numeric value. Parameter behavior: - *supported* if *Display* is set to ”Bar” or ”Indicators”
Thresholds			
Color	1	columns.0.thresholds.color	Hexadecimal color code (e.g., FF0000).
Threshold	1	columns.0.thresholds.threshold	Any numeric value. Suffixes (e.g., ”1d”, ”2w”, ”4K”, ”8G”) are supported.
History data	0	columns.0.history	0 - *(default)* Auto; 1 - History; 2 - Trends.

The following column parameters are supported if the configured *Item* is a character, text, or log type item.

Parameter	type	name	value
Highlights			
Highlight	1	columns.0.highlights.color	Hexadecimal color code (e.g., FF0000).
Threshold	1	columns.0.highlights.pattern	Any regular expression.
Display	0	columns.0.display	1 - *(default)* As is; 4 - HTML; 5 - Single line.
Single line	0	columns.0.max_length	Possible values range from 1 to 500. Default: 100. Parameter behavior: - *supported* if *Display* is set to ”Single line”
Use monospace font	0	columns.0.monospace_font	0 - *(default)* Use default font; 1 - Use monospace font.
Display local time	0	columns.0.local_time	0 - *(default)* Display timestamp; 1 - Display local time. Parameter behavior: - *supported* if *Item* is set to log type item, and *Show timestamp* is set to ”Enabled”

The following column parameters are supported if the configured *Item* is a binary type item.

Parameter	type	name	value
Show thumbnail	1	columns.0.show_thumbnail	0 - *(default)* Disabled; 1 - Enabled.

Advanced configuration

The following advanced configuration parameters are supported for the *Item history* widget.

Parameter	type	name	value
New values	0	sortorder	0 - *(default)* Top; 1 - Bottom.
Show timestamp	0	show_timestamp	0 - *(default)* Disabled; 1 - Enabled.
Show column header	0	show_column_header	0 - Off; 1 - Horizontal; 2 - *(default)* Vertical.
Time period	1	time_period._reference	`DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source. Default: `DASHBOARD._timeperiod` Alternatively, you can set the time period only in the *From* and *To* parameters.
From	1	time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.to` is set
To	1	time_period.to	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.from` is set

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Item history* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring an *Item history* widget

Configure an *Item history* widget that displays latest data for two numeric items "42269" and "42270". In addition, configure the item columns to be displayed vertically, with column names displayed horizontally; limit the display to 15 lines of data and include a separate timestamp column.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "itemhistory",
                        "name": "Item history",
                        "x": "0",
                        "y": "0",
                        "width": "18",
                        "height": "6",
                        "view_mode": "0",
```

```
                    "fields": [
                        {
                            "type": "0",
                            "name": "layout",
                            "value": "1"
                        },
                        {
                            "type": "1",
                            "name": "columns.0.name",
                            "value": "CPU utilization"
                        },
                        {
                            "type": "4",
                            "name": "columns.0.itemid",
                            "value": "42269"
                        },
                        {
                            "type": "1",
                            "name": "columns.1.name",
                            "value": "Memory utilization"
                        },
                        {
                            "type": "4",
                            "name": "columns.1.itemid",
                            "value": "42270"
                        },
                        {
                            "type": "0",
                            "name": "show_lines",
                            "value": "15"
                        },
                        {
                            "type": "0",
                            "name": "show_timestamp",
                            "value": "1"
                        },
                        {
                            "type": "0",
                            "name": "show_column_header",
                            "value": "1"
                        },
                        {
                            "type": "1",
                            "name": "reference",
                            "value": "KIVKD"
                        }
                    ]
                }
            ]
        }
    ],
    "userGroups": [
        {
            "usrgrpid": 7,
            "permission": 2
        }
    ],
    "users": [
        {
            "userid": 1,
            "permission": 3
        }
```

```
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

17 Item navigator

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Item navigator* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Item navigator* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Item navigator* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Hosts	3	hostids.0	Host ID. Note: To configure multiple hosts, create a dashboard widget field object for each host with an incremented number in the property name. For multiple hosts, the parameter *Host groups* must either be not configured at all or configured with at least one host group that the configured hosts belong to. This parameter is not supported if configuring the widget on a template dashboard.
Hosts (Widget/Dashboard)	1	hostids._reference	Instead of Host ID: `DASHBOARD.hostid` - set the dashboard host selector as the data source for hosts; `ABCDE._hostid` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Host tags			
Evaluation type	0	host_tags_evaltype	0 - *(default)* And/Or; 2 - Or. This parameter is not supported if configuring the widget on a template dashboard.
Tag name	1	host_tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Operator	0	host_tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Tag value	1	host_tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Item patterns	1	items.0	Item name or pattern. Note: To configure multiple item patterns, create a dashboard widget field object for each item pattern with an incremented number in the property name.
Item tags			
Evaluation type	0	item_tags_evaltype	0 - *(default)* And/Or; 2 - Or.
Tag name	1	item_tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*
Operator	0	item_tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*
Tag value	1	item_tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*
State	0	state	-1 - *(default)* All; 0 - Normal; 1 - Not supported.
Show problems	0	show_problems	0 - All; 1 - *(default)* Unsuppressed; 2 - None.
Group by			
Attribute	0	group_by.0.attribute	0 - Host group; 1 - Host name; 2 - Host tag value; 3 - Item tag value. Note: The number in the property name references attribute order in the grouping attribute list. Parameter behavior: - *required* if configuring *Group by*
Value	1	group_by.0.tag_name	Any string value. Note: The number in the property name references the grouping attribute set in the *Attribute* parameter. Parameter behavior: - *required* if configuring *Group by* and *Attribute* is set to ”Host tag value” or ”Item tag value”

Parameter	type	name	value
Item limit	0	show_lines	Possible values range from 1 to 9999. Default: 100.
Reference	1	reference	Any string value consisting of 5 characters (e.g., ABCDE or JBPNL). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Item navigator* widget. For more information on configuring a dashboard, see dashboard.create.

Configuring an *Item navigator* widget

Configure an *Item navigator* widget that displays up to 1000 items grouped by their host and, then, by the ”component” item tag value.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": "30",
        "auto_start": "1",
        "pages": [
            {
                "widgets": [
                    {
                        "type": "itemnavigator",
                        "name": "Item navigator",
                        "x": "0",
                        "y": "0",
                        "width": "12",
                        "height": "5",
                        "view_mode": "0",
                        "fields": [
                            {
                                "type": 0,
                                "name": "group_by.0.attribute",
                                "value": 0
                            },
                            {
                                "type": 0,
                                "name": "group_by.1.attribute",
                                "value": 3
                            },
                            {
                                "type": 1,
                                "name": "group_by.1.tag_name",
                                "value": "component"
                            },
                            {
                                "type": 0,
                                "name": "show_lines",
                                "value": 1000
                            },
                            {
                                "type": 1,
```

```
                                "name": "reference",
                                "value": "DFNLK"
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

18 Item value

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Item value* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Item value* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Item value* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Item	4	itemid.0	Item ID. Parameter behavior: - *required* if *Item (Widget)* is not set
Item (Widget)	1	itemid._reference	Instead of Item ID: `ABCDE._itemid` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for items. Parameter behavior: - *required* if *Item* is not set
Show	0	show.0	1 - Description; 2 - Value; 3 - Time; 4 - Change indicator; 5 - Sparkline. Default: 1, 2, 3, 4. Note: To configure multiple values, create a dashboard widget field object for each value with an incremented number in the property name.
Override host	1	override_hostid._reference	`ABCDE._hostid` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for hosts; `DASHBOARD._hostid` - set the dashboard host selector as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.

Advanced configuration

The following advanced configuration parameters are supported for the *Item value* widget.

Note:
The number in the *Thresholds* property name (e.g., thresholds.0.color) references the threshold place in a list, sorted in ascending order. However, if thresholds are configured in a different order, the values will be sorted in ascending order after updating widget configuration in Zabbix frontend (e.g., `"thresholds.0.threshold":"5"` → `"thresholds.0.threshold":"1"`; `"thresholds.1.threshold":"1"` → `"thresholds.1.threshold": "5"`).

Parameter	type	name	value
Background color	1	bg_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).
Thresholds			
Color	1	thresholds.0.color	Hexadecimal color code (e.g., FF0000).
Threshold	1	thresholds.0.threshold	Any string value.

Parameter	type	name	value
Aggregation function	0	aggregate_function	0 - *(default)* not used; 1 - min; 2 - max; 3 - avg; 4 - count; 5 - sum; 6 - first; 7 - last.
Time period	1	time_period._reference	`DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its `reference` parameter equal to `ABCDE`) as the data source. Default: `DASHBOARD._timeperiod` Alternatively, you can set the time period only in the *From* and *To* parameters. Parameter behavior: - *supported* if *Aggregation function* is set to ”min”, ”max”, ”avg”, ”count”, ”sum”, ”first”, ”last”
From	1	time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set and *Aggregation function* is set to ”min”, ”max”, ”avg”, ”count”, ”sum”, ”first”, ”last” - *required* if `time_period.to` is set
To	1	time_period.to	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set and *Aggregation function* is set to ”min”, ”max”, ”avg”, ”count”, ”sum”, ”first”, ”last” - *required* if `time_period.from` is set
History data	0	history	0 - *(default)* Auto; 1 - History; 2 - Trends.

Description

The following advanced configuration parameters are supported if *Show* is set to ”Description”.

Parameter	type	name	value
Description	1	description	Any string value, including macros. Supported macros: `{HOST.*}`, `{ITEM.*}`, `{INVENTORY.*}`, User macros. Default: `{ITEM.NAME}`.
Horizontal position	0	desc_h_pos	0 - Left; 1 - *(default)* Center; 2 - Right. Two or more elements (Description, Value, Time) cannot share the same *Horizontal position* and *Vertical position*.

Parameter	type	name	value
Vertical position	0	desc_v_pos	0 - Top; 1 - Middle; 2 - *(default)* Bottom. Two or more elements (Description, Value, Time) cannot share the same *Horizontal position* and *Vertical position*.
Size	0	desc_size	Possible values range from 1 to 100. Default: 15.
Bold	0	desc_bold	0 - *(default)* Disabled; 1 - Enabled.
Color	1	desc_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).

Value

The following advanced configuration parameters are supported if *Show* is set to ”Value”.

Parameter	type	name	value
Decimal places			
Decimal places	0	decimal_places	Possible values range from 1 to 10. Default: 2.
Size	0	decimal_size	Possible values range from 1 to 100. Default: 35.
Position			
Horizontal position	0	value_h_pos	0 - Left; 1 - *(default)* Center; 2 - Right. Two or more elements (Description, Value, Time) cannot share the same *Horizontal position* and *Vertical position*.
Vertical position	0	value_v_pos	0 - Top; 1 - *(default)* Middle; 2 - Bottom. Two or more elements (Description, Value, Time) cannot share the same *Horizontal position* and *Vertical position*.
Size	0	value_size	Possible values range from 1 to 100. Default: 45.
Bold	0	value_bold	0 - Disabled; 1 - *(default)* Enabled.
Color	1	value_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).
Units			
Units (checkbox)	0	units_show	0 - Disabled; 1 - *(default)* Enabled.
Units (value)	1	units	Any string value.
Position	0	units_pos	0 - Before value; 1 - Above value; 2 - *(default)* After value; 3 - Below value.
Size	0	units_size	Possible values range from 1 to 100. Default: 35.

Parameter	type	name	value
Bold	0	units_bold	0 - Disabled; 1 - *(default)* Enabled.
Color	1	units_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).

Time

The following advanced configuration parameters are supported if *Show* is set to ”Time”.

Parameter	type	name	value
Horizontal position	0	time_h_pos	0 - Left; 1 - *(default)* Center; 2 - Right. Two or more elements (Description, Value, Time) cannot share the same *Horizontal position* and *Vertical position*.
Vertical position	0	time_v_pos	0 - *(default)* Top; 1 - Middle; 2 - Bottom. Two or more elements (Description, Value, Time) cannot share the same *Horizontal position* and *Vertical position*.
Size	0	time_size	Possible values range from 1 to 100. Default: 15.
Bold	0	time_bold	0 - *(default)* Disabled; 1 - Enabled.
Color	1	time_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).

Change indicator

The following advanced configuration parameters are supported if *Show* is set to ”Change indicator”.

Parameter	type	name	value
Change indicator ↑ color	1	up_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).
Change indicator ↓ color	1	down_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).
Change indicator ↕ color	1	updown_color	Hexadecimal color code (e.g., FF0000). Default: "" (empty).

Sparkline

The following advanced configuration parameters are supported if *Show* is set to ”Sparkline”.

Parameter	type	name	value
Width	0	sparkline.width	Possible values range from 0 to 10. Default: 1.
Fill	0	sparkline.fill	Possible values range from 0 to 10. Default: 3.

Parameter	type	name	value
Color	1	sparkline.color	Hexadecimal color code (e.g., FF0000). Default: 42A5F5.
Time period	1	sparkline.time_period.from	DASHBOARD._timeperiod - set the dashboard time period selector as the data source; ABCDE._timeperiod - set a compatible widget (with its reference parameter equal to ABCDE) as the data source. Default: DASHBOARD._timeperiod Alternatively, you can set the time period only in the *From* and *To* parameters.
From	1	sparkline.time_period.from	Valid time string in absolute (YYYY-MM-DD hh:mm:ss) or relative time syntax (now, now/d, now/w-1w, etc.). Parameter behavior: - *required* if sparkline.time_period.to is set
To	1	sparkline.time_period.to	Valid time string in absolute (YYYY-MM-DD hh:mm:ss) or relative time syntax (now, now/d, now/w-1w, etc.). Parameter behavior: - *required* if sparkline.time_period.from is set
History data	0	sparkline.history	0 - *(default)* Auto; 1 - History; 2 - Trends.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Item value* widget. For more information on configuring a dashboard, see dashboard.create.

Configuring an *Item value* widget

Configure an *Item value* widget that displays the item value for the item "42266" (Zabbix agent availability). In addition, visually fine-tune the widget with multiple advanced options, including a dynamic background color that changes based on the availability status of Zabbix agent.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "item",
                        "name": "Item value",
                        "x": 0,
                        "y": 0,
                        "width": 12,
                        "height": 3,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 4,
                                "name": "itemid.0",
                                "value": 42266
                            },
```

```
                {
                    "type": 0,
                    "name": "show.0",
                    "value": 1
                },
                {
                    "type": 0,
                    "name": "show.1",
                    "value": 2
                },
                {
                    "type": 0,
                    "name": "show.2",
                    "value": 3
                },
                {
                    "type": 1,
                    "name": "description",
                    "value": "Agent status"
                },
                {
                    "type": 0,
                    "name": "desc_h_pos",
                    "value": 0
                },
                {
                    "type": 0,
                    "name": "desc_v_pos",
                    "value": 0
                },
                {
                    "type": 0,
                    "name": "desc_bold",
                    "value": 1
                },
                {
                    "type": 1,
                    "name": "desc_color",
                    "value": "F06291"
                },
                {
                    "type": 0,
                    "name": "value_h_pos",
                    "value": 0
                },
                {
                    "type": 0,
                    "name": "value_size",
                    "value": 25
                },
                {
                    "type": 1,
                    "name": "value_color",
                    "value": "FFFF00"
                },
                {
                    "type": 0,
                    "name": "units_show",
                    "value": 0
                },
                {
                    "type": 0,
```

```
                            "name": "time_h_pos",
                            "value": 2
                        },
                        {
                            "type": 0,
                            "name": "time_v_pos",
                            "value": 2
                        },
                        {
                            "type": 0,
                            "name": "time_size",
                            "value": 10
                        },
                        {
                            "type": 0,
                            "name": "time_bold",
                            "value": 1
                        },
                        {
                            "type": 1,
                            "name": "time_color",
                            "value": "9FA8DA"
                        },
                        {
                            "type": 1,
                            "name": "thresholds.0.color",
                            "value": "E1E1E1"
                        },
                        {
                            "type": 1,
                            "name": "thresholds.0.threshold",
                            "value": "0"
                        },
                        {
                            "type": 1,
                            "name": "thresholds.1.color",
                            "value": "D1C4E9"
                        },
                        {
                            "type": 1,
                            "name": "thresholds.1.threshold",
                            "value": "1"
                        }
                    ]
                }
            ]
        }
    ],
    "userGroups": [
        {
            "usrgrpid": 7,
            "permission": 2
        }
    ],
    "users": [
        {
            "userid": 1,
            "permission": 3
        }
    ]
},
"id": 1
```

```
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

19 Map

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Map* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Map* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Map* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - *(default)* 15 minutes.
Map	8	sysmapid.0	Map ID. Parameter behavior: - *required* if *Map (Widget)* is not set
Map (Widget)	1	sysmapid._reference	`ABCDE._mapid` - set a *Map navigation tree* widget (with its *Reference* parameter set to "ABCDE") as the data source for maps. Parameter behavior: - *required* if *Map* is not set
Reference	1	reference	Any string value consisting of 5 characters (e.g., `ABCDE` or `JBPNL`). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Map* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Map* widget

Configure a *Map* widget that displays the map "1".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "map",
                        "name": "Map",
                        "x": 0,
                        "y": 0,
                        "width": 54,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 8,
                                "name": "sysmapid.0",
                                "value": 1
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

Configuring a linked *Map* widget

Configure a *Map* widget that is linked to a *Map navigation tree* widget.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "map",
                        "name": "Map",
                        "x": 0,
                        "y": 5,
                        "width": 54,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 1,
                                "name": "sysmapid._reference",
                                "value": "ABCDE._mapid"
                            }
                        ]
                    },
                    {
                        "type": "navtree",
                        "name": "Map navigation tree",
                        "x": 0,
                        "y": 0,
                        "width": 18,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 1,
                                "name": "navtree.1.name",
                                "value": "Element A"
                            },
                            {
                                "type": 1,
                                "name": "navtree.2.name",
                                "value": "Element B"
                            },
                            {
                                "type": 1,
                                "name": "navtree.3.name",
                                "value": "Element C"
                            },
                            {
                                "type": 1,
                                "name": "navtree.4.name",
                                "value": "Element A1"
                            },
                            {
                                "type": 1,
                                "name": "navtree.5.name",
```

```
            "value": "Element A2"
        },
        {
            "type": 1,
            "name": "navtree.6.name",
            "value": "Element B1"
        },
        {
            "type": 1,
            "name": "navtree.7.name",
            "value": "Element B2"
        },
        {
            "type": 0,
            "name": "navtree.4.parent",
            "value": 1
        },
        {
            "type": 0,
            "name": "navtree.5.parent",
            "value": 1
        },
        {
            "type": 0,
            "name": "navtree.6.parent",
            "value": 2
        },
        {
            "type": 0,
            "name": "navtree.7.parent",
            "value": 2
        },
        {
            "type": 0,
            "name": "navtree.1.order",
            "value": 1
        },
        {
            "type": 0,
            "name": "navtree.2.order",
            "value": 2
        },
        {
            "type": 0,
            "name": "navtree.3.order",
            "value": 3
        },
        {
            "type": 0,
            "name": "navtree.4.order",
            "value": 1
        },
        {
            "type": 0,
            "name": "navtree.5.order",
            "value": 2
        },
        {
            "type": 0,
            "name": "navtree.6.order",
            "value": 1
        },
```

```
                            {
                                "type": 0,
                                "name": "navtree.7.order",
                                "value": 2
                            },
                            {
                                "type": 8,
                                "name": "navtree.6.sysmapid",
                                "value": 1
                            },
                            {
                                "type": 1,
                                "name": "reference",
                                "value": "ABCDE"
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`
- Map navigation tree

20 Map navigation tree

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Map navigation tree* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Map navigation tree* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Map navigation tree* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - *(default)* 15 minutes.
Show unavailable maps	1	show_unavailable	0 - *(default)* Disabled; 1 - Enabled.
Reference	1	reference	Any string value consisting of 5 characters (e.g., `ABCDE` or `JBPNL`). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

The following parameters are supported for configuring map navigation tree elements.

Parameter	type	name	value
Name	1	navtree.1.name	Any string value. Note: The number in the property name sets the element number.
Linked map	8	navtree.1.sysmapid	Map ID. Note: The number in the property name references the element to which the map is linked.
Parameters for creating element hierarchy	0	navtree.1.parent	Parent element number. Note: The number in the property name references the child element. The property value references the parent element.
	0	navtree.1.order	Element position in the map navigation tree. Note: The number in the property name references the element number. The property value references the element position in the map navigation tree. Parent element position is determined within the whole map navigation tree. Child element position is determined within the parent element.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Map navigation tree* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Map navigation tree* widget

Configure a *Map navigation tree* widget that displays the following map navigation tree:

- Element A
  - Element A1
  - Element A2
- Element B
  - Element B1 (contains linked map "1" that can be displayed in a linked *Map widget*)

        - Element B2
    - Element C

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "navtree",
                        "name": "Map navigation tree",
                        "x": 0,
                        "y": 0,
                        "width": 18,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 1,
                                "name": "navtree.1.name",
                                "value": "Element A"
                            },
                            {
                                "type": 1,
                                "name": "navtree.2.name",
                                "value": "Element B"
                            },
                            {
                                "type": 1,
                                "name": "navtree.3.name",
                                "value": "Element C"
                            },
                            {
                                "type": 1,
                                "name": "navtree.4.name",
                                "value": "Element A1"
                            },
                            {
                                "type": 1,
                                "name": "navtree.5.name",
                                "value": "Element A2"
                            },
                            {
                                "type": 1,
                                "name": "navtree.6.name",
                                "value": "Element B1"
                            },
                            {
                                "type": 1,
                                "name": "navtree.7.name",
                                "value": "Element B2"
                            },
                            {
                                "type": 0,
                                "name": "navtree.4.parent",
                                "value": 1
                            },
```

```
                {
                    "type": 0,
                    "name": "navtree.5.parent",
                    "value": 1
                },
                {
                    "type": 0,
                    "name": "navtree.6.parent",
                    "value": 2
                },
                {
                    "type": 0,
                    "name": "navtree.7.parent",
                    "value": 2
                },
                {
                    "type": 0,
                    "name": "navtree.1.order",
                    "value": 1
                },
                {
                    "type": 0,
                    "name": "navtree.2.order",
                    "value": 2
                },
                {
                    "type": 0,
                    "name": "navtree.3.order",
                    "value": 3
                },
                {
                    "type": 0,
                    "name": "navtree.4.order",
                    "value": 1
                },
                {
                    "type": 0,
                    "name": "navtree.5.order",
                    "value": 2
                },
                {
                    "type": 0,
                    "name": "navtree.6.order",
                    "value": 1
                },
                {
                    "type": 0,
                    "name": "navtree.7.order",
                    "value": 2
                },
                {
                    "type": 8,
                    "name": "navtree.6.sysmapid",
                    "value": 1
                },
                {
                    "type": 1,
                    "name": "reference",
                    "value": "HJQXF"
                }
            ]
        }
```

```
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`
- Map

21 Pie chart

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Pie chart* widget in `dashboard.create` and `dashboard.update` methods.

> **Attention:**
> Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Pie chart* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Pie chart* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.

Data set

The following parameters are supported for configuring a *Data set*.

Note:
The first number in the property name (e.g., ds.0.hosts.0, ds.0.items.0) represents the particular data set, while the second number, if present, represents the configured host or item.

Parameter	type	name	value
Data set type	0	ds.0.dataset_type	0 - Item list; 1 - *(default)* Item patterns.
Items	4	ds.0.itemids.0	Item ID. When configuring the widget on a template dashboard, only items configured on the template should be set. Note: To configure multiple items, create a dashboard widget field object for each item with an incremented number in the property name. Parameter behavior: - *required* if *Data set type* is set to ”Item list” and *Items (Widget)* is not set
Items (Widget)	1	ds.0.itemids.0._reference	Instead of Item ID: `ABCDE._itemid` - set a compatible widget (with its *Reference* parameter set to ”ABCDE”) as the data source for items. Note: To configure multiple widgets, create a dashboard widget field object for each widget with an incremented number in the property name. Parameter behavior: - *required* if *Data set type* is set to ”Item list” and *Items* is not set
Color	1	ds.0.color.0	Hexadecimal color code (e.g., `FF0000`). Parameter behavior: - *required* if *Data set type* is set to ”Item list”
Item type	0	ds.0.type.0	0 - *(default)* Normal; 1 - Total. The value ”Total” can be set only for one item in the whole chart. Parameter behavior: - *supported* if *Data set type* is set to ”Item list”
Host patterns	1	ds.0.hosts.0	Host name or pattern (e.g., ”Zabbix*”). Parameter behavior: - *required* if *Data set type* is set to ”Item patterns” This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Item patterns	1	ds.0.items.0	Item name or pattern (e.g., ”*: Number of processed *values per second”). When configuring the widget on a template dashboard, only the patterns for items configured on the template should be set. Parameter behavior: - *required* if *Data set type* is set to ”Item patterns”
Color	1	ds.0.color	Hexadecimal color code (e.g., FF0000). Parameter behavior: - *supported* if *Data set type* is set to ”Item patterns” and *Color palette* is not set
Color palette	0	ds.0.color_palette	Color palette index. Possible values range from 0 to 11. Default: 0. Parameter behavior: - *supported* if *Data set type* is set to ”Item patterns” and *Color* is not set
Aggregation function	0	ds.0.aggregate_function	1 - min; 2 - max; 3 - avg; 4 - count; 5 - sum; 6 - first; 7 - *(default)* last.
Data set aggregation	0	ds.0.dataset_aggregation	0 - *(default)* none; 1 - min; 2 - max; 3 - avg; 4 - count; 5 - sum. Parameter behavior: - *supported* if *Item type* is set to ”Total”
Data set label	1	ds.0.data_set_label	Any string value. Default: "" (empty).
Item tags			
Evaluation type	0	ds.0.item_tags_evaltype	0 - *(default)* And/Or; 2 - Or.
Tag name	1	ds.0.item_tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*

Parameter	type	name	value
Operator	0	ds.0.item_tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*
Tag value	1	ds.0.item_tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*

Displaying options

The following parameters are supported for configuring *Displaying options*.

Parameter	type	name	value
History data selection	0	source	0 - *(default)* Auto; 1 - History; 2 - Trends.
Draw	0	draw_type	0 - *(default)* Pie; 1 - Doughnut.
Width	0	width	20 - 20% of the radius; 30 - 30% of the radius; 40 - 40% of the radius; 50 - *(default)* 50% of the radius. Parameter behavior: - *supported* if *Draw* is set to ”Doughnut”
Stroke width	0	stroke	Possible values range from 0 to 10. Default: 0. Parameter behavior: - *supported* if *Draw* is set to ”Doughnut”
Show total value	0	total_show	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Draw* is set to ”Doughnut”
Size	0	value_size_type	0 - *(default)* Auto; 1 - Custom. Parameter behavior: - *supported* if *Show total value* is set to ”Enabled”

Parameter	type	name	value
Size (value for custom size)	0	value_size	Possible values range from 1 to 100. Default: 20. Parameter behavior: - *supported* if *Show total value* is set to ”Enabled”
Decimal places	0	decimal_places	Possible values range from 0 to 6. Default: 2. Parameter behavior: - *supported* if *Show total value* is set to ”Enabled”
Units (checkbox)	0	units_show	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Show total value* is set to ”Enabled”
Units (value)	1	units	Any string value. Parameter behavior: - *supported* if *Units* (checkbox) is set to ”Enabled”
Bold	0	value_bold	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Show total value* is set to ”Enabled”
Color	1	value_color	Hexadecimal color code (e.g., FF0000). Parameter behavior: - *supported* if *Show total value* is set to ”Enabled”
Space between sectors	0	space	Possible values range from 0 to 10. Default: 1.
Merge sectors smaller than N% (checkbox)	0	merge	0 - *(default)* Disabled; 1 - Enabled.
Merge sectors smaller than N% (value)	0	merge_percent	Possible values range from 1 to 10. Default: 1. Parameter behavior: - *supported* if *Merge sectors smaller than N%* (checkbox) is set to ”Enabled”
Merge sectors smaller than N% (color)	1	merge_color	Hexadecimal color code (e.g., FF0000). Parameter behavior: - *supported* if *Merge sectors smaller than N%* (checkbox) is set to ”Enabled”

Time period

The following parameters are supported for configuring *Time period*.

Parameter	type	name	value
Time period	1	time_period._reference	`DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source. Default: `DASHBOARD._timeperiod` Alternatively, you can set the time period only in the *From* and *To* parameters.
From	1	time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.to` is set
To	1	time_period.to	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.from` is set

Legend

The following parameters are supported for configuring *Legend*.

Parameter	type	name	value
Show legend	0	legend	0 - Disabled; 1 - *(default)* Enabled.
Show value	0	legend_value	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Show legend* is set to "Enabled"
Show aggregation function	0	legend_aggregation	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Show legend* is set to "Enabled"
Rows	0	legend_lines_mode	0 - *(default)* Fixed; 1 - Variable. Parameter behavior: - *supported* if *Show legend* is set to "Enabled"
Number of rows/ Maximum number of rows	0	legend_lines	Possible values range from 1 to 10. Default: 1. Parameter behavior: - *supported* if *Show legend* is set to "Enabled"
Number of columns	0	legend_columns	Possible values range from 1 to 4. Default: 4. Parameter behavior: - *supported* if *Show legend* is set to "Enabled", and *Show value* is set to "Disabled"

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Pie chart* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Pie chart* widget

Configure a *Pie chart* widget in the following way:

- 2 data sets for a total of 9 items on 1 host.
- The first data set is of type "Item list" and consists of 3 items that are all with type "Normal" and represented by a different color.
- The second data set is of type "Item patterns", consists of 6 items, has a configured aggregation for each item, and is represented by a custom color.
- The second data set also has a custom data set label.
- Data in the pie chart are displayed as a doughnut chart with a custom width and total value with units in the center.
- Data in the pie chart are displayed and aggregated for a custom time period of the last 3 hours.
- Pie chart legend displays configured items in 4 rows.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "piechart",
                        "name": "Pie chart",
                        "x": 0,
                        "y": 0,
                        "width": 24,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 0,
                                "name": "ds.0.dataset_type",
                                "value": 0
                            },
                            {
                                "type": 4,
                                "name": "ds.0.itemids.0",
                                "value": 23264
                            },
                            {
                                "type": 1,
                                "name": "ds.0.color.0",
                                "value": "FF0000"
                            },
                            {
                                "type": 0,
                                "name": "ds.0.type.0",
                                "value": 0
                            },
                            {
                                "type": 4,
                                "name": "ds.0.itemids.1",
                                "value": 23269
                            },
                            {
                                "type": 1,
```

```
            "name": "ds.0.color.1",
            "value": "BF00FF"
        },
        {
          "type": 0,
          "name": "ds.0.type.1",
          "value": 0
        },
        {
            "type": 4,
            "name": "ds.0.itemids.2",
            "value": 23257
        },
        {
            "type": 1,
            "name": "ds.0.color.2",
            "value": "0040FF"
        },
        {
          "type": 0,
          "name": "ds.0.type.2",
          "value": 0
        },
        {
            "type": 1,
            "name": "ds.1.hosts.0",
            "value": "Zabbix server"
        },
        {
            "type": 1,
            "name": "ds.1.items.0",
            "value": "*: Number of processed *values per second"
        },
        {
            "type": 1,
            "name": "ds.1.color",
            "value": "000000"
        },
        {
            "type": 0,
            "name": "ds.1.aggregate_function",
            "value": 3
        },
        {
            "type": 1,
            "name": "ds.1.data_set_label",
            "value": "Number of processed values per second"
        },
        {
            "type": 0,
            "name": "draw_type",
            "value": 1
        },
        {
            "type": 0,
            "name": "width",
            "value": 30
        },
        {
            "type": 0,
            "name": "total_show",
            "value": 1
```

```
                            },
                            {
                                "type": 0,
                                "name": "units_show",
                                "value": 1
                            },
                            {
                                "type": 0,
                                "name": "graph_time",
                                "value": 1
                            },
                            {
                                "type": 1,
                                "name": "time_period.from",
                                "value": "now-3h"
                            },
                            {
                                "type": 1,
                                "name": "time_period.to",
                                "value": "now"
                            },
                            {
                                "type": 0,
                                "name": "legend_lines",
                                "value": 4
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- dashboard.create
- dashboard.update

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Problem hosts* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Problem hosts* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Problem hosts* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.
Exclude host groups	2	exclude_groupids.0	Host group ID. Note: To exclude multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Hosts	3	hostids.0	Host ID. Note: To configure multiple hosts, create a dashboard widget field object for each host with an incremented number in the property name. For multiple hosts, the parameter *Host groups* must either be not configured at all or configured with at least one host group that the configured hosts belong to. This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Hosts (Widget/Dashboard)	1	hostids._reference	Instead of Host ID: `DASHBOARD.hostids` - set the dashboard host selector as the data source for hosts; `ABCDE._hostids` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Problem	1	problem	Problem event name (case insensitive, full name or part of it).
Severity	0	severities.0	0 - Not classified; 1 - Information; 2 - Warning; 3 - Average; 4 - High; 5 - Disaster. Default: empty (all enabled). Note: To configure multiple values, create a dashboard widget field object for each value with an incremented number in the property name.
Problem tags			
Evaluation type	0	evaltype	0 - *(default)* And/Or; 2 - Or.
Tag name	1	tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Operator	0	tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Tag value	1	tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Show suppressed problems	0	show_suppressed	0 - *(default)* Disabled; 1 - Enabled.
Hide groups without problems	0	hide_empty_groups	0 - *(default)* Disabled; 1 - Enabled. This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Problem display	0	ext_ack	0 - *(default)* All; 1 - Unacknowledged only; 2 - Separated.
Reference	1	reference	Any string value consisting of 5 characters (e.g., ABCDE or JBPNL). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Problem hosts* widget. For more information on configuring a dashboard, see dashboard.create.

Configuring a *Problem hosts* widget

Configure a *Problem hosts* widget that displays hosts from host groups ”2” and ”4” that have problems with a name that includes the string ”CPU” and that have the following severities: ”Warning”, ”Average”, ”High”, ”Disaster”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "problemhosts",
                        "name": "Problem hosts",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 2,
                                "name": "groupids.0",
                                "value": 2
                            },
                            {
                                "type": 2,
                                "name": "groupids.1",
                                "value": 4
                            },
                            {
                                "type": 1,
                                "name": "problem",
                                "value": "cpu"
                            },
                            {
                                "type": 0,
                                "name": "severities.0",
                                "value": 2
                            },
                            {
                                "type": 0,
```

```
                                "name": "severities.1",
                                "value": 3
                            },
                            {
                                "type": 0,
                                "name": "severities.2",
                                "value": 4
                            },
                            {
                                "type": 0,
                                "name": "severities.3",
                                "value": 5
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

23 Problems

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Problems* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Problems* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Problems* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Show	0	show	1 - *(default)* Recent problems; 2 - History; 3 - Problems.
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.
Exclude host groups	2	exclude_groupids.0	Host group ID. Note: To exclude multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Hosts	3	hostids.0	Host ID. Note: To configure multiple hosts, create a dashboard widget field object for each host with an incremented number in the property name. For multiple hosts, the parameter *Host groups* must either be not configured at all or configured with at least one host group that the configured hosts belong to. This parameter is not supported if configuring the widget on a template dashboard.
Hosts (Widget/Dashboard)	1	hostids._reference	Instead of Host ID: `DASHBOARD.hostids` - set the dashboard host selector as the data source for hosts; `ABCDE._hostids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Problem	1	problem	Problem event name (case insensitive, full name or part of it).
Severity	0	severities.0	0 - Not classified; 1 - Information; 2 - Warning; 3 - Average; 4 - High; 5 - Disaster. Default: empty (all enabled). Note: To configure multiple values, create a dashboard widget field object for each value with an incremented number in the property name.
Problem tags			
Evaluation type	0	evaltype	0 - *(default)* And/Or; 2 - Or.
Tag name	1	tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Operator	0	tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Tag value	1	tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Show tags	0	show_tags	0 - *(default)* None; 1 - 1; 2 - 2; 3 - 3.
Tag name (format)	0	tag_name_format	0 - *(default)* Full; 1 - Shortened; 2 - None. Parameter behavior: - *supported* if *Show tags* is set to ”1”, ”2”, or ”3”
Tag display priority	1	tag_priority	Comma-separated list of tags. Parameter behavior: - *supported* if *Show tags* is set to ”1”, ”2”, or ”3”

Parameter	type	name	value
Show operational data	0	show_opdata	0 - *(default)* None; 1 - Separately; 2 - With problem name.
Show suppressed problems	0	show_suppressed	0 - *(default)* Disabled; 1 - Enabled.
Acknowledgement status	0	acknowledgement_status	0 - *(default)* all; 1 - Unacknowledged; 2 - Acknowledged.
By me	0	acknowledged_by_me	0 - *(default)* Disabled; 1 - Enabled.
Sort entries by	0	sort_triggers	1 - Severity (descending); 2 - Host (ascending); 3 - Time (ascending); 4 - *(default)* Time (descending); 13 - Severity (ascending); 14 - Host (descending); 15 - Problem (ascending); 16 - Problem (descending). For all values, except ”Time (descending)” and ”Time (ascending)”, the *Show timeline* parameter must be set to ”Disabled”. Values ”Host (ascending)” and ”Host (descending)” are not supported if configuring the widget on a template dashboard.
Show timeline	0	show_timeline	0 - Disabled; 1 - *(default)* Enabled. Parameter behavior: - *supported* if *Sort entries by* is set to ”Time (descending)” or ”Time (ascending)”
Highlight whole row	0	highlight_row	0 - *(default)* Disabled; 1 - Enabled. This parameter is not supported in the high-contrast themes.
Show lines	0	show_lines	Possible values range from 1 to 100. Default: 25.
Reference	1	reference	Any string value consisting of 5 characters (e.g., `ABCDE` or `JBPNL`). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Problems* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Problems* widget

Configure a *Problems* widget that displays problems for host group ”4” that satisfy the following conditions:

- Problems that have a tag with the name ”scope” that contains values ”performance” or ”availability”, or ”capacity”.
- Problems that have the following severities: ”Warning”, ”Average”, ”High”, ”Disaster”.

In addition, configure the widget to show tags and operational data.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "problems",
                        "name": "Problems",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 2,
                                "name": "groupids.0",
                                "value": 4
                            },
                            {
                                "type": 1,
                                "name": "tags.0.tag",
                                "value": "scope"
                            },
                            {
                                "type": 0,
                                "name": "tags.0.operator",
                                "value": 0
                            },
                            {
                                "type": 1,
                                "name": "tags.0.value",
                                "value": "performance"
                            },
                            {
                                "type": 1,
                                "name": "tags.1.tag",
                                "value": "scope"
                            },
                            {
                                "type": 0,
                                "name": "tags.1.operator",
                                "value": 0
                            },
                            {
                                "type": 1,
                                "name": "tags.1.value",
                                "value": "availability"
                            },
                            {
                                "type": 1,
                                "name": "tags.2.tag",
                                "value": "scope"
                            },
                            {
                                "type": 0,
```

```
                                "name": "tags.2.operator",
                                "value": 0
                            },
                            {
                                "type": 1,
                                "name": "tags.2.value",
                                "value": "capacity"
                            },
                            {
                                "type": 0,
                                "name": "severities.0",
                                "value": 2
                            },
                            {
                                "type": 0,
                                "name": "severities.1",
                                "value": 3
                            },
                            {
                                "type": 0,
                                "name": "severities.2",
                                "value": 4
                            },
                            {
                                "type": 0,
                                "name": "severities.3",
                                "value": 5
                            },
                            {
                                "type": 0,
                                "name": "show_tags",
                                "value": 1
                            },
                            {
                                "type": 0,
                                "name": "show_opdata",
                                "value": 1
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
```

```
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

24 Problems by severity

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Problems by severity* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Problems by severity* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Problems by severity* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.
Exclude host groups	2	exclude_groupids.0	Host group ID. Note: To exclude multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Hosts	3	hostids.0	Host ID. Note: To configure multiple hosts, create a dashboard widget field object for each host with an incremented number in the property name. For multiple hosts, the parameter *Host groups* must either be not configured at all or configured with at least one host group that the configured hosts belong to. This parameter is not supported if configuring the widget on a template dashboard.
Hosts (Widget/Dashboard)	1	hostids._reference	Instead of Host ID: `DASHBOARD.hostids` - set the dashboard host selector as the data source for hosts; `ABCDE._hostids` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Problem	1	problem	Problem event name (case insensitive, full name or part of it).
Severity	0	severities.0	0 - Not classified; 1 - Information; 2 - Warning; 3 - Average; 4 - High; 5 - Disaster. Default: empty (all enabled). Note: To configure multiple values, create a dashboard widget field object for each value with an incremented number in the property name.
Problem tags			
Evaluation type	0	evaltype	0 - *(default)* And/Or; 2 - Or.
Tag name	1	tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Operator	0	tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Tag value	1	tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*

Parameter	type	name	value
Show	0	show_type	0 - *(default)* Host groups; 1 - Totals. This parameter is not supported if configuring the widget on a template dashboard, and, by default, is set to ”Totals”.
Layout	0	layout	0 - *(default)* Horizontal; 1 - Vertical. Parameter behavior: - *supported* if *Show* is set to ”Totals”
Show operational data	0	show_opdata	0 - *(default)* None; 1 - Separately; 2 - With problem name.
Show suppressed problems	0	show_suppressed	0 - *(default)* Disabled; 1 - Enabled.
Hide groups without problems	0	hide_empty_groups	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Show* is set to ”Host groups” This parameter is not supported if configuring the widget on a template dashboard.
Problem display	0	ext_ack	0 - *(default)* All; 1 - Unacknowledged only; 2 - Separated.
Show timeline	0	show_timeline	0 - Disabled; 1 - *(default)* Enabled.
Reference	1	reference	Any string value consisting of 5 characters (e.g., `ABCDE` or `JBPNL`). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Problems by severity* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Problems by severity* widget

Configure a *Problems by severity* widget that displays problem totals for all host groups.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
```

```
                    {
                        "type": "problemsbysv",
                        "name": "Problems by severity",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 0,
                                "name": "show_type",
                                "value": 1
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

25 Scatter plot

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Scatter plot* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Scatter plot* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Scatter plot* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Reference	1	reference	Any string value consisting of 5 characters (e.g., `ABCDE` or `JBPNL`). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Data set

The following parameters are supported for configuring a *Data set*.

Note:
The first number in the property name (e.g., ds.0.hosts.0, ds.0.items.0) represents the particular data set, while the second number, if present, represents the configured host or item.

Parameter	type	name	value
Data set type ***Data set: Item list***	0	ds.0.dataset_type	0 - Item list; 1 - *(default)* Item patterns.
Color	1	ds.0.color.0	Hexadecimal color code (e.g., `FF0000`). Parameter behavior: - *required* if *Data set type* is set to ”Item list”
X-Axis	4	ds.0.x_axis_itemids.0	Item ID. When configuring the widget on a template dashboard, only items configured on the template should be set. Note: To configure multiple items, create a dashboard widget field object for each item with an incremented number in the property name. Parameter behavior: - *required* if *Data set type* is set to ”Item list” and *X-axis (Widget)* is not set

Parameter	type	name	value
X-Axis (Widget)	1	ds.0.x_axis_itemids.0	Item ID: `ABCDE._itemid` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for items. Note: To configure multiple widgets, create a dashboard widget field object for each widget with an incremented number in the property name. Parameter behavior: - *required* if *Data set type* is set to "Item list" and *X-Axis* is not set
Y-Axis	4	ds.0.y_axis_itemids.0	Item ID. When configuring the widget on a template dashboard, only items configured on the template should be set. Note: To configure multiple items, create a dashboard widget field object for each item with an incremented number in the property name. Parameter behavior: - *required* if *Data set type* is set to "Item list" and *Y-axis (Widget)* is not set
Y-Axis (Widget)	1	ds.0.y_axis_itemids.0	Item ID: `ABCDE._itemid` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for items. Note: To configure multiple widgets, create a dashboard widget field object for each widget with an incremented number in the property name. Parameter behavior: - *required* if *Data set type* is set to "Item list" and *Y-Axis* is not set
Override host	1	ds.0.override_hostid	`ABCDE._hostid`- set a compatible widget (with its "Reference" parameter set to "`ABCDE`") as the data source for hosts; `DASHBOARD._hostid` - set the dashboard host selector as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Data set: Item patterns			
Color	1	ds.0.color	Hexadecimal color code (e.g., `FF0000`). Parameter behavior: - *supported* if *Data set type* is set to "Item patterns" and *Color palette* is not set
Color palette	0	ds.0.color_palette	Color palette index. Possible values range: 0-11. Default: 0. Parameter behavior: - *supported* if *Data set type* is set to "Item patterns" and *Color* is not set

Parameter	type	name	value
Host patterns	1	ds.0.hosts.0	Host name or pattern (e.g., ”Zabbix*”). Parameter behavior: - *required* if *Data set type* is set to ”Item patterns” This parameter is not supported if configuring the widget on a template dashboard.
X-Axis item patterns	1	ds.0.x_axis_items.0	Item name or pattern (e.g., ”*: Number of processed *values per second”). When configuring the widget on a template dashboard, only the patterns for items configured on the template should be set. Parameter behavior: - *required* if *Data set type* is set to ”Item patterns”
Y-Axis item patterns	1	ds.0.y_axis_items.0	Item name or pattern (e.g., ”*: Number of processed *values per second”). When configuring the widget on a template dashboard, only the patterns for items configured on the template should be set. Parameter behavior: - *required* if *Data set type* is set to ”Item patterns”
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to ”`ABCDE`”) as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.
Override host	1	ds.0.override_hostid._reference	Instead of Host ID: `ABCDE._hostid` - set a compatible widget (with its ”Reference” parameter set to ”`ABCDE`”) as the data source for hosts; `DASHBOARD._hostid` - set the dashboard host selector as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Host tags			
Evaluation type	0	evaltype_host	0 - *(default)* And/Or; 2 - Or. This parameter is not supported if configuring the widget on a template dashboard.
Tag name	1	host_tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Operator	0	host_tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Tag value	1	host_tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Marker	0	ds.0.marker	0 - *(default)* Ellipse; 1 - Square; 2 - Triangle; 3 - Diamond; 4 - Star; 5 - Cross.
Marker size	0	ds.0.marker_size	0 - *(default)* Small (6px); 1 - Medium (9px); 2 - Large (15px).
Time shift	1	ds.0.timeshift	Valid time string (e.g., 3600, 1h, etc.). You may use time suffixes. Negative values are also allowed. Default: "" (empty).
Aggregation interval	1	ds.0.aggregate_interval	Valid time string (e.g., 3600, 1h, etc.). You may use time suffixes. Default: 15m.
Aggregation function	0	ds.0.aggregate_function	0 - not used; 1 - min; 2 - max; 3 - *(default)* avg; 4 - count; 5 - sum; 6 - first; 7 - last.

Display options

The following parameters are supported for configuring *Display options*.

Parameter	type	name	value
History data selection	0	source	0 - *(default)* Auto; 1 - History; 2 - Trends.

Parameter	type	name	value
Host names in labels	0	show_hostnames	0 - *(default)* Auto; 1 - Show; 2 - Hide. This parameter is not supported if configuring the widget on a template dashboard.

Time period

The following parameters are supported for configuring *Time period*.

Parameter	type	name	value
Time period	1	time_period._reference	`DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source. Default: `DASHBOARD._timeperiod` Alternatively, you can set the time period only in the *From* and *To* parameters.
From	1	time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.to` is set
To	1	time_period.to	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.from` is set

Axes

The following parameters are supported for configuring *Axes*.

Parameter	type	name	value
X-Axis	0	x_axis	0 - Disabled; 1 - *(default)* Enabled.
Y-Axis	0	y_axis	0 - *(default)* Disabled; 1 - Enabled.
Min	1	x_axis_min	Any numeric value. Default: "" (empty).
		y_axis_min	
Max	1	x_axis_max	Any numeric value. Default: "" (empty).
		y_axis_max	
Units (type)	0	x_axis_units	0 - *(default)* Auto; 1 - Static.
		y_axis_units	
Units (value)	1	x_axis_static_units	Any string value. Default: "" (empty).
		y_axis_static_units	

Legend

The following parameters are supported for configuring *Legend*.

Parameter	type	name	value
Show legend	0	legend	0 - Disabled; 1 - *(default)* Enabled.
Show aggregation function	0	legend_aggregation	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Show legend* is set to ”Enabled”
Rows	0	legend_lines_mode	0 - *(default)* Fixed; 1 - Variable. Parameter behavior: - *supported* if *Show legend* is set to ”Enabled”
Number of rows/ Maximum number of rows	0	legend_lines	Possible values range: 1-10. Default: 1. Parameter behavior: - *supported* if *Show legend* is set to ”Enabled”
Number of columns	0	legend_columns	Possible values range: 1-4. Default: 4. Parameter behavior: - *supported* if *Show legend* is set to ”Enabled”

Thresholds

The following parameters are supported for configuring *Thresholds*.

Note:
The number in the *Thresholds* property name (e.g., thresholds.0.color) references the threshold place in a list, sorted in ascending order. However, if thresholds are configured in a different order, the values will be sorted in ascending order after updating widget configuration in Zabbix frontend (e.g., `"thresholds.0.threshold":"5"` → `"thresholds.0.threshold":"1"`; `"thresholds.1.threshold":"1"` → `"thresholds.1.threshold": "5"`).

Parameter	type	name	value
Color interpolation	0	interpolation	0 - *(default)* Disabled; 1 - Enabled.
Color	1	thresholds.0.color	Hexadecimal color code (e.g., FF0000).
X-Axis	1	thresholds.0.x_axis_threshold	Any numeric value. Suffixes (e.g., ”1d”, ”2w”, ”4K”, ”8G”) are supported.
Y-Axis	1	thresholds.0.y_axis_threshold	Any numeric value. Suffixes (e.g., ”1d”, ”2w”, ”4K”, ”8G”) are supported.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Scatter plot* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Scatter plot* widget

Configure a *Scatter plot* widget in the following way:

- 5 data sets of type ”Item patterns”, each with default aggregation (avg, 15m) and different markers.
- The legend displays configured items in a single row and two columns.
- Thresholds are configured to color a marker red when its value reaches 80 on any axis.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "scatterplot",
                        "name": "Scatter plot",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 1,
                                "name": "ds.0.color",
                                "value": "0C5A87"
                            },
                            {
                                "type": 1,
                                "name": "ds.0.hosts.0",
                                "value": "de-frankfurt*"
                            },
                            {
                                "type": 1,
                                "name": "ds.0.x_axis_items.0",
                                "value": "Memory utilization"
                            },
                            {
                                "type": 1,
                                "name": "ds.0.y_axis_items.0",
                                "value": "CPU utilization"
                            },
                            {
                                "type": 0,
                                "name": "ds.0.marker",
                                "value": 0
                            },
                            {
                                "type": 0,
                                "name": "ds.0.marker_size",
                                "value": 1
                            },
                            {
                                "type": 1,
                                "name": "ds.1.color",
                                "value": "66B0D9"
                            },
                            {
                                "type": 1,
                                "name": "ds.1.hosts.0",
                                "value": "fr-paris*"
                            },
                            {
                                "type": 1,
                                "name": "ds.1.x_axis_items.0",
```

```
            "value": "Memory utilization"
        },
        {
            "type": 1,
            "name": "ds.1.y_axis_items.0",
            "value": "CPU utilization"
        },
        {
            "type": 0,
            "name": "ds.1.marker",
            "value": 1
        },
        {
            "type": 0,
            "name": "ds.1.marker_size",
            "value": 1
        },
        {
            "type": 1,
            "name": "ds.2.color",
            "value": "0A466A"
        },
        {
            "type": 1,
            "name": "ds.2.hosts.0",
            "value": "lv-riga*"
        },
        {
            "type": 1,
            "name": "ds.2.x_axis_items.0",
            "value": "Memory utilization"
        },
        {
            "type": 1,
            "name": "ds.2.y_axis_items.0",
            "value": "CPU utilization"
        },
        {
            "type": 0,
            "name": "ds.2.marker",
            "value": 2
        },
        {
            "type": 0,
            "name": "ds.2.marker_size",
            "value": 1
        },
        {
            "type": 1,
            "name": "ds.3.color",
            "value": "3394C3"
        },
        {
            "type": 1,
            "name": "ds.3.hosts.0",
            "value": "pl-warsaw*"
        },
        {
            "type": 1,
            "name": "ds.3.x_axis_items.0",
            "value": "Memory utilization"
        },
```

```
                {
                    "type": 1,
                    "name": "ds.3.y_axis_items.0",
                    "value": "CPU utilization"
                },
                {
                    "type": 0,
                    "name": "ds.3.marker",
                    "value": 3
                },
                {
                    "type": 0,
                    "name": "ds.3.marker_size",
                    "value": 1
                },
                {
                    "type": 1,
                    "name": "ds.4.color",
                    "value": "1492C8"
                },
                {
                    "type": 1,
                    "name": "ds.4.hosts.0",
                    "value": "se-stockholm*"
                },
                {
                    "type": 1,
                    "name": "ds.4.x_axis_items.0",
                    "value": "Memory utilization"
                },
                {
                    "type": 1,
                    "name": "ds.4.y_axis_items.0",
                    "value": "CPU utilization"
                },
                {
                    "type": 0,
                    "name": "ds.4.marker",
                    "value": 4
                },
                {
                    "type": 0,
                    "name": "ds.4.marker_size",
                    "value": 1
                },
                {
                    "type": 0,
                    "name": "legend_columns",
                    "value": 2
                },
                {
                    "type": 0,
                    "name": "interpolation",
                    "value": 1
                },
                {
                    "type": 1,
                    "name": "thresholds.0.color",
                    "value": "D40000"
                },
                {
                    "type": 1,
```

```
                                "name": "thresholds.0.x_axis_threshold",
                                "value": "80"
                            },
                            {
                                "type": 1,
                                "name": "thresholds.0.y_axis_threshold",
                                "value": ""
                            },
                            {
                                "type": 1,
                                "name": "thresholds.1.color",
                                "value": "D40000"
                            },
                            {
                                "type": 1,
                                "name": "thresholds.1.x_axis_threshold",
                                "value": ""
                            },
                            {
                                "type": 1,
                                "name": "thresholds.1.y_axis_threshold",
                                "value": "80"
                            },
                            {
                                "type": 1,
                                "name": "reference",
                                "value": "JQISY"
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field

- dashboard.create
- dashboard.update

26 SLA report

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *SLA report* widget in dashboard.create and dashboard.update methods.

Attention:
Widget fields properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *SLA report* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *SLA report* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - *(default)* No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
SLA	10	slaid.0	SLA ID. Parameter behavior: - *required*
Service	9	serviceid.0	Service ID.
Show periods	0	show_periods	Possible values range from 1 to 100. Default: 20.
From	1	date_from	Valid date string in format YYYY-MM-DD. Relative dates with modifiers d, w, M, y (e.g., now, now/d, now/w-1w, etc.) are supported.
To	1	date_to	Valid date string in format YYYY-MM-DD. Relative dates with modifiers d, w, M, y (e.g., now, now/d, now/w-1w, etc.) are supported.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *SLA report* widget. For more information on configuring a dashboard, see dashboard.create.

Configuring an *SLA report* widget

Configure an *SLA report* widget that displays the SLA report for SLA "4" service "2" for the period of last 30 days.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
```

```
                "widgets": [
                    {
                        "type": "slareport",
                        "name": "SLA report",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 10,
                                "name": "slaid.0",
                                "value": 4
                            },
                            {
                                "type": 9,
                                "name": "serviceid.0",
                                "value": 2
                            },
                            {
                                "type": 1,
                                "name": "date_from",
                                "value": "now-30d"
                            },
                            {
                                "type": 1,
                                "name": "date_to",
                                "value": "now"
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- dashboard.create
- dashboard.update

27 System information

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *System Information* widget in dashboard.create and dashboard.update methods.

Attention:
Widget fields properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *System information* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *System Information* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - *(default)* 15 minutes.
Show	0	info_type	0 - *(default)* System stats; 1 - High availability nodes.
Show software update check details	0	show_software_update_check_details	0 - (default) Disabled; 1 - Enabled. Parameter behavior: - *supported* if AllowSoftwareUpdateCheck in Zabbix server configuration is enabled and *Show* is set to "System stats"

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *System information* widget. For more information on configuring a dashboard, see dashboard.create.

Configuring a *System information* widget

Configure a *System information* widget that displays system stats with a refresh interval of 10 minutes and software update check enabled.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "systeminfo",
                        "name": "System information",
```

```
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 0,
                                "name": "rf_rate",
                                "value": 600
                            },
                            {
                                "type": 0,
                                "name": "show_software_update_check_details",
                                "value": 1
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

28 Top hosts

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Top hosts* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Top hosts* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Top hosts* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.
Hosts	3	hostids.0	Host ID. Note: To configure multiple hosts, create a dashboard widget field object for each host with an incremented number in the property name. For multiple hosts, the parameter *Host groups* must either be not configured at all or configured with at least one host group that the configured hosts belong to. This parameter is not supported if configuring the widget on a template dashboard.
Hosts (Widget/Dashboard)	1	hostids._reference	Instead of Host ID: `DASHBOARD.hostids` - set the dashboard host selector as the data source for hosts; `ABCDE._hostids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Host tags			
Evaluation type	0	evaltype	0 - *(default)* And/Or; 2 - Or. This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Tag name	1	tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Operator	0	tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Tag value	1	tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Show hosts in maintenance	0	maintenance	0 - *(default)* Disabled; 1 - Enabled.
Columns (see below)			
Order by	0	column	Column numeric value from the configured columns.
Order	0	order	2 - *(default)* Top N; 3 - Bottom N.
Host limit	0	show_lines	Possible values range from 1 to 1000. Default: 10. This parameter is not supported if configuring the widget on a template dashboard.

Columns

Columns have common parameters and additional parameters depending on the configuration of the parameter *Data*.

Note:
For all parameters related to columns the number in the property name (e.g., columns.0.name) references a column for which the parameter is configured.

The following parameters are supported for all columns.

Parameter	type	name	value
Name	1	columns.0.name	Any string value. Parameter behavior: - *required*
Data	0	columns.0.data	1 - Item value; 2 - Host name; 3 - Text. Parameter behavior: - *required*
Base color	1	columns.0.base_color	Hexadecimal color code (e.g., FF0000). Parameter behavior: - *required*

Item value

The following parameters are supported if *Data* is set to ”Item value”.

Note:
The first number in the *Thresholds* property name (e.g., columnsthresholds.0.color.0) references the column for which thresholds are configured, while the second number references threshold place in a list, sorted in ascending order. However, if thresholds are configured in a different order, the values will be sorted in ascending order after updating widget configuration in Zabbix frontend (e.g., `"threshold.0.threshold":"5"` → `"threshold.0.threshold":"1"`; `"threshold.1.threshold":"1"` → `"threshold.1.threshold": "5"`).

Parameter	type	name	value
Item	1	columns.0.item	Valid item name. When configuring the widget on a template dashboard, only items configured on the template should be set.
Display item value as	0	columns.0.display_value_as	0 - *(default)* Numeric; 1 - Text; 2 - Binary.
Display	0	columns.0.display	1 - *(default)* As is; 2 - Bar; 3 - Indicators; 6 - Sparkline. Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric”
Min	1	columns.0.min	Any numeric value. Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric” and *Display* is set to ”Bar” or ”Indicators”
Max	1	columns.0.max	Any numeric value. Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric” and *Display* is set to ”Bar” or ”Indicators”

Parameter	type	name	value
Decimal places	0	columns.0.decimal_places	Possible values range from 0 to 10. Default: 2. Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric”
Sparkline			
Width	0	columns.0.sparkline.width	Possible values range from 0 to 10. Default: 1. Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric” and *Display* is set to ”Sparkline”
Fill	0	columns.0.sparkline.fill	Possible values range from 0 to 10. Default: 3. Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric” and *Display* is set to ”Sparkline”
Color	1	columns.0.sparkline.color	Hexadecimal color code (e.g., FF0000). Default: 42A5F5. Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric” and *Display* is set to ”Sparkline”
Time period	1	columns.0.sparkline.time_period.from columns.0.sparkline.time_period.to	DASHBOARD._timeperiod - set the dashboard time period selector as the data source; ABCDE._timeperiod - set a compatible widget (with its reference parameter equal to ABCDE) as the data source. Default: DASHBOARD._timeperiod Alternatively, you can set the time period only in the *From* and *To* parameters. Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric” and *Display* is set to ”Sparkline”
From	1	columns.0.sparkline.time_period.from	Valid time string in absolute (YYYY-MM-DD hh:mm:ss) or relative time syntax (now, now/d, now/w-1w, etc.). Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric” and *Display* is set to ”Sparkline”
To	1	columns.0.sparkline.time_period.to	Valid time string in absolute (YYYY-MM-DD hh:mm:ss) or relative time syntax (now, now/d, now/w-1w, etc.). Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric” and *Display* is set to ”Sparkline” - *required* if columns.0.sparkline.time_period.from is set
History data	0	columns.0.sparkline.history	0 - *(default)* Auto; 1 - History; 2 - Trends. Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric” and *Display* is set to ”Sparkline”

Parameter	type	name	value
Thresholds			
Color	1	columnsthresholds.0.color	Hexadecimal color code (e.g., `FF0000`). Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric”
Threshold	1	columnsthresholds.0.threshold	Any string value. Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric”
Highlights			
Color	1	columns.0.highlights.0.color	Hexadecimal color code (e.g., `FF0000`). Parameter behavior: - *supported* if *Display item value as* is set to ”Text”
Pattern	1	columns.0.highlights.0.pattern	Any string value. Parameter behavior: - *supported* if *Display item value as* is set to ”Text”
Show thumbnail	0	columns.0.show_thumbnail	0 - *(default)* Disabled; 1 - Enabled. Parameter behavior: - *supported* if *Display item value as* is set to ”Binary”
Aggregation function	0	columns.0.aggregate_function	0 - *(default)* not used; 1 - min; 2 - max; 3 - avg; 4 - count; 5 - sum; 6 - first; 7 - last.
Time period	1	columns.0.time_period.reference	`DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its `reference` parameter equal to `ABCDE`) as the data source. Default: `DASHBOARD._timeperiod` Alternatively, you can set the time period only in the *From* and *To* parameters. Parameter behavior: - *supported* if *Aggregation function* is set to ”min”, ”max”, ”avg”, ”count”, ”sum”, ”first”, ”last”
From	1	columns.0.time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set and *Aggregation function* is set to ”min”, ”max”, ”avg”, ”count”, ”sum”, ”first”, ”last” - *required* if `columns.0.time_period.to` is set
To	1	columns.0.time_period.to	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set and *Aggregation function* is set to ”min”, ”max”, ”avg”, ”count”, ”sum”, ”first”, ”last” - *required* if `columns.0.time_period.from` is set

Parameter	type	name	value
History data	0	columns.0.history	0 - *(default)* Auto; 1 - History; 2 - Trends. Parameter behavior: - *supported* if *Display item value as* is set to ”Numeric”
Reference	1	reference	Any string value consisting of 5 characters (e.g., `ABCDE` or `JBPNL`). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Text

The following parameters are supported if *Data* is set to ”Text”.

Parameter	type	name	value
Text	1	columns.0.text	Any string value, including macros. Supported macros: `{HOST.*}`, `{INVENTORY.*}`. Parameter behavior: - *required* if *Data* is set to ”Text”

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Top hosts* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Top hosts* widget

Configure a *Top hosts* widget that displays top hosts by CPU utilization in host group ”4”. In addition, configure the following custom columns: ”Host name”, ”CPU utilization in %”, ”1m avg”, ”5m avg”, ”15m avg”, ”Processes”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "tophosts",
                        "name": "Top hosts",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 2,
                                "name": "groupids.0",
                                "value": 4
                            },
                            {
                                "type": 1,
```

```
            "name": "columns.0.name",
            "value": "Host"
        },
        {
            "type": 0,
            "name": "columns.0.data",
            "value": 2
        },
        {
            "type": 1,
            "name": "columns.0.base_color",
            "value": "FFFFFF"
        },
        {
            "type": 1,
            "name": "columns.1.name",
            "value": "CPU utilization in %"
        },
        {
            "type": 0,
            "name": "columns.1.data",
            "value": 1
        },
        {
            "type": 1,
            "name": "columns.1.base_color",
            "value": "4CAF50"
        },
        {
            "type": 1,
            "name": "columns.1.item",
            "value": "CPU utilization"
        },
        {
            "type": 0,
            "name": "columns.1.display",
            "value": 3
        },
        {
            "type": 1,
            "name": "columns.1.min",
            "value": "0"
        },
        {
            "type": 1,
            "name": "columns.1.max",
            "value": "100"
        },
        {
            "type": 1,
            "name": "columnsthresholds.1.color.0",
            "value": "FFFF00"
        },
        {
            "type": 1,
            "name": "columnsthresholds.1.threshold.0",
            "value": "50"
        },
        {
            "type": 1,
            "name": "columnsthresholds.1.color.1",
            "value": "FF8000"
```

```
                    },
                    {
                        "type": 1,
                        "name": "columnsthresholds.1.threshold.1",
                        "value": "80"
                    },
                    {
                        "type": 1,
                        "name": "columnsthresholds.1.color.2",
                        "value": "FF4000"
                    },
                    {
                        "type": 1,
                        "name": "columnsthresholds.1.threshold.2",
                        "value": "90"
                    },
                    {
                        "type": 1,
                        "name": "columns.2.name",
                        "value": "1m avg"
                    },
                    {
                        "type": 0,
                        "name": "columns.2.data",
                        "value": 1
                    },
                    {
                        "type": 1,
                        "name": "columns.2.base_color",
                        "value": "FFFFFF"
                    },
                    {
                        "type": 1,
                        "name": "columns.2.item",
                        "value": "Load average (1m avg)"
                    },
                    {
                        "type": 1,
                        "name": "columns.3.name",
                        "value": "5m avg"
                    },
                    {
                        "type": 0,
                        "name": "columns.3.data",
                        "value": 1
                    },
                    {
                        "type": 1,
                        "name": "columns.3.base_color",
                        "value": "FFFFFF"
                    },
                    {
                        "type": 1,
                        "name": "columns.3.item",
                        "value": "Load average (5m avg)"
                    },
                    {
                        "type": 1,
                        "name": "columns.4.name",
                        "value": "15m avg"
                    },
                    {
```

```
                            "type": 0,
                            "name": "columns.4.data",
                            "value": 1
                        },
                        {
                            "type": 1,
                            "name": "columns.4.base_color",
                            "value": "FFFFFF"
                        },
                        {
                            "type": 1,
                            "name": "columns.4.item",
                            "value": "Load average (15m avg)"
                        },
                        {
                            "type": 1,
                            "name": "columns.5.name",
                            "value": "Processes"
                        },
                        {
                            "type": 0,
                            "name": "columns.5.data",
                            "value": 1
                        },
                        {
                            "type": 1,
                            "name": "columns.5.base_color",
                            "value": "FFFFFF"
                        },
                        {
                            "type": 1,
                            "name": "columns.5.item",
                            "value": "Number of processes"
                        },
                        {
                            "type": 0,
                            "name": "columns.5.decimal_places",
                            "value": 0
                        },
                        {
                            "type": 0,
                            "name": "column",
                            "value": 1
                        }
                    ]
                }
            ]
        }
    ],
    "userGroups": [
        {
            "usrgrpid": 7,
            "permission": 2
        }
    ],
    "users": [
        {
            "userid": 1,
            "permission": 3
        }
    ]
},
```

```
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

29 Top items

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Top items* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Top items* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Top items* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Hosts	3	hostids.0	Host ID. Note: To configure multiple hosts, create a dashboard widget field object for each host with an incremented number in the property name. For multiple hosts, the parameter *Host groups* must either be not configured at all or configured with at least one host group that the configured hosts belong to. This parameter is not supported if configuring the widget on a template dashboard.
Hosts (Widget/Dashboard)	1	hostids._reference	Instead of Host ID: `DASHBOARD.hostids` - set the dashboard host selector as the data source for hosts; `ABCDE._hostids` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Host tags			
Evaluation type	0	host_tags_evaltype	0 - *(default)* And/Or; 2 - Or. This parameter is not supported if configuring the widget on a template dashboard.
Tag name	1	host_tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Operator	0	host_tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Tag value	1	host_tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Host tags* This parameter is not supported if configuring the widget on a template dashboard.
Layout	0	layout	0 - *(default)* Horizontal; 1 - Vertical.

Parameter	type	name	value
Show problems *Items* (see below)	0	show_problems	0 - All; 1 - *(default)* Unsuppressed; 2 - None.

Advanced configuration

The following advanced configuration parameters are supported for the *Top items* widget.

Parameter	type	name	value
Host ordering			
Order by	0	host_ordering_order_by	0 - *(default)* Host name; 3 - Item value.
Item patterns	1	host_ordering_item.0	Item name or pattern (e.g., "*: Number of processed *values per second"). Note: To configure multiple item patterns, create a dashboard widget field object for each item pattern with an incremented number in the property name. When configuring the widget on a template dashboard, only the patterns for items configured on the template should be set. Parameter behavior: - *required* if *Order by* is set to "Item value"
Order	0	host_ordering_order	2 - *(default)* Top N; 3 - Bottom N.
Limit	0	host_ordering_limit	Possible values range from 1 to 100. Default: 10.
Item ordering			
Order by	0	item_ordering_order_by	1 - Host; 2 - Item name; 3 - *(default)* Item value.
Host patterns	1	item_ordering_host.0	Host name or pattern. Note: To configure multiple host patterns, create a dashboard widget field object for each host pattern with an incremented number in the property name. Parameter behavior: - *required* if *Order by* is set to "Host"
Order	0	item_ordering_order	2 - *(default)* Top N; 3 - Bottom N.
Limit	0	item_ordering_limit	Possible values range from 1 to 100. Default: 10.

Columns

Note:
The first number in the property name (e.g., columns.0.items.0, columns.0.item_tags_evaltype) represents the particular column, while the second number, if present, represents the configured entity (e.g., item pattern, tag).

Parameter	type	name	value
Item patterns	1	columns.0.items.0	Item name or pattern (e.g., "*: Number of processed *values per second"). Note: To configure multiple item patterns, create a dashboard widget field object for each item pattern with an incremented second number in the property name. When configuring the widget on a template dashboard, only the patterns for items configured on the template should be set. Parameter behavior: - *required*
Item tags			
Evaluation type	0	columns.0.item_tags_evaltype	0 - *(default)* And/Or; 2 - Or.
Tag name	1	columns.0.item_tags.0.tag	Any string value. Note: The second number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*
Operator	0	columns.0.item_tags.0.operator	**0** - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The second number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*
Tag value	1	columns.0.item_tags.0.value	Any string value. Note: The second number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Item tags*
Base color	1	columns.0.base_color	Hexadecimal color code (e.g., FF0000).
Display value as	0	columns.0.display_value_as	1 - *(default)* Numeric; 2 - Text.
Display	0	columns.0.display	1 - *(default)* As is; 2 - Bar; 3 - Indicators; 6 - Sparkline.

Parameter	type	name	value
Min	1	columns.0.min	Any numeric value. Suffixes (e.g., ”1d”, ”2w”, ”4K”, ”8G”) are supported. Parameter behavior: - *supported* if *Display value as* is set to ”Numeric” and *Display* is set to ”Bar” or ”Indicators”
Max	1	columns.0.max	Any numeric value. Suffixes (e.g., ”1d”, ”2w”, ”4K”, ”8G”) are supported. Parameter behavior: - *supported* if *Display value as* is set to ”Numeric” and *Display* is set to ”Bar” or ”Indicators”
Sparkline			
Width	0	columns.0.sparkline.width	Possible values range from 0 to 10. Default: 1. Parameter behavior: - *supported* if *Display* is set to ”Sparkline”
Fill	0	columns.0.sparkline.fill	Possible values range from 0 to 10. Default: 3. Parameter behavior: - *supported* if *Display* is set to ”Sparkline”
Color	1	columns.0.sparkline.color	Hexadecimal color code (e.g., `FF0000`). Default: `42A5F5`. Parameter behavior: - *supported* if *Display* is set to ”Sparkline”
Time period	1	columns.0.sparkline.time_period.reference	`DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its `reference` parameter equal to `ABCDE`) as the data source. Default: `DASHBOARD._timeperiod` Alternatively, you can set the time period only in the *From* and *To* parameters. Parameter behavior: - *supported* if *Display* is set to ”Sparkline”
From	1	columns.0.sparkline.time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Display* is set to ”Sparkline” - *required* if `columns.0.sparkline.time_period.to` is set
To	1	columns.0.sparkline.time_period.to	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Display* is set to ”Sparkline” - *required* if `columns.0.sparkline.time_period.from` is set
History data	0	columns.0.sparkline.history	0 - *(default)* Auto; 1 - History; 2 - Trends. Parameter behavior: - *supported* if *Display* is set to ”Sparkline”

Parameter	type	name	value
Thresholds			
Color	1	columns.0.thresholds.color	Hexadecimal color code (e.g., `FF0000`). Parameter behavior: - *supported* if *Display value as* is set to ”Numeric”
Threshold	1	columns.0.thresholds.threshold	Any numeric value. Suffixes (e.g., ”1d”, ”2w”, ”4K”, ”8G”) are supported. Parameter behavior: - *supported* if *Display value as* is set to ”Numeric”
Highlights			
Highlight	1	columns.0.highlights.color	Hexadecimal color code (e.g., `FF0000`). Parameter behavior: - *supported* if *Display value as* is set to ”Text”
Threshold	1	columns.0.highlights.pattern	Any regular expression. Parameter behavior: - *supported* if *Display value as* is set to ”Text”
Decimal places	0	columns.0.decimal_places	Possible values range from 0 to 10. Default: 2.
Advanced configuration			
Aggregation function	0	columns.0.aggregate_function	If *Display value as* is set to ”Numeric”: 0 - *(default)* not used; 1 - min; 2 - max; 3 - avg; 4 - count; 5 - sum; 6 - first; 7 - last. If *Display value as* is set to ”Text”: 0 - *(default)* not used; 4 - count; 6 - first; 7 - last.
Time period	1	columns.0.time_period.filter	`DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its `reference` parameter equal to `ABCDE`) as the data source. Default: `DASHBOARD._timeperiod` Alternatively, you can set the time period only in the *From* and *To* parameters. Parameter behavior: - *supported* if *Aggregation function* is set to ”min”, ”max”, ”avg”, ”count”, ”sum”, ”first”, ”last”

Parameter	type	name	value
From	1	columns.0.time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set and *Aggregation function* is set to ”min”, ”max”, ”avg”, ”count”, ”sum”, ”first”, ”last” - *required* if `columns.0.time_period.to` is set
To	1	columns.0.time_period.to	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set and *Aggregation function* is set to ”min”, ”max”, ”avg”, ”count”, ”sum”, ”first”, ”last” - *required* if `columns.0.time_period.from` is set
History data	0	columns.0.history	0 - *(default)* Auto; 1 - History; 2 - Trends. Parameter behavior: - *supported* if *Display value as* is set to ”Numeric”
Aggregate	0	columns.0.aggregate_function	0 - *(default)* Each item 1 - Combined. Parameter behavior: - *supported* if *Display value as* is set to ”Numeric” and *Display* is set to ”As is”, ”Bar” or ”Indicators”.
Combined aggregation function	0	columns.0.column_aggregate_function	1 - min; 2 - max; 3 - avg; 4 - count; 5 - sum *(default)*. Parameter behavior: - *required* if configuring *Aggregate*.
Combined column name	1	columns.0.combined_column_name	Any string value. Parameter behavior: - *required* if configuring *Aggregate*.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Top items* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Top items* widget

Configure a *Top items* widget that displays data for host ”10084” and only for items for which the tag with the name ”component” contains the value ”cpu”. In addition, display the data with hosts located at the top and use a colored gauge bar for cell representation.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
```

```
                        "type": "topitems",
                        "name": "Top items",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 3,
                                "name": "hostids.0",
                                "value": 10084
                            },
                            {
                                "type": 1,
                                "name": "columns.0.items.0",
                                "value": "*"
                            },
                            {
                                "type": 1,
                                "name": "columns.0.item_tags.0.tag",
                                "value": "component"
                            },
                            {
                                "type": 0,
                                "name": "columns.0.item_tags.0.operator",
                                "value": 0
                            },
                            {
                                "type": 1,
                                "name": "columns.0.item_tags.0.value",
                                "value": "cpu"
                            },
                            {
                                "type": 0,
                                "name": "columns.0.display",
                                "value": 2
                            },
                            {
                                "type": 0,
                                "name": "layout",
                                "value": 1
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
```

```
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

30 Top triggers

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Top triggers* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Top triggers* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Top triggers* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - *(default)* No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Hosts	3	hostids.0	Host ID. Note: To configure multiple hosts, create a dashboard widget field object for each host with an incremented number in the property name. For multiple hosts, the parameter *Host groups* must either be not configured at all or configured with at least one host group that the configured hosts belong to. This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Problem	1	problem	Problem event name (case insensitive, full name or part of it).
Severity	0	severities.0	Trigger severities. 0 - Not classified; 1 - Information; 2 - Warning; 3 - Average; 4 - High; 5 - Disaster. Default: empty (all enabled). Note: To configure multiple values, create a dashboard widget field object for each value with an incremented number in the property name.
Problem tags			
Evaluation type	0	evaltype	0 - *(default)* And/Or; 2 - Or.
Tag name	1	tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Operator	0	tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Tag value	1	tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Time period	1	time_period._reference	`DASHBOARD._timeperiod` - set the dashboard time period selector as the data source; `ABCDE._timeperiod` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source. Default: `DASHBOARD._timeperiod` Alternatively, you can set the time period only in the *From* and *To* parameters.
From	1	time_period.from	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.to` is set

Parameter	type	name	value
To	1	time_period.to	Valid time string in absolute (`YYYY-MM-DD hh:mm:ss`) or relative time syntax (`now`, `now/d`, `now/w-1w`, etc.). Parameter behavior: - *supported* if *Time period* is not set - *required* if `time_period.from` is set
Trigger limit	0	show_lines	Possible values range from 1 to 100. Default: 10.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Top triggers* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Top triggers* widget

Configure a *Top triggers* widget that displays the top 5 triggers for host group "4" with the count of all problems for each trigger. The widget displays only triggers that have severities "Warning", "Average", "High", or "Disaster", and problems that have a tag with the name "scope" that contains values "performance" or "availability", or "capacity".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "toptriggers",
                        "name": "Top triggers",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 2,
                                "name": "groupids.0",
                                "value": 4
                            },
                            {
                                "type": 0,
                                "name": "severities.0",
                                "value": 2
                            },
                            {
                                "type": 0,
                                "name": "severities.1",
                                "value": 3
                            },
                            {
                                "type": 0,
                                "name": "severities.2",
                                "value": 4
                            },
                            {
```

```
                            "type": 0,
                            "name": "severities.3",
                            "value": 5
                        },
                        {
                            "type": 1,
                            "name": "tags.0.tag",
                            "value": "scope"
                        },
                        {
                            "type": 0,
                            "name": "tags.0.operator",
                            "value": 0
                        },
                        {
                            "type": 1,
                            "name": "tags.0.value",
                            "value": "performance"
                        },
                        {
                            "type": 1,
                            "name": "tags.1.tag",
                            "value": "scope"
                        },
                        {
                            "type": 0,
                            "name": "tags.1.operator",
                            "value": 0
                        },
                        {
                            "type": 1,
                            "name": "tags.1.value",
                            "value": "availability"
                        },
                        {
                            "type": 1,
                            "name": "tags.2.tag",
                            "value": "scope"
                        },
                        {
                            "type": 0,
                            "name": "tags.2.operator",
                            "value": 0
                        },
                        {
                            "type": 1,
                            "name": "tags.2.value",
                            "value": "capacity"
                        },
                        {
                            "type": 0,
                            "name": "show_lines",
                            "value": 5
                        }
                    ]
                }
            ]
        }
    ],
    "userGroups": [
        {
            "usrgrpid": 7,
```

```
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

31 Trigger overview

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Trigger overview* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Trigger overview* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Trigger Overview* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Show	0	show	1 - *(default)* Recent problems; 2 - Any; 3 - Problems.

Parameter	type	name	value
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.
Hosts	3	hostids.0	Host ID. Note: To configure multiple hosts, create a dashboard widget field object for each host with an incremented number in the property name. For multiple hosts, the parameter *Host groups* must either be not configured at all or configured with at least one host group that the configured hosts belong to. This parameter is not supported if configuring the widget on a template dashboard.
Hosts (Widget/Dashboard)	1	hostids._reference	Instead of Host ID: `DASHBOARD.hostids` - set the dashboard host selector as the data source for hosts; `ABCDE._hostids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Problem tags			
Evaluation type	0	evaltype	0 - *(default)* And/Or; 2 - Or.
Tag name	1	tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Operator	0	tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*

Parameter	type	name	value
Tag value	1	tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Problem tags*
Show suppressed problems	0	show_suppressed	0 - *(default)* Disabled; 1 - Enabled.
Layout	0	layout	0 - *(default)* Horizontal; 1 - Vertical.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Trigger overview* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *Trigger overview* widget

Configure a *Trigger overview* widget that displays trigger states for all host groups that have triggers with a tag that has the name "scope" and contains value "availability".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "trigover",
                        "name": "Trigger overview",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 1,
                                "name": "tags.0.tag",
                                "value": "scope"
                            },
                            {
                                "type": 0,
                                "name": "tags.0.operator",
                                "value": 0
                            },
                            {
                                "type": 1,
                                "name": "tags.0.value",
                                "value": "availability"
                            }
                        ]
                    }
```

```
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

32 URL

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *URL* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *URL* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *URL* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - *(default)* No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.

Parameter	type	name	value
URL	1	url	Valid URL string. Parameter behavior: - *required*
Override host	1	override_hostid._reference	`ABCDE._hostid` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for hosts; `DASHBOARD._hostid` - set the dashboard host selector as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *URL* widget. For more information on configuring a dashboard, see `dashboard.create`.

Configuring a *URL* widget

Configure a *URL* widget that displays the home page of Zabbix manual.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "url",
                        "name": "URL",
                        "x": 0,
                        "y": 0,
                        "width": 36,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 1,
                                "name": "url",
                                "value": "https://www.zabbix.com/documentation/8.2/en"
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
```

```
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- `dashboard.create`
- `dashboard.update`

33 Web monitoring

Description

These parameters and the possible property values for the respective dashboard widget field objects allow you to configure the *Web monitoring* widget in `dashboard.create` and `dashboard.update` methods.

Attention:
Widget `fields` properties are not validated during the creation or update of a dashboard. This allows users to modify built-in widgets and create custom widgets, but also introduces the risk of creating or updating widgets incorrectly. To ensure the successful creation or update of the *Web monitoring* widget, please refer to the parameter behavior outlined in the tables below.

Parameters

The following parameters are supported for the *Web monitoring* widget.

Parameter	type	name	value
Refresh interval	0	rf_rate	0 - No refresh; 10 - 10 seconds; 30 - 30 seconds; 60 - *(default)* 1 minute; 120 - 2 minutes; 600 - 10 minutes; 900 - 15 minutes.
Host groups	2	groupids.0	Host group ID. Note: To configure multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Host groups (Widget)	1	groupids._reference	Instead of Host group ID: `ABCDE._hostgroupids` - set a compatible widget (with its *Reference* parameter set to "ABCDE") as the data source for host groups. This parameter is not supported if configuring the widget on a template dashboard.

Parameter	type	name	value
Exclude host groups	2	exclude_groupids.0	Host group ID. Note: To exclude multiple host groups, create a dashboard widget field object for each host group with an incremented number in the property name. This parameter is not supported if configuring the widget on a template dashboard.
Hosts	3	hostids.0	Host ID. Note: To configure multiple hosts, create a dashboard widget field object for each host with an incremented number in the property name. For multiple hosts, the parameter *Host groups* must either be not configured at all or configured with at least one host group that the configured hosts belong to. This parameter is not supported if configuring the widget on a template dashboard.
Hosts (Widget/Dashboard)	1	hostids._reference	Instead of Host ID: `DASHBOARD.hostids` - set the dashboard host selector as the data source for hosts; `ABCDE._hostids` - set a compatible widget (with its *Reference* parameter set to "`ABCDE`") as the data source for hosts. This parameter is not supported if configuring the widget on a template dashboard.
Scenario tags			
Evaluation type	0	evaltype	0 - *(default)* And/Or; 2 - Or.
Tag name	1	tags.0.tag	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Scenario tags*
Operator	0	tags.0.operator	0 - Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Scenario tags*
Tag value	1	tags.0.value	Any string value. Note: The number in the property name references tag order in the tag evaluation list. Parameter behavior: - *required* if configuring *Scenario tags*
Show hosts in maintenance	0	maintenance	0 - Disabled; 1 - *(default)* Enabled.

Parameter	type	name	value
Reference	1	reference	Any string value consisting of 5 characters (e.g., ABCDE or JBPNL). This value must be unique within the dashboard to which the widget belongs. Parameter behavior: - *required*

Examples

The following examples aim to only describe the configuration of the dashboard widget field objects for the *Web monitoring* widget. For more information on configuring a dashboard, see dashboard.create.

Configuring a *Web monitoring* widget

Configure a *Web monitoring* widget that displays a status summary of the active web monitoring scenarios for host group "4".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dashboard.create",
    "params": {
        "name": "My dashboard",
        "display_period": 30,
        "auto_start": 1,
        "pages": [
            {
                "widgets": [
                    {
                        "type": "web",
                        "name": "Web monitoring",
                        "x": 0,
                        "y": 0,
                        "width": 18,
                        "height": 3,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 2,
                                "name": "groupids.0",
                                "value": 4
                            }
                        ]
                    }
                ]
            }
        ],
        "userGroups": [
            {
                "usrgrpid": 7,
                "permission": 2
            }
        ],
        "users": [
            {
                "userid": 1,
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Dashboard widget field
- dashboard.create
- dashboard.update

Device

This class is designed to work with mobile devices for use with Zabbix mobile app.

Object references:

- Device

Available methods:

- device.get - retrieve devices
- device.offboard - remove (offboard) devices

Device object

The Device object represents a mobile device for use with Zabbix mobile app.

Property	Type	Description
deviceid	ID	Device ID (internal, e.g., 1, 2). Property behavior: - Read-only
userid	ID	Device user's ID. Property behavior: - Constant
uuid	string	Device universally unique identifier (UUIDv7) (external, e.g., 019f930e-0ea2-7c1f-9dcc-0fbb60c40664). Property behavior: - Read-only
name	string	Device name. Property behavior: - Read-only
status	integer	Device status. Possible values: 0 - New 1 - Active 2 - Unassigned Property behavior: - Read-only

Property	Type	Description
activated_at	timestamp	Date and time the device was added in the mobile app, or 0 if not yet added. Property behavior: - Read-only
lastaccess	timestamp	Date and time the device was last authenticated in the mobile app, or 0 if not yet authenticated. Property behavior: - Read-only

device.get

This method allows you to retrieve mobile devices for use with Zabbix mobile app.

The method is available to *User*, *Admin*, and *Super admin* users if `EnableMobileDevices` is enabled and the user has access to devices. Access to API methods can be revoked in user role configuration.

User and *Admin* users can only retrieve their own devices.

Method signature: `integer/array device.get(object parameters)`

Method source: `CDevice::get()` in `ui/include/classes/api/services/CDevice.php`

Parameters

Parameter	Type	Description
deviceids	ID/array	Return only devices with the given device IDs.
userids	ID/array	Return only devices belonging to users with the given user IDs.
roleids	ID/array	Return only devices belonging to users with the given user role IDs.
usrgrpids	ID/array	Return only devices belonging to users with the given user group IDs.
sortfield	string/array	Sort the result by the given properties. Possible values: `deviceid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

Returns either an array of objects, or the count of retrieved objects (if `countOutput` is used).

Examples

Retrieve devices

Retrieve devices belonging to users with these user IDs: 1, 2, 3.
 Note that user 2 (guest) has no devices, and device 3 has not yet been added or authenticated in Zabbix mobile app.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "device.get",
    "params": {
        "output": "extend",
```

```
        "userids": [1, 2, 3]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "deviceid": "1",
            "userid": "1",
            "uuid": "019f930e-0ea2-7c1f-9dcc-0fbb60c40664",
            "name": "Google Pixel 10",
            "status": "1",
            "activated_at": "1784533122",
            "lastaccess": "1784879933"
        },
        {
            "deviceid": "2",
            "userid": "1",
            "uuid": "019d00fb-9e87-7ad5-bb4e-978467254aff",
            "name": "Apple iPhone 15 Pro",
            "status": "1",
            "activated_at": "1784533225",
            "lastaccess": "0"
        },
        {
            "deviceid": "3",
            "userid": "3",
            "uuid": "019d00fb-9e87-7ad5-bb4e-6ed88a2297ee",
            "name": "",
            "status": "0",
            "activated_at": "0",
            "lastaccess": "0"
        }
    ],
    "id": 1
}
```

device.offboard

This method allows you to remove (offboard) mobile devices for use with Zabbix mobile app.

The method is available to *User*, *Admin*, and *Super admin* users if `EnableMobileDevices` is enabled and the user has access to devices. Access to API methods can be revoked in user role configuration.

User and *Admin* users can only remove their own devices.

Method signature: `object device.offboard(object parameters)`

Method source: `CDevice::offboard()` in `ui/include/classes/api/services/CDevice.php`

Parameters

Parameter	Type	Description
uuid	string	Device universally unique identifier (UUIDv7). Matches `deviceid`.
force	boolean	Set `true` to ignore the "Device not found" error and forcibly remove the device. Default: `false`

Return values

Returns an object containing the UUID of the removed device.

Examples

Remove device

Remove a device with the given UUID, ignoring the "Device not found" error.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "device.offboard",
    "params": {
        "uuid": "019d0106-5eef-7154-a159-bb29e35d38e5",
        "force": true
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "uuid": "019d0106-5eef-7154-a159-bb29e35d38e5"
    },
    "id": 1
}
```

Discovered host

This class is designed to work with discovered hosts.

Object references:

- Discovered host

Available methods:

- dhost.get - retrieve discovered hosts

Discovered host object

The following objects are directly related to the `dhost` API.

Discovered host

> **Note:**
> Discovered host are created by the Zabbix server and cannot be modified via the API.

The discovered host object contains information about a host discovered by a network discovery rule. It has the following properties.

Property	Type	Description
dhostid	ID	ID of the discovered host.
druleid	ID	ID of the discovery rule that detected the host.
lastdown	timestamp	Time when the discovered host last went down.
lastup	timestamp	Time when the discovered host last went up.
status	integer	Whether the discovered host is up or down. A host is up if it has at least one active discovered service. Possible values: 0 - host up; 1 - host down.

dhost.get

Description

`integer/array dhost.get(object parameters)`

The method allows you to retrieve discovered hosts according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
dhostids	ID/array	Return only discovered hosts with the given IDs.
druleids	ID/array	Return only discovered hosts that have been created by the given discovery rules.
dserviceids	ID/array	Return only discovered hosts that are running the given services.
selectDRules	query	Return a `drules` property with an array of the discovery rules that detected the host.
selectDServices	query	Return a `dservices` property with the discovered services running on the host. Supports `count`.
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectDServices` - results will be sorted by `dserviceid`.
sortfield	string/array	Sort the result by the given properties. Possible values: `dhostid`, `druleid`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve discovered hosts by discovery rule

Retrieve all hosts and the discovered services they are running that have been detected by discovery rule "4".

Request:

```
{
    "jsonrpc": "2.0",
```

```
    "method": "dhost.get",
    "params": {
        "output": "extend",
        "selectDServices": "extend",
        "druleids": "4"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "dservices": [
                {
                    "dserviceid": "1",
                    "dhostid": "1",
                    "type": "4",
                    "key_": "",
                    "value": "",
                    "port": "80",
                    "status": "0",
                    "lastup": "1337697227",
                    "lastdown": "0",
                    "dcheckid": "5",
                    "ip": "192.168.1.1",
                    "dns": "station.company.lan"
                }
            ],
            "dhostid": "1",
            "druleid": "4",
            "status": "0",
            "lastup": "1337697227",
            "lastdown": "0"
        },
        {
            "dservices": [
                {
                    "dserviceid": "2",
                    "dhostid": "2",
                    "type": "4",
                    "key_": "",
                    "value": "",
                    "port": "80",
                    "status": "0",
                    "lastup": "1337697234",
                    "lastdown": "0",
                    "dcheckid": "5",
                    "ip": "192.168.1.4",
                    "dns": "john.company.lan"
                }
            ],
            "dhostid": "2",
            "druleid": "4",
            "status": "0",
            "lastup": "1337697234",
            "lastdown": "0"
        },
        {
            "dservices": [
                {
```

```
                    "dserviceid": "3",
                    "dhostid": "3",
                    "type": "4",
                    "key_": "",
                    "value": "",
                    "port": "80",
                    "status": "0",
                    "lastup": "1337697234",
                    "lastdown": "0",
                    "dcheckid": "5",
                    "ip": "192.168.1.26",
                    "dns": "printer.company.lan"
                }
            ],
            "dhostid": "3",
            "druleid": "4",
            "status": "0",
            "lastup": "1337697234",
            "lastdown": "0"
        },
        {
            "dservices": [
                {
                    "dserviceid": "4",
                    "dhostid": "4",
                    "type": "4",
                    "key_": "",
                    "value": "",
                    "port": "80",
                    "status": "0",
                    "lastup": "1337697234",
                    "lastdown": "0",
                    "dcheckid": "5",
                    "ip": "192.168.1.7",
                    "dns": "mail.company.lan"
                }
            ],
            "dhostid": "4",
            "druleid": "4",
            "status": "0",
            "lastup": "1337697234",
            "lastdown": "0"
        }
    ],
    "id": 1
}
```


See also

- Discovered service
- Discovery rule

Source

CDHost::get() in *ui/include/classes/api/services/CDHost.php*.

Discovered service

This class is designed to work with discovered services.

Object references:

- Discovered service

Available methods:

- dservice.get - retrieve discovered services

Discovered service object

The following objects are directly related to the `dservice` API.

Discovered service

Note:
Discovered services are created by the Zabbix server and cannot be modified via the API.

The discovered service object contains information about a service discovered by a network discovery rule on a host. It has the following properties.

Property	Type	Description
dserviceid	ID	ID of the discovered service.
dcheckid	ID	ID of the discovery check used to detect the service.
dhostid	ID	ID of the discovered host running the service.
dns	string	DNS of the host running the service.
ip	string	IP address of the host running the service.
lastdown	timestamp	Time when the discovered service last went down.
lastup	timestamp	Time when the discovered service last went up.
port	integer	Service port number.
status	integer	Status of the service. Possible values: 0 - service up; 1 - service down.
value	string	Value returned by the service when performing a Zabbix agent, SNMPv1, SNMPv2 or SNMPv3 discovery check.

dservice.get

Description

`integer/array dservice.get(object parameters)`

The method allows you to retrieve discovered services according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
dserviceids	ID/array	Return only discovered services with the given IDs.
dhostids	ID/array	Return only discovered services that belong to the given discovered hosts.
dcheckids	ID/array	Return only discovered services that have been detected by the given discovery checks.
druleids	ID/array	Return only discovered services that have been detected by the given discovery rules.
selectDRules	query	Return a `drules` property with an array of the discovery rules that detected the service.
selectDHosts	query	Return a `dhosts` property with an array the discovered hosts that the service belongs to.

Parameter	Type	Description
selectHosts	query	Return a `hosts` property with the hosts with the same IP address and proxy as the service. Supports `count`.
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectHosts` - result will be sorted by `hostid`.
sortfield	string/array	Sort the result by the given properties. Possible values: `dserviceid`, `dhostid`, `ip`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve services discovered on a host

Retrieve all discovered services detected on discovered host "11".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dservice.get",
    "params": {
        "output": "extend",
        "dhostids": "11"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "dserviceid": "12",
            "dhostid": "11",
            "value": "",
            "port": "80",
            "status": "1",
            "lastup": "0",
            "lastdown": "1348650607",
            "dcheckid": "5",
            "ip": "192.168.1.134",
            "dns": "john.local"
```

```
        },
        {
            "dserviceid": "13",
            "dhostid": "11",
            "value": "",
            "port": "21",
            "status": "1",
            "lastup": "0",
            "lastdown": "1348650610",
            "dcheckid": "6",
            "ip": "192.168.1.134",
            "dns": "john.local"
        }
    ],
    "id": 1
}
```


See also

- Discovered host
- Discovery check
- Host

Source

CDService::get() in *ui/include/classes/api/services/CDService.php*.

Discovery check

This class is designed to work with discovery checks.

Object references:

- Discovery check

Available methods:

- dcheck.get - retrieve discovery checks

Discovery check object

The following objects are directly related to the `dcheck` API.

Discovery check

The discovery check object defines a specific check performed by a network discovery rule. It has the following properties.

Property	Type	Description
dcheckid	ID	ID of the discovery check.
druleid	ID	ID of the discovery rule that the check belongs to.
key_	string	Item key (if `type` is set to "Zabbix agent") or SNMP OID (if `type` is set to "SNMPv1 agent", "SNMPv2 agent", or "SNMPv3 agent"). Property behavior: - *required* if `type` is set to "Zabbix agent", "SNMPv1 agent", "SNMPv2 agent", or "SNMPv3 agent"
ports	string	One or several port ranges to check, separated by commas. Default: 0. Property behavior: - *supported* if `type` is set to "SSH" (0), "LDAP" (1), "SMTP" (2), "FTP" (3), "HTTP" (4), "POP" (5), "NNTP" (6), "IMAP" (7), "TCP" (8), "Zabbix agent" (9), "SNMPv1 agent" (10), "SNMPv2 agent" (11), "SNMPv3 agent" (13), "HTTPS" (14), or "Telnet" (15)

Property	Type	Description
snmp_community	string	SNMP community. Property behavior: - *required* if `type` is set to ”SNMPv1 agent” or ”SNMPv2 agent”
snmpv3_authpassphrase	string	Authentication passphrase. Property behavior: - *supported* if `type` is set to ”SNMPv3 agent” and `snmpv3_securitylevel` is set to ”authNoPriv” or ”authPriv”
snmpv3_authprotocol	integer	Authentication protocol. Possible values: 0 - *(default)* MD5; 1 - SHA1; 2 - SHA224; 3 - SHA256; 4 - SHA384; 5 - SHA512. Property behavior: - *supported* if `type` is set to ”SNMPv3 agent” and `snmpv3_securitylevel` is set to ”authNoPriv” or ”authPriv”
snmpv3_contextname	string	SNMPv3 context name. Property behavior: - *supported* if `type` is set to ”SNMPv3 agent”
snmpv3_privpassphrase	string	Privacy passphrase. Property behavior: - *supported* if `type` is set to ”SNMPv3 agent” and `snmpv3_securitylevel` is set to ”authPriv”
snmpv3_privprotocol	integer	Privacy protocol. Possible values: 0 - *(default)* DES; 1 - AES128; 2 - AES192; 3 - AES256; 4 - AES192C; 5 - AES256C. Property behavior: - *supported* if `type` is set to ”SNMPv3 agent” and `snmpv3_securitylevel` is set to ”authPriv”
snmpv3_securitylevel	string	Security level. Possible values: 0 - noAuthNoPriv; 1 - authNoPriv; 2 - authPriv. Property behavior: - *supported* if `type` is set to ”SNMPv3 agent”
snmpv3_securityname	string	Security name. Property behavior: - *supported* if `type` is set to ”SNMPv3 agent”

Property	Type	Description
type	integer	Type of check. Possible values: 0 - SSH; 1 - LDAP; 2 - SMTP; 3 - FTP; 4 - HTTP; 5 - POP; 6 - NNTP; 7 - IMAP; 8 - TCP; 9 - Zabbix agent; 10 - SNMPv1 agent; 11 - SNMPv2 agent; 12 - ICMP ping; 13 - SNMPv3 agent; 14 - HTTPS; 15 - Telnet. Property behavior: - *required*
uniq	integer	Whether to use this check as a device uniqueness criteria. Only a single unique check can be configured for a discovery rule. Possible values: 0 - *(default)* do not use this check as a uniqueness criteria; 1 - use this check as a uniqueness criteria. Property behavior: - *supported* if `type` is set to ”Zabbix agent”, ”SNMPv1 agent”, ”SNMPv2 agent”, or ”SNMPv3 agent”
host_source	integer	Source for host name. Possible values: 1 - *(default)* DNS; 2 - IP; 3 - discovery value of this check.
name_source	integer	Source for visible name. Possible values: 0 - *(default)* not specified; 1 - DNS; 2 - IP; 3 - discovery value of this check.
allow_redirect	integer	Allow situation where the target being ICMP pinged responds from a different IP address. Possible values: 0 - *(default)* treat redirected responses as if the target host is down (fail); 1 - treat redirected responses as if the target host is up (success). Property behavior: - *supported* if `type` is set to ”ICMP ping”

dcheck.get

Description

`integer/array dcheck.get(object parameters)`

The method allows you to retrieve discovery checks according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
dcheckids	ID/array	Return only discovery checks with the given IDs.
druleids	ID/array	Return only discovery checks that belong to the given discovery rules.
dserviceids	ID/array	Return only discovery checks that have detected the given discovered services.
selectDRules	query	Return discovery rules related to the discovery checks.
sortfield	string/array	Sort the result by the given properties. Possible values: `dcheckid`, `druleid`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve discovery checks for a discovery rule

Retrieve all discovery checks used by discovery rule "6".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "dcheck.get",
    "params": {
        "output": "extend",
        "dcheckids": "6"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "dcheckid": "6",
            "druleid": "4",
```

```
            "type": "3",
            "key_": "",
            "snmp_community": "",
            "ports": "21",
            "snmpv3_securityname": "",
            "snmpv3_securitylevel": "0",
            "snmpv3_authpassphrase": "",
            "snmpv3_privpassphrase": "",
            "uniq": "0",
            "snmpv3_authprotocol": "0",
            "snmpv3_privprotocol": "0",
            "snmpv3_contextname": "",
            "host_source": "1",
            "name_source": "0",
            "allow_redirect": "0"
        }
    ],
    "id": 1
}
```


Source

CDCheck::get() in *ui/include/classes/api/services/CDCheck.php*.

Discovery rule

This class is designed to work with network discovery rules.

> **Note:**
> This API is meant to work with network discovery rules. For low-level discovery rules, see LLD rule API.

Object references:

- Discovery rule

Available methods:

- drule.create - create new network discovery rules
- drule.delete - delete network discovery rules
- drule.get - retrieve network discovery rules
- drule.update - update network discovery rules

Discovery rule object

The following objects are directly related to the `drule` API.

Discovery rule

The discovery rule object defines a network discovery rule. It has the following properties.

Property	Type	Description
druleid	ID	ID of the discovery rule. Property behavior: - *read-only* - *required* for update operations
iprange	string	One or several IP ranges to check, separated by commas. Refer to the network discovery configuration section for more information on supported formats of IP ranges. Property behavior: - *required* for create operations

Property	Type	Description
name	string	Name of the discovery rule. Property behavior: - *required* for create operations
delay	string	Execution interval of the discovery rule. Accepts a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., 30s, 1m, 2h, 1d), or a user macro. Default: 1h.
proxyid	ID	ID of the proxy used for discovery.
status	integer	Whether the discovery rule is enabled. Possible values: 0 - *(default)* enabled; 1 - disabled.
concurrency_max	integer	Maximum number of concurrent checks per discovery rule. Possible values: 0 - *(default)* unlimited number of checks; 1 - one check; 2-999 - custom number of checks.
error	string	Error text if there have been any problems when executing the discovery rule. Property behavior: - *read-only*

drule.create

Description

`object drule.create(object/array discoveryRules)`

This method allows you to create new discovery rules.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Discovery rules to create.

Additionally to the standard discovery rule properties, the method accepts the following parameters.

Parameter	Type	Description
dchecks	array	Discovery checks to create for the discovery rule. Parameter behavior: - *required*

Return values

`(object)` Returns an object containing the IDs of the created discovery rules under the `druleids` property. The order of the returned IDs matches the order of the passed discovery rules.

Examples

Create a discovery rule

Create a discovery rule to find machines running the Zabbix agent in the local network. The rule must use a single Zabbix agent check on port 10050.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "drule.create",
    "params": {
        "name": "Zabbix agent discovery",
        "iprange": "192.168.1.1-255",
        "concurrency_max": "10",
        "dchecks": [
            {
                "type": "9",
                "key_": "system.uname",
                "ports": "10050",
                "uniq": "0"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "druleids": [
            "6"
        ]
    },
    "id": 1
}
```

See also

- Discovery check

Source

CDRule::create() in *ui/include/classes/api/services/CDRule.php*.

drule.delete

Description

`object drule.delete(array discoveryRuleIds)`

This method allows you to delete discovery rules.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the discovery rules to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted discovery rules under the `druleids` property.

Examples

Delete multiple discovery rules

Delete two discovery rules.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "drule.delete",
    "params": [
        "4",
        "6"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "druleids": [
            "4",
            "6"
        ]
    },
    "id": 1
}
```

Source

CDRule::delete() in *ui/include/classes/api/services/CDRule.php*.

drule.get

Description

`integer/array drule.get(object parameters)`

The method allows you to retrieve discovery rules according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
dhostids	ID/array	Return only discovery rules that created the given discovered hosts.
druleids	ID/array	Return only discovery rules with the given IDs.
dserviceids	ID/array	Return only discovery rules that created the given discovered services.
selectDChecks	query	Return a `dchecks` property with the discovery checks used by the discovery rule. Supports `count`.
selectDHosts	query	Return a `dhosts` property with the discovered hosts created by the discovery rule. Supports `count`.
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectDChecks` - results will be sorted by `dcheckid`; `selectDHosts` - results will be sorted by `dhostsid`.
sortfield	string/array	Sort the result by the given properties. Possible values: `druleid`, `name`.

Parameter	Type	Description
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve all discovery rules

Retrieve all configured discovery rules and the discovery checks they use.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "drule.get",
    "params": {
        "output": "extend",
        "selectDChecks": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "druleid": "2",
            "proxyid": "0",
            "name": "Local network",
            "iprange": "192.168.3.1-255",
            "delay": "5s",
            "status": "0",
            "concurrency_max": "0",
            "error": "",
            "dchecks": [
                {
                    "dcheckid": "7",
                    "druleid": "2",
                    "type": "3",
                    "key_": "",
                    "snmp_community": "",
                    "ports": "21",
                    "snmpv3_securityname": "",
                    "snmpv3_securitylevel": "0",
                    "snmpv3_authpassphrase": "",
                    "snmpv3_privpassphrase": "",
                    "uniq": "0",
```

```
                    "snmpv3_authprotocol": "0",
                    "snmpv3_privprotocol": "0",
                    "snmpv3_contextname": "",
                    "host_source": "1",
                    "name_source": "0",
                    "allow_redirect": "0"
                },
                {
                    "dcheckid": "8",
                    "druleid": "2",
                    "type": "4",
                    "key_": "",
                    "snmp_community": "",
                    "ports": "80",
                    "snmpv3_securityname": "",
                    "snmpv3_securitylevel": "0",
                    "snmpv3_authpassphrase": "",
                    "snmpv3_privpassphrase": "",
                    "uniq": "0",
                    "snmpv3_authprotocol": "0",
                    "snmpv3_privprotocol": "0",
                    "snmpv3_contextname": "",
                    "host_source": "1",
                    "name_source": "0",
                    "allow_redirect": "0"
                }
            ]
        },
        {
            "druleid": "6",
            "proxyid": "0",
            "name": "Zabbix agent discovery",
            "iprange": "192.168.1.1-255",
            "delay": "1h",
            "status": "0",
            "concurrency_max": "10",
            "error": "",
            "dchecks": [
                {
                    "dcheckid": "10",
                    "druleid": "6",
                    "type": "9",
                    "key_": "system.uname",
                    "snmp_community": "",
                    "ports": "10050",
                    "snmpv3_securityname": "",
                    "snmpv3_securitylevel": "0",
                    "snmpv3_authpassphrase": "",
                    "snmpv3_privpassphrase": "",
                    "uniq": "0",
                    "snmpv3_authprotocol": "0",
                    "snmpv3_privprotocol": "0",
                    "snmpv3_contextname": "",
                    "host_source": "2",
                    "name_source": "3",
                    "allow_redirect": "0"
                }
            ]
        }
    ],
    "id": 1
}
```

See also

- Discovered host
- Discovery check

Source

CDRule::get() in *ui/include/classes/api/services/CDRule.php*.

drule.update

Description

`object drule.update(object/array discoveryRules)`

This method allows you to update existing discovery rules.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Discovery rule properties to be updated.

The `druleid` property must be defined for each discovery rule, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard discovery rule properties, the method accepts the following parameters.

Parameter	Type	Description
dchecks	array	Discovery checks to replace existing checks.

Return values

`(object)` Returns an object containing the IDs of the updated discovery rules under the `druleids` property.

Examples

Change the IP range of a discovery rule

Change the IP range of a discovery rule to ”192.168.2.1-255”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "drule.update",
    "params": {
        "druleid": "6",
        "iprange": "192.168.2.1-255"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "druleids": [
            "6"
        ]
    },
    "id": 1
}
```

See also

- Discovery check

Source

CDRule::update() in *ui/include/classes/api/services/CDRule.php*.

Event

This class is designed to work with events.

Object references:

- Event
- Event tag
- Media type URL

Available methods:

- event.get - retrieve events
- event.acknowledge - acknowledge events

Event object

The following objects are directly related to the `event` API.

Event

Note:
Events are created by the Zabbix server and cannot be modified via the API.

The event object has the following properties.

Property	Type	Description
eventid	ID	ID of the event.
source	integer	Type of the event. Possible values: 0 - event created by a trigger; 1 - event created by a discovery rule; 2 - event created by active agent autoregistration; 3 - internal event; 4 - event created on service status update.
object	integer	Type of object that is related to the event. Possible values if `source` is set to ”event created by a trigger”: 0 - trigger. Possible values if `source` is set to ”event created by a discovery rule”: 1 - discovered host; 2 - discovered service. Possible values if `source` is set to ”event created by active agent autoregistration”: 3 - auto-registered host. Possible values if `source` is set to ”internal event”: 0 - trigger; 4 - item; 5 - LLD rule. Possible values if `source` is set to ”event created on service status update”: 6 - service.
objectid	ID	ID of the related object.

Property	Type	Description
acknowledged	integer	Whether the event has been acknowledged.
clock	timestamp	Time when the event was created.
ns	integer	Nanoseconds when the event was created.
name	string	Resolved event name.
value	integer	State of the related object. Possible values if `source` is set to ”event created by a trigger” or ”event created on service status update”: 0 - OK; 1 - problem. Possible values if `source` is set to ”event created by a discovery rule”: 0 - host or service up; 1 - host or service down; 2 - host or service discovered; 3 - host or service lost. Possible values if `source` is set to ”internal event”: 0 - ”normal” state; 1 - ”unknown” or ”not supported” state. Property behavior: - *supported* if `source` is set to ”event created by a trigger”, ”event created by a discovery rule”, ”internal event”, or ”event created on service status update”
severity	integer	Event current severity. Possible values: 0 - not classified; 1 - information; 2 - warning; 3 - average; 4 - high; 5 - disaster.
r_eventid	ID	ID of the recovery event.
c_eventid	ID	ID of the event that was used to override (close) current event under global correlation rule. See `correlationid` to identify exact correlation rule. This parameter is only defined when the event is closed by global correlation rule.
cause_eventid	ID	ID of the cause event.
correlationid	ID	ID of the correlation rule that generated closing of the problem. This parameter is only defined when the event is closed by global correlation rule.
userid	ID	ID of the user that closed the event (if the event was closed manually).
suppressed	integer	Whether the event is suppressed. Possible values: 0 - event is in normal state; 1 - event is suppressed.
opdata	string	Operational data with expanded macros.
urls	array	Active media type URLs.

Event tag

The event tag object has the following properties.

Property	Type	Description
tag	string	Event tag name.
value	string	Event tag value.

Media type URL

The media type URL object has the following properties.

Property	Type	Description
name	string	Media type defined URL name.
url	string	Media type defined URL value.

Results will contain entries only for active media types with enabled event menu entry. Macro used in properties will be expanded, but if one of the properties contains an unexpanded macro, both properties will be excluded from results. For supported macros, see *Supported macros*.

event.acknowledge

Description

`object event.acknowledge(object/array parameters)`

This method allows you to update events. The following update actions can be performed:

- Close event. If the event is already resolved, this action will be skipped.
- Acknowledge event. If the event is already acknowledged, this action will be skipped.
- Unacknowledge event. If the event is not acknowledged, this action will be skipped.
- Add message.
- Change event severity. If the event already has the same severity, this action will be skipped.
- Suppress event. If the event is already suppressed, this action will be skipped.
- Unsuppress event. If the event is not suppressed, this action will be skipped.
- Change event rank to cause. If the event is already ranked as cause, this action will be skipped.
- Change event rank to symptom. If the event is already ranked as symptom, this action will be skipped.

Attention:
Only trigger events can be updated.
 Only problem events can be updated.
 Read/Write rights for trigger are required to close the event or to change event's severity.
 To close an event, manual close should be allowed in the trigger.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Parameters containing the IDs of the events and update operations that should be performed.

Parameter	Type	Description
eventids	ID/array	IDs of the events to acknowledge. Parameter behavior: *- required*

Parameter	Type	Description
action	integer	Event update action(s). Possible bitmap values: 1 - close problem; 2 - acknowledge event; 4 - add message; 8 - change severity; 16 - unacknowledge event; 32 - suppress event; 64 - unsuppress event; 128 - change event rank to cause; 256 - change event rank to symptom. This is a bitmask field; any sum of possible bitmap values is acceptable (for example, 6 for acknowledging an event and leaving a message, or 34 for acknowledging and suppressing an event). Parameter behavior: - *required*
cause_eventid	ID	Cause event ID. Parameter behavior: - *required* if `action` contains the ”change event rank to symptom” bit
message	string	Text of the message. Parameter behavior: - *required* if `action` contains the ”add message” bit
severity	integer	New severity for events. Possible values: 0 - not classified; 1 - information; 2 - warning; 3 - average; 4 - high; 5 - disaster. Parameter behavior: - *required* if `action` contains the ”change severity” bit
suppress_until	integer	Unix timestamp until which event must be suppressed. If set to ”0”, the suppression will be indefinite. Parameter behavior: - *required* if `action` contains the ”suppress event” bit

Return values

`(object)` Returns an object containing the IDs of the updated events under the `eventids` property.

Examples

Acknowledging an event

Acknowledge a single event and leave a message.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "event.acknowledge",
    "params": {
        "eventids": "20427",
        "action": 6,
```

```
        "message": "Problem resolved."
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "eventids": [
            "20427"
        ]
    },
    "id": 1
}
```

Changing event's severity

Change severity for multiple events and leave a message.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "event.acknowledge",
    "params": {
        "eventids": ["20427", "20428"],
        "action": 12,
        "message": "Maintenance required to fix it.",
        "severity": 4
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "eventids": [
            "20427",
            "20428"
        ]
    },
    "id": 1
}
```

Suppressing an event indefinitely

Suppress a single event without an expiration time.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "event.acknowledge",
    "params": {
        "eventids": "20427",
        "action": 32,
        "suppress_until": 0
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
```

```
    "result": {
        "eventids": [
            20427
        ]
    },
    "id": 1
}
```

Unsuppressing an event

Unsuppress a single previously suppressed event.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "event.acknowledge",
    "params": {
        "eventids": "20427",
        "action": 64
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "eventids": [
            20427
        ]
    },
    "id": 1
}
```

Acknowledge and suppress an event

Acknowledge one or more events, add a message, and suppress them in the same request.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "event.acknowledge",
    "params": {
        "eventids": [
            "20427",
            "20428"
        ],
        "action": 38,
        "message": "Suppressing during planned maintenance.",
        "suppress_until": 0
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "eventids": [
            20427,
            20428
        ]
    },
```

```
    "id": 1
}
```

Source

CEvent::acknowledge() in *ui/include/classes/api/services/CEvent.php*.

event.get

Description

`integer/array event.get(object parameters)`

The method allows you to retrieve events according to the given parameters.

Attention:
This method may return events of a deleted entity if these events have not been removed by the housekeeper yet.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
eventids	ID/array	Return only events with the given IDs.
groupids	ID/array	Return only events created by objects that belong to the given host groups.
hostids	ID/array	Return only events created by objects that belong to the given hosts.
objectids	ID/array	Return only events created by the given objects.
source	integer	Return only events with the given type. Refer to the event object page for a list of supported event types. Default: 0 - trigger events.
object	integer	Return only events created by objects of the given type. Refer to the event object page for a list of supported object types. Default: 0 - trigger.
acknowledged	boolean	If set to `true` return only acknowledged events.
action	integer	Return only events for which the given event update actions have been performed. For multiple actions, use a sum of any acceptable bitmap values as bitmask (for example, 34 for acknowledge and suppress event).
action_userids	ID/array	Return only events with the given IDs of users who performed the event update actions.
suppressed	boolean	If set to `true` return only suppressed events.
symptom	boolean	If set to `true` return only symptom events.
severities	integer/array	Return only events with the given event severities. Applies only if object is trigger.
trigger_severities	integer/array	Return only events with the given trigger severities. Applies only if object is trigger.
evaltype	integer	Tag evaluation method. Possible values: 0 - *(default)* And/Or; 2 - Or.

Parameter	Type	Description
tags	array	Return only events with the given tags. Format: `[{"tag": "<tag>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all events. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
eventid_from	string	Return only events with IDs greater or equal to the given ID.
eventid_till	string	Return only events with IDs less or equal to the given ID.
time_from	timestamp	Return only events created at or after the given time.
time_till	timestamp	Return only events created at or before the given time.
problem_time_from	timestamp	Returns only events that were in the problem state from `problem_time_from` regardless of their current status. Applies only when the source is a trigger event and the object is a trigger. This parameter is mandatory if `problem_time_till` is specified.
problem_time_till	timestamp	Returns only events that were in the problem state up to `problem_time_till` regardless of their current status. Applies only when the source is a trigger event and the object is a trigger. This parameter is mandatory if `problem_time_from` is specified.
value	integer/array	Return only events with the given values.
selectAcknowledges	query	Return an `acknowledges` property with event updates. Event updates are sorted in reverse chronological order. The event update object has the following properties: `acknowledgeid` - `(ID)` acknowledgment's ID; `userid` - `(ID)` ID of the user that updated the event; `clock` - `(timestamp)` time when the event was updated; `message` - `(string)` text of the message; `action` - `(integer)` update action that was performed, see `event.acknowledge`; `old_severity` - `(integer)` event severity before this update action; `new_severity` - `(integer)` event severity after this update action; `suppress_until` - `(timestamp)` time till event will be suppressed; `taskid` - `(ID)` ID of task if current event is undergoing a rank change; `maintenanceid` - `(ID)` ID of the maintenance; `username` - `(string)` username of the user that updated the event; `name` - `(string)` name of the user that updated the event; `surname` - `(string)` surname of the user that updated the event. Supports `count`.
selectAlerts	query	Return an `alerts` property with alerts generated by the event. Alerts are sorted in reverse chronological order.
selectHosts	query	Return a `hosts` property with hosts containing the object that created the event. Supported only for events generated by triggers, items or LLD rules.
selectRelatedObject	query	Return a `relatedObject` property with the object that created the event. The type of object returned depends on the event type.
selectSuppressionData	query	Return a `suppression_data` property with the list of active maintenances and manual suppressions: `maintenanceid` - `(ID)` ID of the maintenance; `userid` - `(ID)` ID of user who suppressed the event; `suppress_until` - `(integer)` time until the event is suppressed.
selectTags	query	Return a `tags` property with event tags.

Parameter	Type	Description
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type.
sortfield	string/array	Sort the result by the given properties. Possible values: `eventid`, `objectid`, `clock`. Possible values when used together with `groupBy`: `objectid`. Possible values when used together with `countOutput` and `groupBy`: `objectid`, `rowscount`.
groupBy	string/array	Group the results by the given properties. The specified properties will be returned in the results. Possible values: `objectid`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects
- the count of retrieved objects, if the `countOutput` parameter has been used, but the `groupBy` parameter has not been used
- an array of objects with aggregation results, if the `groupBy` parameter has been used

Examples

Retrieving trigger events

Retrieve the latest events from trigger "22395".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "event.get",
    "params": {
        "output": "extend",
        "selectAcknowledges": "extend",
        "selectSuppressionData": "extend",
        "selectTags": "extend",
        "objectids": "22395",
        "sortfield": ["clock", "eventid"],
        "sortorder": "DESC"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "eventid": "20",
            "source": "0",
            "object": "0",
            "objectid": "22395",
            "clock": "1728658089",
            "value": "0",
            "acknowledged": "0",
            "ns": "461809482",
            "name": "Load average is too high (per CPU load over 1.5 for 5m)",
            "severity": "0",
            "r_eventid": "0",
            "c_eventid": "0",
            "correlationid": "0",
            "userid": "0",
            "suppressed": "0",
            "cause_eventid": "0",
            "acknowledges": [],
            "opdata": "Load averages(1m 5m 15m): (1.6328 3.0522 2.2515), # of CPUs: 2",
            "suppression_data": [],
            "tags": [
                {
                    "tag": "scope",
                    "value": "capacity"
                },
                {
                    "tag": "scope",
                    "value": "performance"
                },
                {
                    "tag": "component",
                    "value": "cpu"
                },
                {
                    "tag": "class",
                    "value": "os"
                },
                {
                    "tag": "target",
                    "value": "linux"
                }
            ],
            "urls": []
        },
        {
            "eventid": "4",
            "source": "0",
            "object": "0",
            "objectid": "22395",
            "clock": "1728657737",
            "value": "1",
            "acknowledged": "1",
            "ns": "460759366",
            "name": "Load average is too high (per CPU load over 1.5 for 5m)",
            "severity": "3",
            "r_eventid": "20",
            "c_eventid": "0",
            "correlationid": "0",
            "userid": "0",
```

```
            "suppressed": "1",
            "cause_eventid": "0",
            "acknowledges": [
                {
                    "acknowledgeid": "1",
                    "userid": "1",
                    "clock": "1728657938",
                    "message": "Testing environment. Please disregard this alert.",
                    "action": "38",
                    "old_severity": "0",
                    "new_severity": "0",
                    "suppress_until": "1728744338",
                    "taskid": "0",
                    "maintenanceid": "0",
                    "username": "Admin",
                    "name": "Zabbix",
                    "surname": "Administrator"
                }
            ],
            "opdata": "Load averages(1m 5m 15m): (1.6328 3.0522 2.2515), # of CPUs: 2",
            "suppression_data": [
                {
                    "maintenanceid": "0",
                    "suppress_until": "1728744338",
                    "userid": "1"
                }
            ],
            "tags": [
                {
                    "tag": "scope",
                    "value": "capacity"
                },
                {
                    "tag": "scope",
                    "value": "performance"
                },
                {
                    "tag": "component",
                    "value": "cpu"
                },
                {
                    "tag": "class",
                    "value": "os"
                },
                {
                    "tag": "target",
                    "value": "linux"
                }
            ],
            "urls": []
        }
    ],
    "id": 1
}
```


Retrieving events by time period

Retrieve all events that have been created between October 17 and 18, 2012, in reverse chronological order.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "event.get",
```

```
    "params": {
        "output": "extend",
        "time_from": "1350432000",
        "time_till": "1350518400",
        "sortfield": ["clock", "eventid"],
        "sortorder": "DESC"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "eventid": "20617",
            "source": "0",
            "object": "0",
            "objectid": "14282",
            "clock": "1350477816",
            "value": "1",
            "acknowledged": "0",
            "ns": "0",
            "name": "Less than 25% free in the history cache",
            "severity": "3",
            "r_eventid": "0",
            "c_eventid": "0",
            "correlationid": "0",
            "userid": "0",
            "cause_eventid": "0",
            "opdata": "",
            "suppressed": "0",
            "urls": []
        },
        {
            "eventid": "20616",
            "source": "0",
            "object": "0",
            "objectid": "14281",
            "clock": "1350477814",
            "value": "0",
            "acknowledged": "0",
            "ns": "0",
            "name": "Zabbix trapper processes more than 75% busy",
            "severity": "0",
            "r_eventid": "0",
            "c_eventid": "0",
            "correlationid": "0",
            "userid": "0",
            "cause_eventid": "0",
            "opdata": "",
            "suppressed": "0",
            "urls": []
        },
        {
            "eventid": "20615",
            "source": "0",
            "object": "0",
            "objectid": "14281",
            "clock": "1350477541",
            "value": "1",
            "acknowledged": "0",
```

```
            "ns": "0",
            "name": "Zabbix trapper processes more than 75% busy",
            "severity": "3",
            "r_eventid": "20616",
            "c_eventid": "0",
            "correlationid": "0",
            "userid": "0",
            "cause_eventid": "0",
            "opdata": "",
            "suppressed": "0",
            "urls": []
        }
    ],
    "id": 1
}
```

Retrieving events acknowledged by specified user

Retrieving events acknowledged by user with ID=10

Request:

```
{
    "jsonrpc": "2.0",
    "method": "event.get",
    "params": {
        "output": "extend",
        "action": 2,
        "action_userids": [10],
        "selectAcknowledges": ["userid", "action"],
        "sortfield": ["eventid"],
        "sortorder": "DESC"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "eventid": "503",
            "source": "0",
            "object": "0",
            "objectid": "23162",
            "clock": "1747212236",
            "value": "1",
            "acknowledged": "1",
            "ns": "413470863",
            "name": "Number of installed packages has been changed",
            "severity": "2",
            "r_eventid": "0",
            "c_eventid": "0",
            "correlationid": "0",
            "userid": "0",
            "suppressed": "0",
            "cause_eventid": "0",
            "acknowledges": [
                {
                    "userid": "10",
                    "action": "2"
                }
            ],
            "opdata": "",
```

```
            "urls": []
        }
    ],
    "id": 1
}
```

Retrieving top triggers with problem event count

Retrieve the top 5 triggers that have severities ”Warning”, ”Average”, ”High”, or ”Disaster”, together with the number of problem events within a specified time period.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "event.get",
    "params": {
        "countOutput": true,
        "groupBy": "objectid",
        "source": 0,
        "object": 0,
        "value": 1,
        "time_from": 1672531200,
        "time_till": 1677628800,
        "trigger_severities": [2, 3, 4, 5],
        "sortfield": ["rowscount"],
        "sortorder": "DESC",
        "limit": 5
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "objectid": "232124",
            "rowscount": "27"
        },
        {
            "objectid": "29055",
            "rowscount": "23"
        },
        {
            "objectid": "253731",
            "rowscount": "18"
        },
        {
            "objectid": "254062",
            "rowscount": "11"
        },
        {
            "objectid": "23216",
            "rowscount": "7"
        }
    ],
    "id": 1
}
```

See also

- Alert
- Item
- Host

- LLD rule
- Service
- Trigger

Source

CEvent::get() in *ui/include/classes/api/services/CEvent.php*.

Graph

This class is designed to work with graphs.

Object references:

- Graph

Available methods:

- graph.create - create new graphs
- graph.delete - delete graphs
- graph.get - retrieve graphs
- graph.update - update graphs

Graph object

The following objects are directly related to the `graph` API.

Graph

The graph object has the following properties.

Property	Type	Description
graphid	ID	ID of the graph. Property behavior: - *read-only* - *required* for update operations
height	integer	Height of the graph in pixels. Property behavior: - *required* for create operations
name	string	Name of the graph. Property behavior: - *required* for create operations
width	integer	Width of the graph in pixels. Property behavior: - *required* for create operations
flags	integer	Origin of the graph. Possible values: 0 - *(default)* a plain graph; 4 - a graph converted from prototype. Property behavior: - *read-only*
graphtype	integer	Graph's layout type. Possible values: 0 - *(default)* normal; 1 - stacked; 2 - pie; 3 - exploded.

Property	Type	Description
percent_left	float	Left percentile. Default: 0.
percent_right	float	Right percentile. Default: 0.
show_3d	integer	Whether to show pie and exploded graphs in 3D. Possible values: 0 - *(default)* show in 2D; 1 - show in 3D.
show_legend	integer	Whether to show the legend on the graph. Possible values: 0 - hide; 1 - *(default)* show.
show_work_period	integer	Whether to show the working time on the graph. Possible values: 0 - hide; 1 - *(default)* show.
show_triggers	integer	Whether to show the trigger line on the graph. Possible values: 0 - hide; 1 - *(default)* show.
templateid	ID	ID of the parent template graph. Property behavior: - *read-only*
yaxismax	float	The fixed maximum value for the Y axis. Default: 100.
yaxismin	float	The fixed minimum value for the Y axis. Default: 0.
ymax_itemid	ID	ID of the item that is used as the maximum value for the Y axis. If a user has no access to the specified item, the graph is rendered as if `ymax_type` is set to ”calculated”.
ymax_type	integer	Maximum value calculation method for the Y axis. Possible values: 0 - *(default)* calculated; 1 - fixed; 2 - item.
ymin_itemid	ID	ID of the item that is used as the minimum value for the Y axis. If a user has no access to the specified item, the graph is rendered as if `ymin_type` is set to ”calculated”.
ymin_type	integer	Minimum value calculation method for the Y axis. Possible values: 0 - *(default)* calculated; 1 - fixed; 2 - item.
uuid	string	Universal unique identifier, used for linking imported graphs to already existing ones. Auto-generated, if not given. Property behavior: - *supported* if the graph belongs to a template

graph.create

Description

`object graph.create(object/array graphs)`

This method allows you to create new graphs.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Graphs to create.

Additionally to the standard graph properties, the method accepts the following parameters.

Parameter	Type	Description
gitems	array	Graph items to be created for the graph. Parameter behavior: *- required*

Return values

`(object)` Returns an object containing the IDs of the created graphs under the `graphids` property. The order of the returned IDs matches the order of the passed graphs.

Examples

Creating a graph

Create a graph with two items.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "graph.create",
    "params": {
        "name": "MySQL bandwidth",
        "width": 900,
        "height": 200,
        "gitems": [
            {
                "itemid": "22828",
                "color": "00AA00"
            },
            {
                "itemid": "22829",
                "color": "3333FF"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "graphids": [
            "652"
        ]
    },
```

```
    "id": 1
}
```

See also

- Graph item

Source

CGraph::create() in *ui/include/classes/api/services/CGraph.php*.

graph.delete

Description

`object graph.delete(array graphIds)`

This method allows you to delete graphs.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the graphs to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted graphs under the `graphids` property.

Examples

Deleting multiple graphs

Delete two graphs.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "graph.delete",
    "params": [
        "652",
        "653"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "graphids": [
            "652",
            "653"
        ]
    },
    "id": 1
}
```

Source

CGraph::delete() in *ui/include/classes/api/services/CGraph.php*.

graph.get

Description

`integer/array graph.get(object parameters)`

The method allows you to retrieve graphs according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
graphids	ID/array	Return only graphs with the given IDs.
groupids	ID/array	Return only graphs that belong to hosts or templates in the given host groups or template groups.
templateids	ID/array	Return only graph that belong to the given templates.
hostids	ID/array	Return only graphs that belong to the given hosts.
itemids	ID/array	Return only graphs that contain the given items.
templated	boolean	If set to `true` return only graphs that belong to templates.
inherited	boolean	If set to `true` return only graphs inherited from a template.
expandName	flag	Expand macros in the graph name.
selectDiscoveryData	query	Return a `discoveryData` property with the graph discovery object data. The graph discovery object links a discovered graph to a graph prototype from which it was discovered. It has the following properties: `parent_graphid` - `(ID)` ID of the graph prototype from which the graph has been created; `status` - `(int)` graph discovery status: 0 - *(default)* graph is discovered, 1 - graph is not discovered anymore; `ts_delete` - `(timestamp)` time when a graph that is no longer discovered will be deleted.
selectDiscoveryRule	query	Return a `discoveryRule` property with the low-level discovery rule that created the graph.
selectHostGroups	query	Return a `hostgroups` property with the host groups that the graph belongs to.
selectTemplateGroups	query	Return a `templategroups` property with the template groups that the graph belongs to.
selectTemplates	query	Return a `templates` property with the templates that the graph belongs to.
selectHosts	query	Return a `hosts` property with the hosts that the graph belongs to.
selectItems	query	Return an `items` property with the items used in the graph.
selectGraphItems	query	Return a `gitems` property with the items used in the graph.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type. Supports additional properties: `host` - technical name of the host that the graph belongs to; `hostid` - ID of the host that the graph belongs to.
sortfield	string/array	Sort the result by the given properties. Possible values: `graphid`, `name`, `graphtype`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
limit	integer	

Parameter	Type	Description
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	
selectGraphDiscovery	query	Return a `graphDiscovery` property with the graph discovery object. The graph discovery objects links the graph to a graph prototype from which it was created. This query is **deprecated**, please use `selectDiscoveryData` instead.

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving graphs from hosts

Retrieve all graphs from host "10107" and sort them by name.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "graph.get",
    "params": {
        "output": "extend",
        "hostids": 10107,
        "sortfield": "name"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "graphid": "612",
            "name": "CPU jumps",
            "width": "900",
            "height": "200",
            "yaxismin": "0",
            "yaxismax": "100",
            "templateid": "439",
            "show_work_period": "1",
            "show_triggers": "1",
            "graphtype": "0",
            "show_legend": "1",
            "show_3d": "0",
            "percent_left": "0",
            "percent_right": "0",
            "ymin_type": "0",
            "ymax_type": "0",
            "ymin_itemid": "0",
            "ymax_itemid": "0",
            "flags": "0"
```

```
        },
        {
            "graphid": "613",
            "name": "CPU load",
            "width": "900",
            "height": "200",
            "yaxismin": "0",
            "yaxismax": "100",
            "templateid": "433",
            "show_work_period": "1",
            "show_triggers": "1",
            "graphtype": "0",
            "show_legend": "1",
            "show_3d": "0",
            "percent_left": "0",
            "percent_right": "0",
            "ymin_type": "1",
            "ymax_type": "0",
            "ymin_itemid": "0",
            "ymax_itemid": "0",
            "flags": "0"
        },
        {
            "graphid": "614",
            "name": "CPU utilization",
            "width": "900",
            "height": "200",
            "yaxismin": "0",
            "yaxismax": "100",
            "templateid": "387",
            "show_work_period": "1",
            "show_triggers": "0",
            "graphtype": "1",
            "show_legend": "1",
            "show_3d": "0",
            "percent_left": "0",
            "percent_right": "0",
            "ymin_type": "1",
            "ymax_type": "1",
            "ymin_itemid": "0",
            "ymax_itemid": "0",
            "flags": "0"
        },
        {
            "graphid": "645",
            "name": "Disk space usage /",
            "width": "600",
            "height": "340",
            "yaxismin": "0",
            "yaxismax": "0",
            "templateid": "0",
            "show_work_period": "0",
            "show_triggers": "0",
            "graphtype": "2",
            "show_legend": "1",
            "show_3d": "1",
            "percent_left": "0",
            "percent_right": "0",
            "ymin_type": "0",
            "ymax_type": "0",
            "ymin_itemid": "0",
            "ymax_itemid": "0",
```

```
            "flags": "4"
        }
    ],
    "id": 1
}
```

See also

- Discovery rule
- Graph item
- Item
- Host
- Host group
- Template
- Template group

Source

CGraph::get() in *ui/include/classes/api/services/CGraph.php*.

graph.update

Description

`object graph.update(object/array graphs)`

This method allows you to update existing graphs.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Graph properties to be updated.

The `graphid` property must be defined for each graph, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard graph properties the method accepts the following parameters.

Parameter	Type	Description
gitems	array	Graph items to replace existing graph items. If a graph item has the `gitemid` property defined it will be updated; otherwise, a new graph item will be created.

Return values

`(object)` Returns an object containing the IDs of the updated graphs under the `graphids` property.

Examples

Setting the maximum for the Y scale

Set the maximum of the Y scale to a fixed value of 100.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "graph.update",
    "params": {
        "graphid": "439",
        "ymax_type": 1,
        "yaxismax": 100
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "graphids": [
            "439"
        ]
    },
    "id": 1
}
```

Source

CGraph::update() in *ui/include/classes/api/services/CGraph.php*.

Graph item

This class is designed to work with graph items.

Object references:

- Graph item

Available methods:

- graphitem.get - retrieve graph items

Graph item object

The following objects are directly related to the `graphitem` API.

Graph item

Note:
Graph items can only be modified via the `graph` API.

The graph item object has the following properties.

Property	Type	Description
gitemid	ID	ID of the graph item. Property behavior: - *read-only*
color	string	Graph item's draw color as a hexadecimal color code. Property behavior: - *required* for create operations
itemid	ID	ID of the item. Property behavior: - *required* for create operations
calc_fnc	integer	Value of the item that will be displayed. Possible values: 1 - minimum value; 2 - *(default)* average value; 4 - maximum value; 7 - all values; 9 - last value, used only by pie and exploded graphs.

Property	Type	Description
drawtype	integer	Draw style of the graph item. Possible values: 0 - *(default)* line; 1 - filled region; 2 - bold line; 3 - dot; 4 - dashed line; 5 - gradient line.
graphid	ID	ID of the graph that the graph item belongs to.
sortorder	integer	Position of the item in the graph. Default: starts with "0" and increases by one with each entry.
type	integer	Type of graph item. Possible values: 0 - *(default)* simple; 2 - graph sum, used only by pie and exploded graphs.
yaxisside	integer	Side of the graph where the graph item's Y scale will be drawn. Possible values: 0 - *(default)* left side; 1 - right side.

graphitem.get

Description

`integer/array graphitem.get(object parameters)`

The method allows you to retrieve graph items according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
graphids	ID/array	Return only graph items that belong to the given graphs.
itemids	ID/array	Return only graph items with the given item IDs.
type	integer	Return only graph items with the given type. Refer to the graph item object page for a list of supported graph item types.
selectGraphs	query	Return a `graphs` property with an array of graphs that the item belongs to.
sortfield	string/array	Sort the result by the given properties. Possible values: `gitemid`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
limit	integer	
output	query	
preservekeys	boolean	
sortorder	string/array	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving graph items from a graph

Retrieve all graph items used in a graph with additional information about the item and the host.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "graphitem.get",
    "params": {
        "output": "extend",
        "graphids": "387"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "gitemid": "1242",
            "graphid": "387",
            "itemid": "22665",
            "drawtype": "1",
            "sortorder": "1",
            "color": "FF5555",
            "yaxisside": "0",
            "calc_fnc": "2",
            "type": "0"
        },
        {
            "gitemid": "1243",
            "graphid": "387",
            "itemid": "22668",
            "drawtype": "1",
            "sortorder": "2",
            "color": "55FF55",
            "yaxisside": "0",
            "calc_fnc": "2",
            "type": "0"
        },
        {
            "gitemid": "1244",
            "graphid": "387",
            "itemid": "22671",
            "drawtype": "1",
            "sortorder": "3",
            "color": "009999",
            "yaxisside": "0",
            "calc_fnc": "2",
            "type": "0"
        }
    ],
    "id": 1
}
```

See also

- Graph

Source

CGraphItem::get() in *ui/include/classes/api/services/CGraphItem.php*.

Graph prototype

This class is designed to work with graph prototypes.

Object references:

- Graph prototype

Available methods:

- graphprototype.create - create new graph prototypes
- graphprototype.delete - delete graph prototypes
- graphprototype.get - retrieve graph prototypes
- graphprototype.update - update graph prototypes

Graph prototype object

The following objects are directly related to the `graphprototype` API.

Graph prototype

The graph prototype object has the following properties.

Property	Type	Description
graphid	ID	ID of the graph prototype. Property behavior: - *read-only* - *required* for update operations
height	integer	Height of the graph prototype in pixels. Property behavior: - *required* for create operations
name	string	Name of the graph prototype. Property behavior: - *required* for create operations
width	integer	Width of the graph prototype in pixels. Property behavior: - *required* for create operations
flags	integer	Origin of the graph prototype. Possible values: 2 - a graph prototype; 6 - a discovered graph prototype Property behavior: - *read-only*
graphtype	integer	Graph prototypes's layout type. Possible values: 0 - *(default)* normal; 1 - stacked; 2 - pie; 3 - exploded.
percent_left	float	Left percentile. Default: 0.

Property	Type	Description
percent_right	float	Right percentile. Default: 0.
show_3d	integer	Whether to show discovered pie and exploded graphs in 3D. Possible values: 0 - *(default)* show in 2D; 1 - show in 3D.
show_legend	integer	Whether to show the legend on the discovered graph. Possible values: 0 - hide; 1 - *(default)* show.
show_work_period	integer	Whether to show the working time on the discovered graph. Possible values: 0 - hide; 1 - *(default)* show.
templateid	ID	ID of the parent template graph prototype. Property behavior: - *read-only*
yaxismax	float	The fixed maximum value for the Y axis.
yaxismin	float	The fixed minimum value for the Y axis.
ymax_itemid	ID	ID of the item that is used as the maximum value for the Y axis. If a user has no access to the specified item, the graph is rendered as if `ymax_type` is set to ”calculated”.
ymax_type	integer	Maximum value calculation method for the Y axis. Possible values: 0 - *(default)* calculated; 1 - fixed; 2 - item.
ymin_itemid	ID	ID of the item that is used as the minimum value for the Y axis. If a user has no access to the specified item, the graph is rendered as if `ymin_type` is set to ”calculated”.
ymin_type	integer	Minimum value calculation method for the Y axis. Possible values: 0 - *(default)* calculated; 1 - fixed; 2 - item.
discover	integer	Graph prototype discovery status. Possible values: 0 - *(default)* new graphs will be discovered; 1 - new graphs will not be discovered and existing graphs will be marked as lost.
uuid	string	Universal unique identifier, used for linking imported graph prototypes to already existing ones. Auto-generated, if not given. Property behavior: - *supported* if the graph prototype belongs to a template

graphprototype.create

Description

`object graphprototype.create(object/array graphPrototypes)`

This method allows you to create new graph prototypes.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Graph prototypes to create.

Additionally to the standard graph prototype properties, the method accepts the following parameters.

Parameter	Type	Description
gitems	array	Graph items to be created for the graph prototypes. Graph items can reference both items and item prototypes, but at least one item prototype must be present. Parameter behavior: - *required*

Return values

`(object)` Returns an object containing the IDs of the created graph prototypes under the `graphids` property. The order of the returned IDs matches the order of the passed graph prototypes.

Examples

Creating a graph prototype

Create a graph prototype with two items.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "graphprototype.create",
    "params": {
        "name": "Disk space usage {#FSNAME}",
        "width": 900,
        "height": 200,
        "gitems": [
            {
                "itemid": "22828",
                "color": "00AA00"
            },
            {
                "itemid": "22829",
                "color": "3333FF"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "graphids": [
            "652"
        ]
    },
    "id": 1
}
```

See also

- Graph item

Source

CGraphPrototype::create() in *ui/include/classes/api/services/CGraphPrototype.php*.

graphprototype.delete

Description

`object graphprototype.delete(array graphPrototypeIds)`

This method allows you to delete graph prototypes.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the graph prototypes to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted graph prototypes under the `graphids` property.

Examples

Deleting multiple graph prototypes

Delete two graph prototypes.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "graphprototype.delete",
    "params": [
        "652",
        "653"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "graphids": [
            "652",
            "653"
        ]
    },
    "id": 1
}
```

Source

CGraphPrototype::delete() in *ui/include/classes/api/services/CGraphPrototype.php*.

graphprototype.get

Description

`integer/array graphprototype.get(object parameters)`

The method allows you to retrieve graph prototypes according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
discoveryids	ID/array	Return only graph prototypes that belong to the given discovery rules.
graphids	ID/array	Return only graph prototypes with the given IDs.
groupids	ID/array	Return only graph prototypes that belong to hosts or templates in the given host groups or template groups.
hostids	ID/array	Return only graph prototypes that belong to the given hosts.
inherited	boolean	If set to `true` return only graph prototypes inherited from a template.
itemids	ID/array	Return only graph prototypes that contain the given item prototypes.
templated	boolean	If set to `true` return only graph prototypes that belong to templates.
templateids	ID/array	Return only graph prototypes that belong to the given templates.
selectDiscoveryData	query	Return a `discoveryData` property with the graph prototype discovery object data. The graph prototype discovery object links a discovered graph prototype to a graph prototype from which it was discovered. It has the following properties: `parent_graphid` - `(ID)` ID of the graph prototype from which the graph has been created; `status` - `(int)` graph discovery status: 0 - *(default)* graph prototype is discovered, 1 - graph prototype is not discovered anymore; `ts_delete` - `(timestamp)` time when a graph prototype that is no longer discovered will be deleted.
selectDiscoveryRule	query	Return a `discoveryRule` property with the LLD rule that the graph prototype belongs to.
selectDiscoveryRulePrototype	query	Return a `discoveryRulePrototype` property with the parent LLD rule prototype that the graph prototype belongs to.
selectGraphItems	query	Return a `gitems` property with the graph items used in the graph prototype.
selectHostGroups	query	Return a `hostgroups` property with the host groups that the graph prototype belongs to.
selectHosts	query	Return a `hosts` property with the hosts that the graph prototype belongs to.
selectItems	query	Return an `items` property with the items and item prototypes used in the graph prototype.
selectTemplateGroups	query	Return a `templategroups` property with the template groups that the graph prototype belongs to.
selectTemplates	query	Return a `templates` property with the templates that the graph prototype belongs to.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type. Supports additional properties: `host` - technical name of the host that the graph prototype belongs to; `hostid` - ID of the host that the graph prototype belongs to.
sortfield	string/array	Sort the result by the given properties. Possible values: `graphid`, `name`, `graphtype`, `discovered`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	

Parameter	Type	Description
excludeSearch	boolean	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving graph prototypes from an LLD rule

Retrieve all graph prototypes from an LLD rule.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "graphprototype.get",
    "params": {
        "output": "extend",
        "discoveryids": "27426"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "graphid": "1017",
            "name": "Disk space usage {#FSNAME}",
            "width": "600",
            "height": "340",
            "yaxismin": "0.0000",
            "yaxismax": "0.0000",
            "templateid": "442",
            "show_work_period": "0",
            "show_triggers": "0",
            "graphtype": "2",
            "show_legend": "1",
            "show_3d": "1",
            "percent_left": "0.0000",
            "percent_right": "0.0000",
            "ymin_type": "0",
            "ymax_type": "0",
            "ymin_itemid": "0",
            "ymax_itemid": "0",
            "flags": "2"
            "discover": "0"
        }
    ],
    "id": 1
}
```

See also

- Discovery rule
- Graph item
- Item
- Host
- Host group
- Template
- Template group

Source

CGraphPrototype::get() in *ui/include/classes/api/services/CGraphPrototype.php*.

graphprototype.update

Description

`object graphprototype.update(object/array graphPrototypes)`

This method allows you to update existing graph prototypes.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Graph prototype properties to be updated.

The `graphid` property must be defined for each graph prototype, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard graph prototype properties, the method accepts the following parameters.

Parameter	Type	Description
gitems	array	Graph items to replace existing graph items. If a graph item has the `gitemid` property defined it will be updated; otherwise, a new graph item will be created.

Return values

`(object)` Returns an object containing the IDs of the updated graph prototypes under the `graphids` property.

Examples

Changing the size of a graph prototype

Change the size of a graph prototype to 1100 to 400 pixels.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "graphprototype.update",
    "params": {
        "graphid": "439",
        "width": 1100,
        "height": 400
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
```

```
        "graphids": [
            "439"
        ]
    },
    "id": 1
}
```

Source

CGraphPrototype::update() in *ui/include/classes/api/services/CGraphPrototype.php*.

High availability node

This class is designed to work with server nodes that are part of a high availability cluster or a standalone server instance.

Object references:

- High availability node

Available methods:

- hanode.get - retrieve nodes

High availability node object

The following object is related to operating a High availability cluster of Zabbix servers.

High availability node

Note:
Nodes are created by the Zabbix server and cannot be modified via the API.

The High availability node object has the following properties.

Property	Type	Description
ha_nodeid	ID	ID of the node.
name	string	Name assigned to the node, using the HANodeName configuration entry of `zabbix_server.conf`. Empty for a server running in standalone mode.
address	string	IP or DNS name where the node connects from.
port	integer	Port on which the node is running.
lastaccess	integer	Heartbeat time, that is, time of last update from the node. UTC timestamp.
status	integer	State of the node. Possible values: 0 - standby; 1 - stopped manually; 2 - unavailable; 3 - active.

hanode.get

Description

`integer/array hanode.get(object parameters)`

The method allows you to retrieve a list of High availability cluster nodes according to the given parameters.

Note:
This method is only available to *Super admin* user types. See User roles for more information.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
ha_nodeids	ID/array	Return only nodes with the given node IDs.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Supports properties: name, address, status.
sortfield	string/array	Sort the result by the given properties. Possible values: name, lastaccess, status.
countOutput	flag	These parameters are described in the reference commentary.
limit	integer	
output	query	
preservekeys	boolean	
sortorder	string/array	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Get a list of nodes ordered by status

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hanode.get",
    "params": {
        "preservekeys": true,
        "sortfield": "status",
        "sortorder": "DESC"
    },
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "ckuo7i1nw000h0sajj3l3hh8u": {
      "ha_nodeid": "ckuo7i1nw000h0sajj3l3hh8u",
      "name": "node-active",
      "address": "192.168.1.13",
      "port": "10051",
      "lastaccess": "1635335704",
      "status": "3"
    },
    "ckuo7i1nw000e0sajwfttc1mp": {
      "ha_nodeid": "ckuo7i1nw000e0sajwfttc1mp",
      "name": "node6",
      "address": "192.168.1.10",
      "port": "10053",
      "lastaccess": "1635332902",
      "status": "2"
    },
```

```
    "ckuo7i1nv000c0sajz85xcrtt": {
      "ha_nodeid": "ckuo7i1nv000c0sajz85xcrtt",
      "name": "node4",
      "address": "192.168.1.8",
      "port": "10052",
      "lastaccess": "1635334214",
      "status": "1"
    },
    "ckuo7i1nv000a0saj1fcdkeu4": {
      "ha_nodeid": "ckuo7i1nv000a0saj1fcdkeu4",
      "name": "node2",
      "address": "192.168.1.6",
      "port": "10051",
      "lastaccess": "1635335705",
      "status": "0"
    }
  },
  "id": 1
}
```

Get a list of specific nodes by their IDs

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hanode.get",
    "params": {
        "ha_nodeids": ["ckuo7i1nw000e0sajwfttc1mp", "ckuo7i1nv000c0sajz85xcrtt"]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "ha_nodeid": "ckuo7i1nv000c0sajz85xcrtt",
            "name": "node4",
            "address": "192.168.1.8",
            "port": "10052",
            "lastaccess": "1635334214",
            "status": "1"
        },
        {
            "ha_nodeid": "ckuo7i1nw000e0sajwfttc1mp",
            "name": "node6",
            "address": "192.168.1.10",
            "port": "10053",
            "lastaccess": "1635332902",
            "status": "2"
        }
    ],
    "id": 1
}
```

Get a list of stopped nodes

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hanode.get",
    "params": {
```

```
        "output": ["ha_nodeid", "address", "port"],
        "filter": {
            "status": 1
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "ha_nodeid": "ckuo7i1nw000g0sajjsjre7e3",
            "address": "192.168.1.12",
            "port": "10051"
        },
        {
            "ha_nodeid": "ckuo7i1nv000c0sajz85xcrtt",
            "address": "192.168.1.8",
            "port": "10052"
        },
        {
            "ha_nodeid": "ckuo7i1nv000d0sajd95y1b6x",
            "address": "192.168.1.9",
            "port": "10053"
        }
    ],
    "id": 1
}
```

Get a count of standby nodes

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hanode.get",
    "params": {
        "countOutput": true,
        "filter": {
            "status": 0
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": "3",
    "id": 1
}
```

Check status of nodes at specific IP addresses

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hanode.get",
    "params": {
        "output": ["name", "status"],
        "filter": {
            "address": ["192.168.1.7", "192.168.1.13"]
```

```
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "name": "node3",
            "status": "0"
        },
        {
            "name": "node-active",
            "status": "3"
        }
    ],
    "id": 1
}
```

Source

CHaNode::get() in *ui/include/classes/api/services/CHaNode.php*.

History

This class is designed to work with history data.

Object references:

- Float history
- Integer history
- String history
- Text history
- Log history

Available methods:

- history.clear - clear history data
- history.get - retrieve history data
- history.push - send history data to Zabbix server

History object

The following objects are directly related to the `history` API.

Note:
History objects differ depending on the item's type of information. They are created by Zabbix server and cannot be modified via the API.

Float history

The float history object has the following properties.

Property	Type	Description
clock	timestamp	Time when that value was received.
itemid	ID	ID of the related item.
ns	integer	Nanoseconds when the value was received.
value	float	Received value.

Integer history

The integer history object has the following properties.

Property	Type	Description
clock	timestamp	Time when that value was received.
itemid	ID	ID of the related item.
ns	integer	Nanoseconds when the value was received.
value	integer	Received value.

String history

The string history object has the following properties.

Property	Type	Description
clock	timestamp	Time when that value was received.
itemid	ID	ID of the related item.
ns	integer	Nanoseconds when the value was received.
value	string	Received value.

Text history

The text history object has the following properties.

Property	Type	Description
id	ID	ID of the history entry.
clock	timestamp	Time when that value was received.
itemid	ID	ID of the related item.
ns	integer	Nanoseconds when the value was received.
value	text	Received value.

Log history

The log history object has the following properties.

Property	Type	Description
id	ID	ID of the history entry.
clock	timestamp	Time when that value was received.
itemid	ID	ID of the related item.
logeventid	integer	Windows event log entry ID.
ns	integer	Nanoseconds when the value was received.
severity	integer	Windows event log entry level.
source	string	Windows event log entry source.
timestamp	timestamp	Windows event log entry time.
value	text	Received value.

Binary history

The binary history object has the following properties.

Property	Type	Description
clock	timestamp	Time when that value was received.
itemid	ID	ID of the related item.
ns	integer	Nanoseconds when the value was received.
value	text	Received value.

JSON history

The JSON history object has the following properties.

Property	Type	Description
clock	timestamp	Time when that value was received.
itemid	ID	ID of the related item.
ns	integer	Nanoseconds when the value was received.
value	text	Received value.

history.clear

Description

`object history.clear(array itemids)`

This method allows you to clear item history.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of items to clear.

Return values

`(object)` Returns an object containing the IDs of the cleared items under the `itemids` property.

Examples

Clear history

Request:

```
{
    "jsonrpc": "2.0",
    "method": "history.clear",
    "params": [
        "10325",
        "13205"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "10325",
            "13205"
        ]
    },
    "id": 1
}
```

Source

CHistory::clear() in *ui/include/classes/api/services/CHistory.php*.

history.get

Description

`integer/array history.get(object parameters)`

The method allows you to retrieve history data according to the given parameters.

Attention:
This method may return historical data of a deleted entity if this data has not been removed by the housekeeper yet.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
history	integer	History object types to return. Possible values: 0 - numeric float; 1 - character; 2 - log; 3 - *(default)* numeric unsigned; 4 - text; 5 - binary; 6 - JSON.
hostids	ID/array	Return only history from the given hosts.
itemids	ID/array	Return only history from the given items.
maxValueSize	integer	Maximum number of bytes to return in the `value` property. Possible values range: 1-134217728 (1 byte to 128MiB). If set to `null`, no limit is applied. Default: 65536 (64KiB). Parameter behavior: - *supported* if `history` is set to "binary" or "JSON"
time_from	timestamp	Return only values that have been received after or at the given time.
time_till	timestamp	Return only values that have been received before or at the given time.
sortfield	string/array	Sort the result by the given properties. Possible values: `itemid`, `clock`, `ns`. Does not support sorting by `ns` if history data is stored in ClickHouse.
search	object	Return results that match the given pattern (case-insensitive). Accepts an object, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a `LIKE "%…%"` search. Does not support searching by `value` if `history` is set to 6 (JSON).
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving item history data

Return 10 latest values received from a numeric(float) item.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "history.get",
    "params": {
        "output": "extend",
        "history": 0,
        "itemids": "23296",
        "sortfield": "clock",
        "sortorder": "DESC",
        "limit": 10
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "itemid": "23296",
            "clock": "1351090996",
            "value": "0.085",
            "ns": "563157632"
        },
        {
            "itemid": "23296",
            "clock": "1351090936",
            "value": "0.16",
            "ns": "549216402"
        },
        {
            "itemid": "23296",
            "clock": "1351090876",
            "value": "0.18",
            "ns": "537418114"
        },
        {
            "itemid": "23296",
            "clock": "1351090816",
            "value": "0.21",
            "ns": "522659528"
        },
        {
            "itemid": "23296",
            "clock": "1351090756",
            "value": "0.215",
            "ns": "507809457"
        },
        {
            "itemid": "23296",
            "clock": "1351090696",
            "value": "0.255",
            "ns": "495509699"
```

```
        },
        {
            "itemid": "23296",
            "clock": "1351090636",
            "value": "0.36",
            "ns": "477708209"
        },
        {
            "itemid": "23296",
            "clock": "1351090576",
            "value": "0.375",
            "ns": "463251343"
        },
        {
            "itemid": "23296",
            "clock": "1351090516",
            "value": "0.315",
            "ns": "447947017"
        },
        {
            "itemid": "23296",
            "clock": "1351090456",
            "value": "0.275",
            "ns": "435307141"
        }
    ],
    "id": 1
}
```


Source

CHistory::get() in *ui/include/classes/api/services/CHistory.php*.

history.push

Description

`object history.push(object/array itemHistoryData)`

This method allows sending item history data to Zabbix server.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Item history data to send.

The method supports the following parameters.

Parameter	Type	Description
itemid	ID	ID of the related item. Parameter behavior: - *required* if `host` and `key` are not set
host	string	Technical name of the host. Parameter behavior: - *required* if `itemid` is not set
key	string	Item key. Parameter behavior: - *required* if `itemid` is not set

Parameter	Type	Description
value	mixed	Item value. Parameter behavior: - *required*
clock	timestamp	Time when the value was received.
ns	integer	Nanoseconds when the value was received.

Return values

`(object)` Returns the result of the data sending operation.

Examples

Send item history data

Send item history data to Zabbix server for items ”10600”, ”10601”, and ”999999”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "history.push",
    "params": [
        {
            "itemid": 10600,
            "value": 0.5,
            "clock": 1690891294,
            "ns": 45440940
        },
        {
            "itemid": 10600,
            "value": 0.6,
            "clock": 1690891295,
            "ns": 312431
        },
        {
            "itemid": 10601,
            "value": "[Tue Aug 01 15:01:35 2023] [error] [client 1.2.3.4] File does not exist: /var/www/ht
        },
        {
            "itemid": 999999,
            "value": 123
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "response": "success",
        "data": [
            {
                "itemid": "10600"
            },
            {
                "itemid": "10600"
            },
            {
                "itemid": "10601",
                "error": "Item is disabled."
            },
```

```
            {
                "error": "No permissions to referred object or it does not exist."
            }
        ]
    },
    "id": 1
}
```

See also

- Trapper items
- HTTP agent items
- Host
- Item

Source

CHistory::push() in *ui/include/classes/api/services/CHistory.php*.

Host

This class is designed to work with hosts.

Object references:

- Host
- Host inventory
- Host tag

Available methods:

- host.create - create new hosts
- host.delete - delete hosts
- host.get - retrieve hosts
- host.massadd - add related objects to hosts
- host.massremove - remove related objects from hosts
- host.update - update hosts

Host object

The following objects are directly related to the host API.

Host

The host object has the following properties.

Property	Type	Description
hostid	ID	ID of the host. Property behavior: - *read-only* - *required* for update operations
host	string	Technical name of the host. Property behavior: - *required* for create operations
description	text	Description of the host.
flags	integer	Origin of the host. Possible values: 0 - a plain host; 4 - a host converted from prototype. Property behavior: - *read-only*

Property	Type	Description
inventory_mode	integer	Host inventory population mode. Possible values: -1 - *(default)* disabled; 0 - manual; 1 - automatic.
ipmi_authtype	integer	IPMI authentication algorithm. Possible values: -1 - *(default)* default; 0 - none; 1 - MD2; 2 - MD5 4 - straight; 5 - OEM; 6 - RMCP+.
ipmi_password	string	IPMI password.
ipmi_privilege	integer	IPMI privilege level. Possible values: 1 - callback; 2 - *(default)* user; 3 - operator; 4 - admin; 5 - OEM.
ipmi_username	string	IPMI username.
maintenance_from	timestamp	Starting time of the effective maintenance. Property behavior: - *read-only*
maintenance_status	integer	Effective maintenance status. Possible values: 0 - *(default)* no maintenance; 1 - maintenance in effect. Property behavior: - *read-only*
maintenance_type	integer	Effective maintenance type. Possible values: 0 - *(default)* maintenance with data collection; 1 - maintenance without data collection. Property behavior: - *read-only*
maintenanceid	ID	ID of the maintenance that is currently in effect on the host. Property behavior: - *read-only*
name	string	Visible name of the host. Default: `host` property value.
monitored_by	integer	Source that is used to monitor the host. Possible values: 0 - *(default)* Zabbix server; 1 - Proxy; 2 - Proxy group. Property behavior: - *required* if `select_server_for_monitoring` is set to ”Disabled”.

Property	Type	Description
proxyid	ID	ID of the proxy that is used to monitor the host. Property behavior: - *required* if `monitored_by` is set to ”Proxy”
proxy_groupid	ID	ID of the proxy group that is used to monitor the host. Property behavior: - *required* if `monitored_by` is set to ”Proxy group”
status	integer	Status and function of the host. Possible values: 0 - *(default)* enabled; 1 - disabled.
tls_connect	integer	Connections to host. Possible values: 1 - *(default)* No encryption; 2 - PSK; 4 - certificate.
tls_accept	integer	Connections from host. Possible bitmap values: 1 - *(default)* No encryption; 2 - PSK; 4 - certificate. This is a bitmask field; any sum of possible bitmap values is acceptable (for example, 6 for PSK and certificate).
tls_issuer	string	Certificate issuer.
tls_subject	string	Certificate subject.
tls_psk_identity	string	PSK identity; must be paired with only one PSK (across autoregistration, hosts, and proxies). Do not include sensitive information in the PSK identity, as it is sent unencrypted over the network to inform the receiver which PSK to use. Property behavior: - *write-only* - *required* if `tls_connect` is set to ”PSK”, or `tls_accept` contains the ”PSK” bit
tls_psk	string	Pre-shared key (PSK); must be at least 32 hex digits. Property behavior: - *write-only* - *required* if `tls_connect` is set to ”PSK”, or `tls_accept` contains the ”PSK” bit
active_available	integer	Host active interface availability status. Possible values: 0 - interface status is unknown; 1 - interface is available; 2 - interface is not available. Property behavior: - *read-only*
assigned_proxyid	ID	ID of the proxy assigned by Zabbix server, if the host is monitored by a proxy group. Property behavior: - *read-only*

Host inventory

The host inventory object has the following properties.

Note:
Each property has it's own unique ID number, which is used to associate host inventory fields with items.

ID	Property	Type	Description	Maximum length
4	alias	string	Alias.	128 characters
11	asset_tag	string	Asset tag.	64 characters
28	chassis	string	Chassis.	64 characters
23	contact	string	Contact person.	65535 characters for SQL databases
32	contract_number	string	Contract number.	64 characters
47	date_hw_decomm	string	HW decommissioning date.	64 characters
46	date_hw_expiry	string	HW maintenance expiry date.	64 characters
45	date_hw_install	string	HW installation date.	64 characters
44	date_hw_purchase	string	HW purchase date.	64 characters
34	deployment_status	string	Deployment status.	64 characters
14	hardware	string	Hardware.	255 characters
15	hardware_full	string	Detailed hardware.	65535 characters for SQL databases
39	host_netmask	string	Host subnet mask.	39 characters
38	host_networks	string	Host networks.	65535 characters for SQL databases
40	host_router	string	Host router.	39 characters
30	hw_arch	string	HW architecture.	32 characters
33	installer_name	string	Installer name.	64 characters
24	location	string	Location.	65535 characters for SQL databases
25	location_lat	string	Location latitude.	16 characters
26	location_lon	string	Location longitude.	16 characters
12	macaddress_a	string	MAC address A.	64 characters
13	macaddress_b	string	MAC address B.	64 characters
29	model	string	Model.	64 characters
3	name	string	Name.	128 characters
27	notes	string	Notes.	65535 characters for SQL databases
41	oob_ip	string	OOB IP address.	39 characters
42	oob_netmask	string	OOB host subnet mask.	39 characters
43	oob_router	string	OOB router.	39 characters
5	os	string	OS name.	128 characters
6	os_full	string	Detailed OS name.	255 characters
7	os_short	string	Short OS name.	128 characters
61	poc_1_cell	string	Primary POC mobile number.	64 characters
58	poc_1_email	string	Primary email.	128 characters
57	poc_1_name	string	Primary POC name.	128 characters
63	poc_1_notes	string	Primary POC notes.	65535 characters for SQL databases
59	poc_1_phone_a	string	Primary POC phone A.	64 characters
60	poc_1_phone_b	string	Primary POC phone B.	64 characters
62	poc_1_screen	string	Primary POC screen name.	64 characters
68	poc_2_cell	string	Secondary POC mobile number.	64 characters
65	poc_2_email	string	Secondary POC email.	128 characters
64	poc_2_name	string	Secondary POC name.	128 characters
70	poc_2_notes	string	Secondary POC notes.	65535 characters for SQL databases
66	poc_2_phone_a	string	Secondary POC phone A.	64 characters
67	poc_2_phone_b	string	Secondary POC phone B.	64 characters
69	poc_2_screen	string	Secondary POC screen name.	64 characters
8	serialno_a	string	Serial number A.	64 characters
9	serialno_b	string	Serial number B.	64 characters
48	site_address_a	string	Site address A.	128 characters
49	site_address_b	string	Site address B.	128 characters
50	site_address_c	string	Site address C.	128 characters
51	site_city	string	Site city.	128 characters
53	site_country	string	Site country.	64 characters
56	site_notes	string	Site notes.	65535 characters for SQL databases
55	site_rack	string	Site rack location.	128 characters

ID	Property	Type	Description	Maximum length
52	site_state	string	Site state.	64 characters
54	site_zip	string	Site ZIP/postal code.	64 characters
16	software	string	Software.	255 characters
18	software_app_a	string	Software application A.	64 characters
19	software_app_b	string	Software application B.	64 characters
20	software_app_c	string	Software application C.	64 characters
21	software_app_d	string	Software application D.	64 characters
22	software_app_e	string	Software application E.	64 characters
17	software_full	string	Software details.	65535 characters for SQL databases
10	tag	string	Tag.	64 characters
1	type	string	Type.	64 characters
2	type_full	string	Type details.	64 characters
35	url_a	string	URL A.	2048 characters
36	url_b	string	URL B.	2048 characters
37	url_c	string	URL C.	2048 characters
31	vendor	string	Vendor.	64 characters

Host tag

The host tag object has the following properties.

Property	Type	Description
tag	string	Host tag name. Property behavior: - *required*
value	string	Host tag value. Property behavior: - *read-only*.
automatic	integer	Type of host tag. Possible values: 0 - *(default)* manual (tag created by user); 1 - automatic (tag created by low-level discovery) Property behavior: - *read-only*
object	integer	Type of object the tag was inherited from. Possible values: 0 - template. Property behavior: - *read-only*.
objectid	ID	ID of the object the tag was inherited from. Property behavior: - *read-only*.

host.create

Description

`object host.create(object/array hosts)`

This method allows you to create new hosts.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Hosts to create.

Additionally to the standard host properties, the method accepts the following parameters.

Parameter	Type	Description
groups	object/array	Host groups to add the host to. The host groups must have only the `groupid` property defined. Parameter behavior: - *required*
interfaces	object/array	Interfaces to be created for the host.
tags	object/array	Host tags.
templates	object/array	Templates to be linked to the host. The templates must have only the `templateid` property defined.
macros	object/array	User macros to be created for the host.
inventory	object	Host inventory properties.

Return values

`(object)` Returns an object containing the IDs of the created hosts under the `hostids` property. The order of the returned IDs matches the order of the passed hosts.

Examples

Creating a host

Create a host called "Linux server" with an IP interface and tags, add it to a group, link a template to it and set the MAC addresses in the host inventory.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.create",
    "params": {
        "host": "Linux server",
        "interfaces": [
            {
                "type": 1,
                "main": 1,
                "useip": 1,
                "ip": "192.168.3.1",
                "dns": "",
                "port": "10050"
            }
        ],
        "groups": [
            {
                "groupid": "50"
            }
        ],
        "tags": [
            {
                "tag": "host-name",
                "value": "linux-server"
            }
        ],
```

```
        "templates": [
            {
                "templateid": "20045"
            }
        ],
        "macros": [
            {
                "macro": "{$USER_ID}",
                "value": "123321"
            },
            {
                "macro": "{$USER_LOCATION}",
                "value": "0:0:0",
                "description": "latitude, longitude and altitude coordinates"
            }
        ],
        "inventory_mode": 0,
        "inventory": {
            "macaddress_a": "01234",
            "macaddress_b": "56768"
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "107819"
        ]
    },
    "id": 1
}
```

Creating a host with SNMP interface

Create a host called "SNMP host" with an SNMPv3 interface with details.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.create",
    "params": {
        "host": "SNMP host",
        "interfaces": [
            {
                "type": 2,
                "main": 1,
                "useip": 1,
                "ip": "127.0.0.1",
                "dns": "",
                "port": "161",
                "details": {
                    "version": 3,
                    "bulk": 0,
                    "securityname": "mysecurityname",
                    "contextname": "",
                    "securitylevel": 1
                }
            }
        ],
```

```
        "groups": [
            {
                "groupid": "4"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10658"
        ]
    },
    "id": 1
}
```

Creating a host with PSK encryption

Create a host called "PSK host" with PSK encryption configured. Note that the host has to be pre-configured to use PSK.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.create",
    "params": {
        "host": "PSK host",
        "interfaces": [
            {
                "type": 1,
                "ip": "192.168.3.1",
                "dns": "",
                "port": "10050",
                "useip": 1,
                "main": 1
            }
        ],
        "groups": [
            {
                "groupid": "2"
            }
        ],
        "tls_accept": 2,
        "tls_connect": 2,
        "tls_psk_identity": "PSK 001",
        "tls_psk": "1f87b595725ac58dd977beef14b97461a7c1045b9a1c963065002c5473194952"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10590"
        ]
    },
    "id": 1
}
```

Creating a host monitored by a proxy

Create a host that will be monitored by proxy with ID ”1”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.create",
    "params": {
        "host": "Host monitored by proxy",
        "groups": [
            {
                "groupid": "2"
            }
        ],
        "monitored_by": 1,
        "proxyid": 1
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10591"
        ]
    },
    "id": 1
}
```

Creating a host monitored by a proxy group

Create a host that will be monitored by proxy group with ID ”1”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.create",
    "params": {
        "host": "Host monitored by proxy group",
        "groups": [
            {
                "groupid": "2"
            }
        ],
        "monitored_by": 2,
        "proxy_groupid": 1
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10592"
        ]
    },
    "id": 1
}
```

See also

- Host group
- Template
- User macro
- Host interface
- Host inventory
- Host tag
- Proxy
- Proxy group

Source

CHost::create() in *ui/include/classes/api/services/CHost.php*.

host.delete

Description

`object host.delete(array hosts)`

This method allows you to delete hosts.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of hosts to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted hosts under the `hostids` property.

Examples

Deleting multiple hosts

Delete two hosts.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.delete",
    "params": [
        "13",
        "32"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "13",
            "32"
        ]
    },
    "id": 1
}
```

Source

CHost::delete() in *ui/include/classes/api/services/CHost.php*.

host.get

Description

`integer/array host.get(object parameters)`

The method allows you to retrieve hosts according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
groupids	ID/array	Return only hosts that belong to the given groups.
dserviceids	ID/array	Return only hosts that are related to the given discovered services.
graphids	ID/array	Return only hosts that have the given graphs.
hostids	ID/array	Return only hosts with the given host IDs.
httptestids	ID/array	Return only hosts that have the given web checks.
interfaceids	ID/array	Return only hosts that use the given interfaces.
itemids	ID/array	Return only hosts that have the given items.
maintenanceids	ID/array	Return only hosts that are affected by the given maintenances.
monitored_hosts	flag	Return only monitored hosts.
proxyids	ID/array	Return only hosts that are monitored by the given proxies.
proxy_groupids	ID/array	Return only hosts that are monitored by the given proxy groups.
templated_hosts	flag	Return both hosts and templates.
templateids	ID/array	Return only hosts that are linked to the given templates.
triggerids	ID/array	Return only hosts that have the given triggers.
with_items	flag	Return only hosts that have items. Overrides the `with_monitored_items` and `with_simple_graph_items` parameters.
with_item_prototypes	flag	Return only hosts that have item prototypes. Overrides the `with_simple_graph_item_prototypes` parameter.
with_simple_graph_item_prototypes	flag	Return only hosts that have item prototypes, which are enabled for creation and have numeric type of information.
with_graphs	flag	Return only hosts that have graphs.
with_graph_prototypes	flag	Return only hosts that have graph prototypes.
with_httptests	flag	Return only hosts that have web checks. Overrides the `with_monitored_httptests` parameter.
with_monitored_httptests	flag	Return only hosts that have enabled web checks.
with_monitored_items	flag	Return only hosts that have enabled items. Overrides the `with_simple_graph_items` parameter.
with_monitored_triggers	flag	Return only hosts that have enabled triggers. All of the items used in the trigger must also be enabled.
with_simple_graph_items	flag	Return only hosts that have items with numeric type of information.
with_triggers	flag	Return only hosts that have triggers. Overrides the `with_monitored_triggers` parameter.
withProblemsSuppressed	boolean	If set to `true` return only hosts with suppressed problems.
evaltype	integer	Tag evaluation method. Possible values: 0 - *(default)* And/Or; 2 - Or.

Parameter	Type	Description
severities	integer/array	Return hosts that have only problems with given severities. Applies only if problem object is trigger.
tags	object/array	Return only hosts with the given tags. Format: `[{"tag": "<tag>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all hosts. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
inheritedTags	boolean	Return hosts that have given `tags` also in all of their linked templates. Possible values: `true` - linked templates must also have given `tags`; `false` - *(default)* linked template tags are ignored.
selectDiscoveryData	query	Return a `discoveryData` property with the host discovery object data. The host discovery object links a discovered host to a host prototype from which it was discovered. It has the following properties: `host` - `(string)` ID of the host; `parent_hostid` - `(string)` ID of the host prototype from which the host has been created; `status` - `(int)` host discovery status: 0 - *(default)* host is discovered, 1 - host is not discovered anymore; `ts_delete` - `(timestamp)` time when a host that is no longer discovered will be deleted; `ts_disable` - `(timestamp)` time when a host that is no longer discovered will be disabled; `disable_source` - `(int)` indicator of whether host was disabled by an LLD rule or manually: 0 - *(default)* disabled automatically, 1 - disabled by an LLD rule.
selectDiscoveryRule	query	Return a `discoveryRule` property with the low-level discovery rule that created the host (from host prototype in VMware monitoring).
selectDiscoveryRules	query	Return a `discoveryRules` property with host LLD rules. Supports `count`.
selectGraphs	query	Return a `graphs` property with host graphs. Supports `count`.
selectHostGroups	query	Return a `hostgroups` property with host groups data that the host belongs to.
selectHttpTests	query	Return an `httpTests` property with host web scenarios. Supports `count`.
selectInterfaces	query	Return an `interfaces` property with host interfaces. Supports `count`.
selectInventory	query	Return an `inventory` property with host inventory data.
selectItems	query	Return an `items` property with host items. Supports `count`.
selectMacros	query	Return a `macros` property with host macros.

Parameter	Type	Description
selectParentTemplates	query	Return a `parentTemplates` property with templates that the host is linked to. In addition to Template object fields, it contains `link_type` - `(integer)` the way that the template is linked to host. Possible values: 0 - *(default)* manually linked; 1 - automatically linked by LLD. Supports `count`.
selectDashboards	query	Return a `dashboards` property. Supports `count`.
selectTags	query	Return a `tags` property with host tags.
selectInheritedTags	query	Return an `inheritedTags` property with tags that are on all templates which are linked to host.
selectTriggers	query	Return a `triggers` property with host triggers. Supports `count`.
selectValueMaps	query	Return a `valuemaps` property with host value maps.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type. Supports additional properties: Host interface properties.
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectParentTemplates` - results will be sorted by `host`; `selectInterfaces`; `selectItems` - sorted by `name`; `selectDiscoveryRules` - sorted by `name`; `selectTriggers` - sorted by `description`; `selectGraphs` - sorted by `name`; `selectDashboards` - sorted by `name`.
search	object	Return results that match the given pattern (case-insensitive). Accepts an object, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a `LIKE "%…%"` search. Supports only properties of `string` and `text` data type. Supports additional properties: Host interface properties.
searchInventory	object	Return hosts that have inventory data that match the given pattern (case-insensitive). Accepts an object, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a `LIKE "%…%"` search. Supports only properties of `string` and `text` data type.
sortfield	string/array	Sort the result by the given properties. Possible values: `hostid`, `host`, `name`, `status`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	

Parameter	Type	Description
limit	integer	
output	query	
preservekeys	boolean	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	
selectDiscoveries	query	Return a `discoveries` property with host low-level discovery rules. Supports `count`. This query is **deprecated**, please use `selectDiscoveryRules` instead.
selectHostDiscovery	query	Return a `hostDiscovery` property with host discovery object data. This query is **deprecated**, please use `selectDiscoveryData` instead.

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving data by name

Retrieve all data about two hosts named "Zabbix server" and "Linux server".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "filter": {
            "host": [
                "Zabbix server",
                "Linux server"
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostid": "10160",
            "proxyid": "0",
            "host": "Zabbix server",
            "status": "0",
            "ipmi_authtype": "-1",
            "ipmi_privilege": "2",
            "ipmi_username": "",
            "ipmi_password": "",
            "maintenanceid": "0",
            "maintenance_status": "0",
            "maintenance_type": "0",
            "maintenance_from": "0",
```

```
            "name": "Zabbix server",
            "flags": "0",
            "description": "The Zabbix monitoring server.",
            "tls_connect": "1",
            "tls_accept": "1",
            "tls_issuer": "",
            "tls_subject": "",
            "proxy_groupid": "0",
            "monitored_by": "0",
            "inventory_mode": "1",
            "active_available": "1",
            "assigned_proxyid": "0"
        },
        {
            "hostid": "10167",
            "proxyid": "0",
            "host": "Linux server",
            "status": "0",
            "ipmi_authtype": "-1",
            "ipmi_privilege": "2",
            "ipmi_username": "",
            "ipmi_password": "",
            "maintenanceid": "0",
            "maintenance_status": "0",
            "maintenance_type": "0",
            "maintenance_from": "0",
            "name": "Linux server",
            "flags": "0",
            "description": "",
            "tls_connect": "1",
            "tls_accept": "1",
            "tls_issuer": "",
            "tls_subject": "",
            "proxy_groupid": "0",
            "monitored_by": "0",
            "inventory_mode": "1",
            "active_available": "1",
            "assigned_proxyid": "0"
        }
    ],
    "id": 1
}
```


Retrieving host groups

Retrieve host groups that the host ”Zabbix server” is a member of.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["hostid"],
        "selectHostGroups": "extend",
        "filter": {
            "host": [
                "Zabbix server"
            ]
        }
    },
    "id": 1
}
```


Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostid": "10085",
            "hostgroups": [
                {
                    "groupid": "2",
                    "name": "Linux servers",
                    "flags": "0",
                    "uuid": "dc579cd7a1a34222933f24f52a68bcd8"
                },
                {
                    "groupid": "4",
                    "name": "Zabbix servers",
                    "flags": "0",
                    "uuid": "6f6799aa69e844b4b3918f779f2abf08"
                }
            ]
        }
    ],
    "id": 1
}
```


Retrieving linked templates

Retrieve the IDs and names of templates linked to host ”10084”.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["hostid"],
        "selectParentTemplates": [
            "templateid",
            "name"
        ],
        "hostids": "10084"
    },
    "id": 1
}
```


Response:


```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostid": "10084",
            "parentTemplates": [
                {
                    "name": "Linux",
                    "templateid": "10001"
                },
                {
                    "name": "Zabbix Server",
                    "templateid": "10047"
                }
            ]
        }
    ],
    "id": 1
}
```

Retrieving hosts by template

Retrieve hosts that have the ”10001” (*Linux by Zabbix agent*) template linked to them.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["hostid", "name"],
        "templateids": "10001"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "templateid": "10001",
            "hosts": [
                {
                    "hostid": "10084",
                    "name": "Zabbix server"
                },
                {
                    "hostid": "10603",
                    "name": "Host 1"
                },
                {
                    "hostid": "10604",
                    "name": "Host 2"
                }
            ]
        }
    ],
    "id": 1
}
```

Searching by host inventory data

Retrieve hosts that contain ”Linux” in the host inventory ”OS” field.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": [
            "host"
        ],
        "selectInventory": [
            "os"
        ],
        "searchInventory": {
            "os": "Linux"
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostid": "10084",
            "host": "Zabbix server",
            "inventory": {
                "os": "Linux Ubuntu"
            }
        },
        {
            "hostid": "10107",
            "host": "Linux server",
            "inventory": {
                "os": "Linux Mint"
            }
        }
    ],
    "id": 1
}
```

Searching by host tags

Retrieve hosts that have tag ”host-name” equal to ”linux-server”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["hostid"],
        "selectTags": "extend",
        "evaltype": 0,
        "tags": [
            {
                "tag": "host-name",
                "value": "linux-server",
                "operator": 1
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostid": "10085",
            "tags": [
                {
                    "tag": "host-name",
                    "value": "linux-server",
                    "automatic": "0"
                },
                {
                    "tag": "os",
                    "value": "rhel-7",
                    "automatic": "0"
                }
            ]
        }
```

```
    ],
    "id": 1
}
```

Retrieve hosts that have these tags not only on host level but also in their linked parent templates.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["name"],
        "tags": [
            {
                "tag": "os",
                "value": "rhel-7",
                "operator": 1
            }
        ],
        "inheritedTags": true
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostid": "10623",
            "name": "PC room 1"
        },
        {
            "hostid": "10601",
            "name": "Office"
        }
    ],
    "id": 1
}
```

Retrieve hosts excluding those with the tag ”component”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["hostid"],
        "selectTags": "extend",
        "evaltype": 0,
        "tags": [
            {
                "tag": "component",
                "value": "",
                "operator": 5
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostid": "10084",
            "tags": [
                {
                    "tag": "class",
                    "value": "os",
                    "automatic": "0"
                },
                {
                    "tag": "target",
                    "value": "linux",
                    "automatic": "0"
                }
            ]
        },
        {
            "hostid": "10629",
            "tags": [
                {
                    "tag": "class",
                    "value": "application",
                    "automatic": "0"
                },
                {
                    "tag": "target",
                    "value": "browser",
                    "automatic": "0"
                }
            ]
        }
    ],
    "id": 1
}
```


Searching host with tags and template tags

Retrieve a host with tags and all tags that are linked to parent templates.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["name"],
        "hostids": 10502,
        "selectTags": ["tag", "value"],
        "selectInheritedTags": ["tag", "value"]
    },
    "id": 1
}
```

Response:


```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostid": "10502",
            "name": "Desktop",
            "tags": [
                {
```

```
                    "tag": "host-name",
                    "value": "linux-server"
                },
                {
                    "tag": "os",
                    "value": "rhel-7"
                }
            ],
            "inheritedTags": [
                {
                    "tag": "class",
                    "value": "os"
                },
                {
                    "tag": "target",
                    "value": "linux"
                },
                {
                    "tag": "os",
                    "value": "rhel-7"
                }
            ]
        }
    ],
    "id": 1
}
```

Searching hosts by problem severity

Retrieve hosts that have ”Disaster” problems.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["name"],
        "severities": 5
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostid": "10160",
            "name": "Zabbix server"
        }
    ],
    "id": 1
}
```

Retrieve hosts that have ”Average” and ”High” problems.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["name"],
        "severities": [3, 4]
    },
```

```
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostid": "20170",
            "name": "Database"
        },
        {
            "hostid": "20183",
            "name": "workstation"
        }
    ],
    "id": 1
}
```

Retrieving host items

Retrieve items for the "Zabbix server" host. The request is limited to the item properties `itemid`, `name`, and `status`. Due to the large response size, only a subset of items is shown in the example.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["hostid"],
        "selectItems": ["itemid","name","status"],
        "filter": {
            "host": [
                "Zabbix server"
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "hostid": "10084",
      "items": [
        {
          "itemid": "42227",
          "name": "Host name of Zabbix agent running",
          "status": "0"
        },
        {
          "itemid": "42237",
          "name": "Zabbix agent ping",
          "status": "0"
        },
        {
          "itemid": "42250",
          "name": "Version of Zabbix agent running",
          "status": "0"
        },
        {
```

```
          "itemid": "42251",
          "name": "Maximum number of open file descriptors",
          "status": "0"
        },
        {
          "itemid": "42252",
          "name": "Maximum number of processes",
          "status": "0"
        },
        {
          "itemid": "69869",
          "name": "Interface enp0s3: Inbound packets discarded",
          "status": "0"
        },
        {
          "itemid": "69870",
          "name": "Interface enp0s3: Inbound packets with errors",
          "status": "0"
        },
        {
          "itemid": "69871",
          "name": "Interface enp0s3: Bits received",
          "status": "0"
        },
        {
          "itemid": "69872",
          "name": "Interface enp0s3: Outbound packets discarded",
          "status": "0"
        },
        {
          "itemid": "69873",
          "name": "Interface enp0s3: Outbound packets with errors",
          "status": "0"
        },
        {
          "itemid": "69874",
          "name": "Interface enp0s3: Bits sent",
          "status": "0"
        },
        {
          "itemid": "42253",
          "name": "Number of processes",
          "status": "0"
        },
        {
          "itemid": "42254",
          "name": "Number of running processes",
          "status": "0"
        },
        {
          "itemid": "42255",
          "name": "System boot time",
          "status": "0"
        }
      ]
    }
  ],
  "id": 1
}
```


See also

- Host group

- Template
- User macro
- Host interface
- Proxy
- Proxy group

Source

CHost::get() in *ui/include/classes/api/services/CHost.php*.

host.massadd

Description

`object host.massadd(object parameters)`

This method allows you to simultaneously add multiple related objects to all the given hosts.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters containing the IDs of the hosts to update and the objects to add to all the hosts.

The method accepts the following parameters.

Parameter	Type	Description
hosts	object/array	Hosts to be updated. The hosts must have only the `hostid` property defined. Parameter behavior: - *required*
groups	object/array	Host groups to add to the given hosts. The host groups must have only the `groupid` property defined.
interfaces	object/array	Host interfaces to be created for the given hosts.
macros	object/array	User macros to be created for the given hosts.
templates	object/array	Templates to link to the given hosts. The templates must have only the `templateid` property defined.

Return values

`(object)` Returns an object containing the IDs of the updated hosts under the `hostids` property.

Examples

Adding macros

Add two new macros to two hosts.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.massadd",
    "params": {
        "hosts": [
            {
                "hostid": "10160"
            },
            {
                "hostid": "10167"
            }
```

```
        ],
        "macros": [
            {
                "macro": "{$TEST1}",
                "value": "MACROTEST1"
            },
            {
                "macro": "{$TEST2}",
                "value": "MACROTEST2",
                "description": "Test description"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10160",
            "10167"
        ]
    },
    "id": 1
}
```

See also

- host.update
- Host group
- Template
- User macro
- Host interface

Source

CHost::massAdd() in *ui/include/classes/api/services/CHost.php*.

host.massremove

Description

`object host.massremove(object parameters)`

This method allows you to remove related objects from multiple hosts.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters containing the IDs of the hosts to update and the objects that should be removed.

Parameter	Type	Description
hostids	ID/array	IDs of the hosts to be updated. Parameter behavior: - *required*
groupids	ID/array	IDs of the host groups to remove the given hosts from.

Parameter	Type	Description
interfaces	object/array	Host interfaces to remove from the given hosts. The host interface object must have only the `ip`, `dns` and `port` properties defined.
macros	string/array	User macros to delete from the given hosts.
templateids	ID/array	IDs of the templates to unlink from the given hosts.
templateids_clear	ID/array	IDs of the templates to unlink and clear from the given hosts.

Return values

`(object)` Returns an object containing the IDs of the updated hosts under the `hostids` property.

Examples

Unlinking templates

Unlink a template from two hosts and delete all of the templated entities.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.massremove",
    "params": {
        "hostids": ["69665", "69666"],
        "templateids_clear": "325"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "69665",
            "69666"
        ]
    },
    "id": 1
}
```

See also

- host.update
- User macro
- Host interface

Source

CHost::massRemove() in *ui/include/classes/api/services/CHost.php*.

host.update

Description

`object host.update(object/array hosts)`

This method allows you to update existing hosts.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

(object/array) Host properties to be updated.

The hostid property must be defined for each host, all other properties are optional. Only the given properties will be updated, all others will remain unchanged.

Note, however, that updating the host technical name will also update the host's visible name (if not given separately) by the host's technical name value.

Additionally to the standard host properties, the method accepts the following parameters.

Parameter	Type	Description
groups	object/array	Host groups to replace the current host groups the host belongs to. All host groups that are not listed in the request will be unlinked. The host groups must have only the groupid property defined.
interfaces	object/array	Host interfaces to replace the current host interfaces. All interfaces that are not listed in the request will be removed.
tags	object/array	Host tags to replace the current host tags. All tags that are not listed in the request will be removed.
inventory	object	Host inventory properties.
macros	object/array	User macros to replace the current user macros. All macros that are not listed in the request will be removed.
templates	object/array	Templates to replace the currently linked templates. All templates that are not listed in the request will be only unlinked. The templates must have only the templateid property defined.
templates_clear	object/array	Templates to unlink and clear from the host. The templates must have only the templateid property defined.

Return values

(object) Returns an object containing the IDs of the updated hosts under the hostids property.

Examples

Enabling a host

Enable host monitoring, that is, set its status to "0".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.update",
    "params": {
        "hostid": "10126",
        "status": 0
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10126"
        ]
    },
    "id": 1
}
```

Enabling multiple hosts

Enable monitoring of two hosts, that is, set their status to "0".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.update",
    "params": [
        {
            "hostid": "10127",
            "status": 0
        },
        {
            "hostid": "10128",
            "status": 0
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10127",
            "10128"
        ]
    },
    "id": 1
}
```

Unlinking templates

Unlink and clear two templates from host.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.update",
    "params": {
        "hostid": "10126",
        "templates_clear": [
            {
                "templateid": "10124"
            },
            {
                "templateid": "10125"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10126"
        ]
    },
    "id": 1
}
```

Updating host macros

Replace all host macros with two new ones.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.update",
    "params": {
        "hostid": "10126",
        "macros": [
            {
                "macro": "{$PASS}",
                "value": "password"
            },
            {
                "macro": "{$DISC}",
                "value": "sda",
                "description": "Updated description"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10126"
        ]
    },
    "id": 1
}
```

Updating host inventory

Change inventory mode and add location

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.update",
    "params": {
        "hostid": "10387",
        "inventory_mode": 0,
        "inventory": {
            "location": "Latvia, Riga"
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10387"
        ]
    },
    "id": 1
}
```

Updating host tags

Replace all host tags with a new one.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.update",
    "params": {
        "hostid": "10387",
        "tags": {
            "tag": "os",
            "value": "rhel-7"
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10387"
        ]
    },
    "id": 1
}
```

Updating discovered host macros

Convert discovery rule created ”automatic” macro to ”manual” and change its value to ”new-value”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.update",
    "params": {
        "hostid": "10387",
        "macros": {
            "hostmacroid": "5541",
            "value": "new-value",
            "automatic": "0"
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10387"
        ]
    },
    "id": 1
}
```

Updating host encryption

Update the host ”10590” to use PSK encryption only for connections from host to Zabbix server, and change the PSK identity and PSK key. Note that the host has to be pre-configured to use PSK.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.update",
```

```
    "params": {
        "hostid": "10590",
        "tls_connect": 1,
        "tls_accept": 2,
        "tls_psk_identity": "PSK 002",
        "tls_psk": "e560cb0d918d26d31b4f642181f5f570ad89a390931102e5391d08327ba434e9"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10590"
        ]
    },
    "id": 1
}
```

See also

- host.massadd
- host.massremove
- Host group
- Template
- User macro
- Host interface
- Host inventory
- Host tag
- Proxy
- Proxy group

Source

CHost::update() in *ui/include/classes/api/services/CHost.php*.

Host group

This class is designed to work with host groups.

Object references:

- Host group

Available methods:

- hostgroup.create - create new host groups
- hostgroup.delete - delete host groups
- hostgroup.get - retrieve host groups
- hostgroup.massadd - add related objects to host groups
- hostgroup.massremove - remove related objects from host groups
- hostgroup.propagate - propagate permissions and tag filters to host groups' subgroups
- hostgroup.update - update host groups

Host group object

The following objects are directly related to the `hostgroup` API.

Host group

The host group object has the following properties.

Property	Type	Description
groupid	ID	ID of the host group. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the host group. Property behavior: - *required* for create operations
flags	integer	Origin of the host group. Possible values: 0 - a plain host group; 4 - a discovered host group. Property behavior: - *read-only*
uuid	string	Universal unique identifier, used for linking imported host groups to already existing ones. Auto-generated, if not given.

hostgroup.create

Description

`object hostgroup.create(object/array hostGroups)`

This method allows you to create new host groups.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Host groups to create.

The method accepts host groups with the standard host group properties.

Return values

`(object)` Returns an object containing the IDs of the created host groups under the `groupids` property. The order of the returned IDs matches the order of the passed host groups.

Examples

Creating a host group

Create a host group called ”Linux servers”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostgroup.create",
    "params": {
        "name": "Linux servers"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
```

```
            "107819"
        ]
    },
    "id": 1
}
```

Source

CHostGroup::create() in *ui/include/classes/api/services/CHostGroup.php*.

hostgroup.delete

Description

`object hostgroup.delete(array hostGroupIds)`

This method allows you to delete host groups.

A host group cannot be deleted if:

- it contains hosts that belong to this group only;
- it is marked as internal;
- it is used by a host prototype;
- it is used in a global script;
- it is used in a correlation condition.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the host groups to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted host groups under the `groupids` property.

Examples

Deleting multiple host groups

Delete two host groups.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostgroup.delete",
    "params": [
        "107824",
        "107825"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
            "107824",
            "107825"
        ]
    },
    "id": 1
}
```

Source

CHostGroup::delete() in *ui/include/classes/api/services/CHostGroup.php*.

hostgroup.get

Description

`integer/array hostgroup.get(object parameters)`

The method allows you to retrieve host groups according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
graphids	ID/array	Return only host groups that contain hosts with the given graphs.
groupids	ID/array	Return only host groups with the given host group IDs.
hostids	ID/array	Return only host groups that contain the given hosts.
maintenanceids	ID/array	Return only host groups that are affected by the given maintenances.
triggerids	ID/array	Return only host groups that contain hosts with the given triggers.
with_graphs	boolean	Return only host groups that contain hosts with graphs.
with_graph_prototypes	boolean	Return only host groups that contain hosts with graph prototypes.
with_hosts	boolean	Return only host groups that contain hosts.
with_httptests	boolean	Return only host groups that contain hosts with web checks. Overrides the `with_monitored_httptests` parameter.
with_items	boolean	Return only host groups that contain hosts with items. Overrides the `with_monitored_items` and `with_simple_graph_items` parameters.
with_item_prototypes	boolean	Return only host groups that contain hosts with item prototypes. Overrides the `with_simple_graph_item_prototypes` parameter.
with_simple_graph_item_prototypes	boolean	Return only host groups that contain hosts with item prototypes, which are enabled for creation and have numeric type of information.
with_monitored_httptests	boolean	Return only host groups that contain hosts with enabled web checks.
with_monitored_hosts	boolean	Return only host groups that contain monitored hosts.
with_monitored_items	boolean	Return only host groups that contain hosts with enabled items. Overrides the `with_simple_graph_items` parameter.
with_monitored_triggers	boolean	Return only host groups that contain hosts with enabled triggers. All of the items used in the trigger must also be enabled.
with_simple_graph_items	boolean	Return only host groups that contain hosts with numeric items.
with_triggers	boolean	Return only host groups that contain hosts with triggers. Overrides the `with_monitored_triggers` parameter.
selectDiscoveryRules	query	Return a `discoveryRules` property with the LLD rules that discovered the host group.

Parameter	Type	Description
selectDiscoveryData	query	Return a `discoveryData` property with the host group discovery objects. Each host group discovery object is a host group prototype linked to the discovered host group and has the following properties: `parent_group_prototypeid` - `(ID)` ID of the host group prototype from which the host group was discovered; `name` - `(string)` name of the host group prototype; `ts_delete` - `(timestamp)` time when the host group that is no longer discovered will be deleted; `status` - `(int)` host group discovery status: 0 - *(default)* host group is discovered, 1 - host group is not discovered anymore.
selectHostPrototypes	query	Return a `hostPrototypes` property with host prototypes that discovered this host group.
selectHosts	query	Return a `hosts` property with the hosts that belong to the host group. Supports `count`.
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectHosts` - results will be sorted by `host`.
sortfield	string/array	Sort the result by the given properties. Possible values: `groupid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	
selectGroupDiscoveries	query	Return a `groupDiscoveries` property with the host group discovery objects. Each host group discovery object is a host group prototype linked to the discovered host group. This query is **deprecated**, please use `selectDiscoveryData` instead.

Return values

`(integer/array)` Returns either:

- an array of objects
- the count of retrieved objects, if the `countOutput` parameter has been used

Examples

Retrieving data by name

Retrieve all data about two host groups named ”Zabbix servers” and ”Linux servers”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostgroup.get",
```

```
    "params": {
        "output": "extend",
        "filter": {
            "name": [
                "Zabbix servers",
                "Linux servers"
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "groupid": "2",
            "name": "Linux servers",
            "flags": "0",
            "uuid": "dc579cd7a1a34222933f24f52a68bcd8"
        },
        {
            "groupid": "4",
            "name": "Zabbix servers",
            "flags": "0",
            "uuid": "6f6799aa69e844b4b3918f779f2abf08"
        }
    ],
    "id": 1
}
```

See also

- Host

Source

CHostGroup::get() in *ui/include/classes/api/services/CHostGroup.php*.

hostgroup.massadd

Description

`object hostgroup.massadd(object parameters)`

This method allows you to simultaneously add multiple related objects to all the given host groups.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters containing the IDs of the host groups to update and the objects to add to all the host groups.

The method accepts the following parameters.

Parameter	Type	Description
groups	object/array	Host groups to be updated. The host groups must have only the `groupid` property defined. Parameter behavior: - *required*
hosts	object/array	Hosts to add to all host groups. The hosts must have only the `hostid` property defined.

Return values

`(object)` Returns an object containing the IDs of the updated host groups under the `groupids` property.

Examples

Adding hosts to host groups

Add two hosts to host groups with IDs 5 and 6.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostgroup.massadd",
    "params": {
        "groups": [
            {
                "groupid": "5"
            },
            {
                "groupid": "6"
            }
        ],
        "hosts": [
            {
                "hostid": "30050"
            },
            {
                "hostid": "30001"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
            "5",
            "6"
        ]
    },
    "id": 1
}
```

See also

- Host

Source

CHostGroup::massAdd() in *ui/include/classes/api/services/CHostGroup.php*.

hostgroup.massremove

Description

`object hostgroup.massremove(object parameters)`

This method allows you to remove related objects from multiple host groups.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters containing the IDs of the host groups to update and the objects that should be removed.

Parameter	Type	Description
groupids	ID/array	IDs of the host groups to be updated. Parameter behavior: *- required*
hostids	ID/array	IDs of the hosts to remove from all host groups.

Return values

`(object)` Returns an object containing the IDs of the updated host groups under the `groupids` property.

Examples

Removing hosts from host groups

Remove two hosts from the given host groups.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostgroup.massremove",
    "params": {
        "groupids": [
            "5",
            "6"
        ],
        "hostids": [
            "30050",
            "30001"
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
            "5",
            "6"
        ]
    },
    "id": 1
}
```

Source

CHostGroup::massRemove() in *ui/include/classes/api/services/CHostGroup.php*.

hostgroup.propagate

Description

`object hostgroup.propagate(object parameters)`

This method allows you to apply permissions and tag filters to all subgroups of a host group.

> **Note:**
> This method is only available to *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
groups	object/array	Host groups to propagate. The host groups must have the `groupid` property defined. Parameter behavior: - *required*
permissions	boolean	Set to "true" to propagate permissions. Parameter behavior: - *required* if `tag_filters` is not set
tag_filters	boolean	Set to "true" to propagate tag filters. Parameter behavior: - *required* if `permissions` is not set

Return values

`(object)` Returns an object containing the IDs of the propagated host groups under the `groupids` property.

Examples

Propagating host group permissions and tag filters to its subgroups.

Propagate host group permissions and tag filters to its subgroups.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostgroup.propagate",
    "params": {
        "groups": [
            {
                "groupid": "6"
            }
        ],
        "permissions": true,
        "tag_filters": true
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
```

```
            "6",
        ]
    },
    "id": 1
}
```

See also

- hostgroup.update
- hostgroup.massadd
- Host

Source

CHostGroup::propagate() in *ui/include/classes/api/services/CHostGroup.php*.

hostgroup.update

Description

`object hostgroup.update(object/array hostGroups)`

This method allows you to update existing hosts groups.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Host group properties to be updated.

The `groupid` property must be defined for each host group, all other properties are optional. Only the given properties will be updated, all others will remain unchanged.

Return values

`(object)` Returns an object containing the IDs of the updated host groups under the `groupids` property.

Examples

Renaming a host group

Rename a host group to "Linux hosts".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostgroup.update",
    "params": {
        "groupid": "7",
        "name": "Linux hosts"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
            "7"
        ]
    },
    "id": 1
}
```

Renaming multiple host groups

Rename two host groups to include region information in their names.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostgroup.update",
    "params": [
        {
            "groupid": "8",
            "name": "Europe/Linux hosts"
        },
        {
            "groupid": "9",
            "name": "Europe/Windows hosts"
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
            "8",
            "9"
        ]
    },
    "id": 1
}
```

Source

CHostGroup::update() in *ui/include/classes/api/services/CHostGroup.php*.

Host interface

This class provides access to host interfaces.

Object references:

- Host interface
  - Details

Available method:

- hostinterface.get - retrieve host interfaces

Host interfaces are created, updated, and deleted using host methods.

Host interface object

The following objects are directly related to the `hostinterface` API.

Host interface

The host interface object has the following properties.

Property	Type	Description
interfaceid	ID	ID of the interface. Property behavior: - *read-only* - *required* for update operations

Property	Type	Description
available	integer	Availability of host interface. Possible values: 0 - *(default)* unknown; 1 - available; 2 - unavailable. Property behavior: - *read-only*
hostid	ID	ID of the host that the interface belongs to. Property behavior: - *constant* - *required* for create operations
type	integer	Interface type. Possible values: 1 - Agent; 2 - SNMP; 3 - IPMI; 4 - JMX. Property behavior: - *required* for create operations
ip	string	IP address used by the interface. Can be empty if the connection is made via DNS. Property behavior: - *required* for create operations
dns	string	DNS name used by the interface. Can be empty if the connection is made via IP. Property behavior: - *required* for create operations
port	string	Port number used by the interface. Can contain user macros. Property behavior: - *required* for create operations
useip	integer	Whether the connection should be made via IP. Possible values: 0 - connect using host DNS name; 1 - connect using host IP address. Property behavior: - *required* for create operations
main	integer	Whether the interface is used as default on the host. Only one interface of some type can be set as default on a host. Possible values: 0 - not default; 1 - default. Property behavior: - *required* for create operations
details	object	Additional details object for interface. Property behavior: - *required* if `type` is set to ”SNMP”

Property	Type	Description
disable_until	timestamp	The next polling time of an unavailable host interface. Property behavior: - *read-only*
error	string	Error text if host interface is unavailable. Property behavior: - *read-only*
errors_from	timestamp	Time when host interface became unavailable. Property behavior: - *read-only*

Details

The details object has the following properties.

Property	Type	Description
version	integer	SNMP interface version. Possible values: 1 - SNMPv1; 2 - SNMPv2c; 3 - SNMPv3. Property behavior: - *required*
bulk	integer	Whether to use bulk SNMP requests. Possible values: 0 - don't use bulk requests; 1 - (default) - use bulk requests.
community	string	SNMP community. Used only by SNMPv1 and SNMPv2 interfaces. Property behavior: - *required* if `version` is set to "SNMPv1" or "SNMPv2c"
max_repetitions	string	Max repetition value for native SNMP bulk requests (GetBulkRequest-PDUs). Used only for the `walk[]` item in SNMPv2 and SNMPv3. User macros are supported. Possible values: 1–2147483647 (INT32_MAX). Default: 10.
retries	string	Maximum number of retries if a request times out. Used for `walk[]` and `get[]` items. For `walk[]`, retries are shared across the entire walk. User macros are supported. Possible values: 0–100 (0 - disable retries). Default: 5.
securityname	string	SNMPv3 security name. Used only by SNMPv3 interfaces.
securitylevel	integer	SNMPv3 security level. Used only by SNMPv3 interfaces. Possible values: 0 - (default) - noAuthNoPriv; 1 - authNoPriv; 2 - authPriv.
authpassphrase	string	SNMPv3 authentication passphrase. Used only by SNMPv3 interfaces.
privpassphrase	string	SNMPv3 privacy passphrase. Used only by SNMPv3 interfaces.

Property	Type	Description
authprotocol	integer	SNMPv3 authentication protocol. Used only by SNMPv3 interfaces. Possible values: 0 - (default) - MD5; 1 - SHA1; 2 - SHA224; 3 - SHA256; 4 - SHA384; 5 - SHA512.
privprotocol	integer	SNMPv3 privacy protocol. Used only by SNMPv3 interfaces. Possible values: 0 - (default) - DES; 1 - AES128; 2 - AES192; 3 - AES256; 4 - AES192C; 5 - AES256C.
contextname	string	SNMPv3 context name. Used only by SNMPv3 interfaces.

hostinterface.get

Description

`integer/array hostinterface.get(object parameters)`

The method allows you to retrieve host interfaces according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
hostids	ID/array	Return only host interfaces used by the given hosts.
interfaceids	ID/array	Return only host interfaces with the given IDs.
itemids	ID/array	Return only host interfaces used by the given items.
triggerids	ID/array	Return only host interfaces used by items in the given triggers.
selectItems	query	Return an `items` property with the items that use the interface. Supports `count`.
selectHosts	query	Return a `hosts` property with an array of hosts that use the interface.
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectItems`.
sortfield	string/array	Sort the result by the given properties. Possible values: `interfaceid`, `dns`, `ip`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	

Parameter	Type	Description
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve host interfaces

Retrieve all data about the interfaces used by host ”10809”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostinterface.get",
    "params": {
        "output": "extend",
        "hostids": "10809"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "interfaceid": "84",
            "hostid": "10809",
            "main": "1",
            "type": "1",
            "useip": "0",
            "ip": "",
            "dns": "localhost",
            "port": "10050",
            "available": "0",
            "error": "",
            "errors_from": "0",
            "disable_until": "0",
            "details": []
        },
        {
            "interfaceid": "83",
            "hostid": "10809",
            "main": "1",
            "type": "2",
            "useip": "1",
            "ip": "127.0.0.1",
            "dns": "",
            "port": "161",
            "available": "0",
            "error": "",
            "errors_from": "0",
            "disable_until": "0",
            "details": {
```

```
                "version": "2",
                "bulk": "1",
                "community": "{$SNMP_COMMUNITY}",
                "max_repetitions": "10",
                "retries": "5"
            }
        }
    ],
    "id": 1
}
```

See also

- Host
- Item

Source

CHostInterface::get() in *ui/include/classes/api/services/CHostInterface.php*.

Host prototype

This class is designed to work with host prototypes.

Object references:

- Host prototype
- Group link
- Group prototype
- Host prototype tag
- Custom interface
    - Custom interface details

Available methods:

- hostprototype.create - create new host prototypes
- hostprototype.delete - delete host prototypes
- hostprototype.get - retrieve host prototypes
- hostprototype.update - update host prototypes

Host prototype object

The following objects are directly related to the `hostprototype` API.

Host prototype

The host prototype object has the following properties.

Property	Type	Description
hostid	ID	ID of the host prototype. Property behavior: - *read-only* - *required* for update operations
host	string	Technical name of the host prototype. Property behavior: - *required* for create operations - *read-only* for inherited objects
name	string	Visible name of the host prototype. Default: `host` property value. Property behavior: - *read-only* for inherited objects

Property	Type	Description
status	integer	Status of the host prototype. Possible values: 0 - *(default)* monitored host; 1 - unmonitored host.
flags	integer	Origin of the host prototype. Possible values: 2 - a host prototype; 6 - a discovered host prototype Property behavior: - *read-only*
inventory_mode	integer	Host inventory population mode. Possible values: -1 - *(default)* disabled; 0 - manual; 1 - automatic.
templateid	ID	ID of the parent template host prototype. Property behavior: - *read-only*
discover	integer	Host prototype discovery status. Possible values: 0 - *(default)* new hosts will be discovered; 1 - new hosts will not be discovered and existing hosts will be marked as lost.
custom_interfaces	integer	Source of custom interfaces for hosts created by the host prototype. Possible values: 0 - *(default)* inherit interfaces from parent host; 1 - use host prototypes custom interfaces. Property behavior: - *read-only* for inherited objects
uuid	string	Universal unique identifier, used for linking imported host prototypes to already existing ones. Auto-generated, if not given. Property behavior: - *supported* if the host prototype belongs to a template

Group link

The group link object links a host prototype with a host group. It has the following properties.

Property	Type	Description
groupid	ID	ID of the host group. Property behavior: - *required*

Group prototype

The group prototype object defines a group that will be created for a discovered host. It has the following properties.

Property	Type	Description
group_prototypeid	ID	ID of the group prototype. Property behavior: - *read-only*
name	string	Name of the group prototype. Property behavior: - *required* for create operations

Host prototype tag

The host prototype tag object has the following properties.

Property	Type	Description
tag	string	Host prototype tag name. Property behavior: - *required*
value	string	Host prototype tag value. Property behavior: - *read-only*.
object	integer	Type of object the tag was inherited from. Possible values: 0 - template. Property behavior: - *read-only*.
objectid	ID	ID of the object the tag was inherited from. Property behavior: - *read-only*.

Custom interface

Custom interfaces are supported if `custom_interfaces` of Host prototype object is set to "use host prototypes custom interfaces". The custom interface object has the following properties.

Property	Type	Description
type	integer	Interface type. Possible values: 1 - Agent; 2 - SNMP; 3 - IPMI; 4 - JMX. Property behavior: - *required*
useip	integer	Whether the connection should be made via IP. Possible values: 0 - connect using host DNS name; 1 - connect using host IP address. Property behavior: - *required*

Property	Type	Description
ip	string	IP address used by the interface. Can contain macros. Property behavior: - *required* if `useip` is set to ”connect using host IP address”
dns	string	DNS name used by the interface. Can contain macros. Property behavior: - *required* if `useip` is set to ”connect using host DNS name”
port	string	Port number used by the interface. Can contain user and LLD macros. Property behavior: - *required*
main	integer	Whether the interface is used as default on the host. Only one interface of some type can be set as default on a host. Possible values: 0 - not default; 1 - default. Property behavior: - *required*
details	object	Additional object for custom interface details. Property behavior: - *required* if `type` is set to ”SNMP”

Custom interface details

The details object has the following properties.

Property	Type	Description
version	integer	SNMP interface version. Possible values: 1 - SNMPv1; 2 - SNMPv2c; 3 - SNMPv3. Property behavior: - *required*
bulk	integer	Whether to use bulk SNMP requests. Possible values: 0 - don’t use bulk requests; 1 - *(default)* - use bulk requests.
community	string	SNMP community. Property behavior: - *required* if `version` is set to ”SNMPv1” or ”SNMPv2c”
max_repetitions	string	Max repetition value for native SNMP bulk requests (GetBulkRequest-PDUs). Used only for the `walk[]` item in SNMPv2 and SNMPv3. LLD macros and user macros are supported. Possible values: 1–2147483647 (INT32_MAX). Default: 10.

Property	Type	Description
retries	string	Maximum number of retries if a request times out. Used for `walk[]` and `get[]` items. For `walk[]`, retries are shared across the entire walk. LLD macros and user macros are supported. Possible values: 0–100 (0 - disable retries). Default: 5.
securityname	string	SNMPv3 security name. Property behavior: - *supported* if `version` is set to ”SNMPv3”
securitylevel	integer	SNMPv3 security level. Possible values: 0 - *(default)* - noAuthNoPriv; 1 - authNoPriv; 2 - authPriv. Property behavior: - *supported* if `version` is set to ”SNMPv3”
authpassphrase	string	SNMPv3 authentication passphrase. Property behavior: - *supported* if `version` is set to ”SNMPv3” and `securitylevel` is set to ”authNoPriv” or ”authPriv”
privpassphrase	string	SNMPv3 privacy passphrase. Property behavior: - *supported* if `version` is set to ”SNMPv3” and `securitylevel` is set to ”authPriv”
authprotocol	integer	SNMPv3 authentication protocol. Possible values: 0 - *(default)* - MD5; 1 - SHA1; 2 - SHA224; 3 - SHA256; 4 - SHA384; 5 - SHA512. Property behavior: - *supported* if `version` is set to ”SNMPv3” and `securitylevel` is set to ”authNoPriv” or ”authPriv”
privprotocol	integer	SNMPv3 privacy protocol. Used only by SNMPv3 interfaces. Possible values: 0 - *(default)* - DES; 1 - AES128; 2 - AES192; 3 - AES256; 4 - AES192C; 5 - AES256C. Property behavior: - *supported* if `version` is set to ”SNMPv3” and `securitylevel` is set to ”authPriv”
contextname	string	SNMPv3 context name. Property behavior: - *supported* if `version` is set to ”SNMPv3”

hostprototype.create

Description

`object hostprototype.create(object/array hostPrototypes)`

This method allows you to create new host prototypes.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Host prototypes to create.

Additionally to the standard host prototype properties, the method accepts the following parameters.

Parameter	Type	Description
groupLinks	array	Group links to be created for the host prototype. Parameter behavior: - *required*
ruleid	ID	ID of the LLD rule that the host prototype belongs to. Parameter behavior: - *required*
groupPrototypes	array	Group prototypes to be created for the host prototype.
macros	object/array	User macros to be created for the host prototype.
tags	object/array	Host prototype tags.
interfaces	object/array	Host prototype custom interfaces.
templates	object/array	Templates to be linked to the host prototype. The templates must have only the `templateid` property defined.

Return values

`(object)` Returns an object containing the IDs of the created host prototypes under the `hostids` property. The order of the returned IDs matches the order of the passed host prototypes.

Examples

Creating a host prototype

Create a host prototype "{#VM.NAME}" on LLD rule "23542" with a group prototype "{#HV.NAME}", tag pair "datacenter": "{#DATACENTER.NAME}" and custom SNMPv2 interface 127.0.0.1:161 with community {$SNMP_COMMUNITY}. Link it to host group "2".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostprototype.create",
    "params": {
        "host": "{#VM.NAME}",
        "ruleid": "23542",
        "custom_interfaces": "1",
        "groupLinks": [
            {
                "groupid": "2"
            }
        ],
        "groupPrototypes": [
            {
                "name": "{#HV.NAME}"
            }
```

```
        ],
        "tags": [
            {
                "tag": "datacenter",
                "value": "{#DATACENTER.NAME}"
            }
        ],
        "interfaces": [
            {
                "main": "1",
                "type": "2",
                "useip": "1",
                "ip": "127.0.0.1",
                "dns": "",
                "port": "161",
                "details": {
                    "version": "2",
                    "bulk": "1",
                    "community": "{$SNMP_COMMUNITY}"
                }
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10103"
        ]
    },
    "id": 1
}
```

See also

- Group link
- Group prototype
- Host prototype tag
- Custom interface
- User macro

Source

CHostPrototype::create() in *ui/include/classes/api/services/CHostPrototype.php*.

hostprototype.delete

Description

`object hostprototype.delete(array hostPrototypeIds)`

This method allows you to delete host prototypes.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the host prototypes to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted host prototypes under the `hostids` property.

Examples

Deleting multiple host prototypes

Delete two host prototypes.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostprototype.delete",
    "params": [
        "10103",
        "10105"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10103",
            "10105"
        ]
    },
    "id": 1
}
```

Source

CHostPrototype::delete() in *ui/include/classes/api/services/CHostPrototype.php*.

hostprototype.get

Description

`integer/array hostprototype.get(object parameters)`

The method allows you to retrieve host prototypes according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
hostids	ID/array	Return only host prototypes with the given IDs.
discoveryids	ID/array	Return only host prototypes that belong to the given LLD rules.
inherited	boolean	If set to `true` return only items inherited from a template.

Parameter	Type	Description
selectDiscoveryData	query	Return a `discoveryData` property with the host prototype discovery object data. The host prototype discovery object links a discovered host prototype to a host prototype from which it was discovered. It has the following properties: `host` - `(string)` ID of the host; `parent_hostid` - `(string)` ID of the host prototype from which the host prototype has been created; `status` - `(int)` host prototype discovery status: 0 - *(default)* host prototype is discovered, 1 - host prototype is not discovered anymore; `ts_delete` - `(timestamp)` time when a host prototype that is no longer discovered will be deleted; `ts_disable` - `(timestamp)` time when a host prototype that is no longer discovered will be disabled; `disable_source` - `(int)` indicator of whether host prototype was disabled by an LLD rule or manually: 0 - *(default)* disabled automatically, 1 - disabled by an LLD rule.
selectDiscoveryRule	query	Return a `discoveryRule` property with the LLD rule that the host prototype belongs to.
selectDiscoveryRulePrototype	query	Return a `discoveryRulePrototype` property with the parent LLD rule prototype that the host prototype belongs to.
selectInterfaces	query	Return an `interfaces` property with host prototype custom interfaces.
selectGroupLinks	query	Return a `groupLinks` property with the group links of the host prototype.
selectGroupPrototypes	query	Return a `groupPrototypes` property with the group prototypes of the host prototype.
selectInheritedTags	query	Return an `inheritedTags` property with tags that are on linked templates.
selectMacros	query	Return a `macros` property with host prototype macros.
selectParentHost	query	Return a `parentHost` property with the host that the host prototype belongs to.
selectTags	query	Return a `tags` property with host prototype tags.
selectTemplates	query	Return a `templates` property with the templates linked to the host prototype. Supports `count`.
sortfield	string/array	Sort the result by the given properties. Possible values: `hostid`, `host`, `name`, `status`, `discovered`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving host prototypes from an LLD rule

Retrieve all host prototypes, their group links, group prototypes and tags from an LLD rule.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostprototype.get",
    "params": {
        "output": "extend",
        "selectInterfaces": "extend",
        "selectGroupLinks": "extend",
        "selectGroupPrototypes": "extend",
        "selectTags": "extend",
        "discoveryids": "23554"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostid": "10092",
            "host": "{#HV.UUID}",
            "name": "{#HV.UUID}",
            "status": "0",
            "templateid": "0",
            "discover": "0",
            "custom_interfaces": "1",
            "uuid": "051a1469d4d045cbbf818fcc843a352e",
            "flags": "2",
            "inventory_mode": "-1",
            "tags": [
                {
                    "tag": "datacenter",
                    "value": "{#DATACENTER.NAME}"
                },
                {
                    "tag": "instance-type",
                    "value": "{#INSTANCE_TYPE}"
                }
            ],
            "groupLinks": [
                {
                    "group_prototypeid": "4",
                    "hostid": "10092",
                    "groupid": "7",
                    "templateid": "0"
                }
            ],
            "groupPrototypes": [
                {
                    "group_prototypeid": "7",
                    "hostid": "10092",
                    "name": "{#CLUSTER.NAME}",
                    "templateid": "0"
                }
            ],
            "interfaces": [
                {
```

```
                    "main": "1",
                    "type": "2",
                    "useip": "1",
                    "ip": "127.0.0.1",
                    "dns": "",
                    "port": "161",
                    "available": "0",
                    "error": "",
                    "errors_from": "0",
                    "disable_until": "0",
                    "details": {
                        "version": "2",
                        "bulk": "1",
                        "community": "{$SNMP_COMMUNITY}",
                        "max_repetitions": "10",
                        "retries": "5"
                    }
                }
            ]
        }
    ],
    "id": 1
}
```


See also

- Group link
- Group prototype
- User macro

Source

CHostPrototype::get() in *ui/include/classes/api/services/CHostPrototype.php*.

hostprototype.update

Description

`object hostprototype.update(object/array hostPrototypes)`

This method allows you to update existing host prototypes.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Host prototype properties to be updated.

The `hostid` property must be defined for each host prototype, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard host prototype properties, the method accepts the following parameters.

Parameter	Type	Description
groupLinks	array	Group links to replace the current group links on the host prototype. Parameter behavior: - *read-only* for inherited objects

Parameter	Type	Description
groupPrototypes	array	Group prototypes to replace the existing group prototypes on the host prototype. All group prototypes that are not listed in the request will be removed. Parameter behavior: - *read-only* for inherited objects
macros	object/array	User macros to replace the current user macros. All macros that are not listed in the request will be removed.
tags	object/array	Host prototype tags to replace the current tags. All tags that are not listed in the request will be removed. Parameter behavior: - *read-only* for inherited objects
interfaces	object/array	Host prototype custom interfaces to replace the current interfaces. Custom interface object should contain all its parameters. All interfaces that are not listed in the request will be removed. Parameter behavior: - *supported* if `custom_interfaces` of Host prototype object is set to "use host prototypes custom interfaces" - *read-only* for inherited objects
templates	object/array	Templates to replace the currently linked templates. The templates must have only the `templateid` property defined.

Return values

`(object)` Returns an object containing the IDs of the updated host prototypes under the `hostids` property.

Examples

Disabling a host prototype

Disable a host prototype, that is, set its status to "1".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostprototype.update",
    "params": {
        "hostid": "10092",
        "status": 1
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10092"
        ]
    },
    "id": 1
}
```

Updating host prototype tags

Replace host prototype tags with new ones.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostprototype.update",
    "params": {
        "hostid": "10092",
        "tags": [
            {
                "tag": "datacenter",
                "value": "{#DATACENTER.NAME}"
            },
            {
                "tag": "instance-type",
                "value": "{#INSTANCE_TYPE}"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostids": [
            "10092"
        ]
    },
    "id": 1
}
```

Updating host prototype custom interfaces

Replace inherited interfaces with host prototype custom interfaces.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "hostprototype.update",
    "params": {
        "hostid": "10092",
        "custom_interfaces": "1",
        "interfaces": [
            {
                "main": "1",
                "type": "2",
                "useip": "1",
                "ip": "127.0.0.1",
                "dns": "",
                "port": "161",
                "details": {
                    "version": "2",
                    "bulk": "1",
                    "community": "{$SNMP_COMMUNITY}"
                }
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
```

```
    "result": {
        "hostids": [
            "10092"
        ]
    },
    "id": 1
}
```

See also

- Group link
- Group prototype
- Host prototype tag
- Custom interface
- User macro

Source

CHostPrototype::update() in *ui/include/classes/api/services/CHostPrototype.php*.

Housekeeping

This class is designed to work with housekeeping.

Object references:

- Housekeeping

Available methods:

- housekeeping.get - retrieve housekeeping
- housekeeping.update - update housekeeping

Housekeeping object

The following objects are directly related to the `housekeeping` API.

Housekeeping

The settings object has the following properties.

Property	Type	Description
hk_events_mode	integer	Enable internal housekeeping for events and alerts. Possible values: 0 - Disable; 1 - *(default)* Enable.
hk_events_trigger	string	Trigger data storage period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Default: 365d.
hk_events_service	string	Service data storage period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Default: 1d.
hk_events_internal	string	Internal data storage period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Default: 1d.
hk_events_discovery	string	Network discovery data storage period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Default: 1d.

Property	Type	Description
hk_events_autoreg	string	Autoregistration data storage period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Default: 1d.
hk_services_mode	integer	Enable internal housekeeping for services. Possible values: 0 - Disable; 1 - *(default)* Enable.
hk_services	string	Services data storage period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Default: 365d.
hk_audit_mode	integer	Enable internal housekeeping for audit. Possible values: 0 - Disable; 1 - *(default)* Enable.
hk_audit	string	Audit data storage period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Default: 31d.
hk_sessions_mode	integer	Enable internal housekeeping for sessions. Possible values: 0 - Disable; 1 - *(default)* Enable.
hk_sessions	string	Sessions data storage period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Default: 31d.
hk_history_mode	integer	Enable internal housekeeping for history. Possible values: 0 - Disable; 1 - *(default)* Enable.
hk_history_global	integer	Override item history period. Possible values: 0 - Do not override; 1 - *(default)* Override.
hk_history	string	History data storage period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Default: 31d.
hk_trends_mode	integer	Enable internal housekeeping for trends. Possible values: 0 - Disable; 1 - *(default)* Enable.
hk_trends_global	integer	Override item trend period. Possible values: 0 - Do not override; 1 - *(default)* Override.
hk_trends	string	Trends data storage period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Default: 365d.

Property	Type	Description
db_extension	string	Configuration flag DB extension. If this flag is set to ”timescaledb” then the server changes its behavior for housekeeping and item deletion. Property behavior: - *read-only*
compression_availability	integer	Whether data compression is supported by the database (or its extension). Possible values: 0 - Unavailable; 1 - Available. Property behavior: - *read-only*
compression_status	integer	Enable TimescaleDB compression for history and trends. Possible values: 0 - *(default)* Off; 1 - On.
compress_older	string	Compress history and trends records older than specified period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Default: 7d.

housekeeping.get

Description

`object housekeeping.get(object parameters)`

The method allows you to retrieve housekeeping object according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports only one parameter.

Parameter	Type	Description
output	query	This parameter is described in the reference commentary.

Return values

`(object)` Returns housekeeping object.

Examples

Request:

```
{
    "jsonrpc": "2.0",
    "method": "housekeeping.get",
    "params": {
        "output": "extend"
    },
    "id": 1
}
```

Response:


```
{
    "jsonrpc": "2.0",
    "result": {
        "hk_events_mode": "1",
        "hk_events_trigger": "365d",
        "hk_events_service": "1d",
        "hk_events_internal": "1d",
        "hk_events_discovery": "1d",
        "hk_events_autoreg": "1d",
        "hk_services_mode": "1",
        "hk_services": "365d",
        "hk_audit_mode": "1",
        "hk_audit": "31d",
        "hk_sessions_mode": "1",
        "hk_sessions": "365d",
        "hk_history_mode": "1",
        "hk_history_global": "0",
        "hk_history": "31d",
        "hk_trends_mode": "1",
        "hk_trends_global": "0",
        "hk_trends": "365d",
        "db_extension": "",
        "compression_status": "0",
        "compress_older": "7d"
    },
    "id": 1
}
```


Source

CHousekeeping ::get() in *ui/include/classes/api/services/CHousekeeping.php*.

housekeeping.update

Description

`object housekeeping.update(object housekeeping)`

This method allows you to update existing housekeeping settings.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Housekeeping properties to be updated.

Return values

`(array)` Returns an array with the names of updated parameters.

Examples

Request:


```
{
    "jsonrpc": "2.0",
    "method": "housekeeping.update",
    "params": {
        "hk_events_mode": "1",
        "hk_events_trigger": "200d",
        "hk_events_internal": "2d",
        "hk_events_discovery": "2d"
    },
```

```
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        "hk_events_mode",
        "hk_events_trigger",
        "hk_events_internal",
        "hk_events_discovery"
    ],
    "id": 1
}
```

Source

CHousekeeping::update() in *ui/include/classes/api/services/CHousekeeping.php*.

Icon map

This class is designed to work with icon maps.

Object references:

- Icon map
- Icon mapping

Available methods:

- iconmap.create - create new icon maps
- iconmap.delete - delete icon maps
- iconmap.get - retrieve icon maps
- iconmap.update - update icon maps

Icon map object

The following objects are directly related to the `iconmap` API.

Icon map

The icon map object has the following properties.

Property	Type	Description
iconmapid	ID	ID of the icon map. Property behavior: - *read-only* - *required* for update operations
default_iconid	ID	ID of the default icon. Property behavior: - *required* for create operations
name	string	Name of the icon map. Property behavior: - *required* for create operations

Icon mapping

The icon mapping object defines a specific icon to be used for hosts with a certain inventory field value. It has the following properties.

Property	Type	Description
iconid	ID	ID of the icon used by the icon mapping. Property behavior: - *required*
expression	string	Expression to match the inventory field against. Property behavior: - *required*
inventory_link	integer	ID of the host inventory field. Refer to the host inventory object for a list of supported inventory fields. Property behavior: - *required*
sortorder	integer	Position of the icon mapping in the icon map. Property behavior: - *read-only*

iconmap.create

Description

`object iconmap.create(object/array iconMaps)`

This method allows you to create new icon maps.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Icon maps to create.

Additionally to the standard icon map properties, the method accepts the following parameters.

Parameter	Type	Description
mappings	array	Icon mappings to be created for the icon map. Parameter behavior: - *required*

Return values

`(object)` Returns an object containing the IDs of the created icon maps under the `iconmapids` property. The order of the returned IDs matches the order of the passed icon maps.

Examples

Create an icon map

Create an icon map to display hosts of different types.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "iconmap.create",
    "params": {
        "name": "Type icons",
        "default_iconid": "2",
```

```
        "mappings": [
            {
                "inventory_link": 1,
                "expression": "server",
                "iconid": "3"
            },
            {
                "inventory_link": 1,
                "expression": "switch",
                "iconid": "4"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "iconmapids": [
            "2"
        ]
    },
    "id": 1
}
```

See also

- Icon mapping

Source

CIconMap::create() in *ui/include/classes/api/services/CIconMap.php*.

iconmap.delete

Description

`object iconmap.delete(array iconMapIds)`

This method allows you to delete icon maps.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the icon maps to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted icon maps under the `iconmapids` property.

Examples

Delete multiple icon maps

Delete two icon maps.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "iconmap.delete",
    "params": [
        "2",
```

```
        "5"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "iconmapids": [
            "2",
            "5"
        ]
    },
    "id": 1
}
```

Source

CIconMap::delete() in *ui/include/classes/api/services/CIconMap.php*.

iconmap.get

Description

`integer/array iconmap.get(object parameters)`

The method allows you to retrieve icon maps according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
iconmapids	ID/array	Return only icon maps with the given IDs.
sysmapids	ID/array	Return only icon maps that are used in the given maps.
selectMappings	query	Return a `mappings` property with the icon mappings used.
sortfield	string/array	Sort the result by the given properties. Possible values: `iconmapid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;

- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve an icon map

Retrieve all data about icon map "3".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "iconmap.get",
    "params": {
        "iconmapids": "3",
        "output": "extend",
        "selectMappings": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "mappings": [
                {
                    "iconmappingid": "3",
                    "iconmapid": "3",
                    "iconid": "6",
                    "inventory_link": "1",
                    "expression": "server",
                    "sortorder": "0"
                },
                {
                    "iconmappingid": "4",
                    "iconmapid": "3",
                    "iconid": "10",
                    "inventory_link": "1",
                    "expression": "switch",
                    "sortorder": "1"
                }
            ],
            "iconmapid": "3",
            "name": "Host type icons",
            "default_iconid": "2"
        }
    ],
    "id": 1
}
```

See also

- Icon mapping

Source

CIconMap::get() in *ui/include/classes/api/services/CIconMap.php*.

iconmap.update

Description

`object iconmap.update(object/array iconMaps)`

This method allows you to update existing icon maps.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Icon map properties to be updated.

The `iconmapid` property must be defined for each icon map, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard icon map properties, the method accepts the following parameters.

Parameter	Type	Description
mappings	array	Icon mappings to replace the existing icon mappings.

Return values

`(object)` Returns an object containing the IDs of the updated icon maps under the `iconmapids` property.

Examples

Rename icon map

Rename an icon map to "OS icons".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "iconmap.update",
    "params": {
        "iconmapid": "1",
        "name": "OS icons"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "iconmapids": [
            "1"
        ]
    },
    "id": 1
}
```

See also

- Icon mapping

Source

CIconMap::update() in *ui/include/classes/api/services/CIconMap.php*.

Image

This class is designed to work with images.

Object references:

- Image

Available methods:

- image.create - create new images
- image.delete - delete images

- image.get - retrieve images
- image.update - update images

Image object

The following objects are directly related to the `image` API.

Image

The image object has the following properties.

Property	Type	Description
imageid	ID	ID of the image. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the image. Property behavior: - *required* for create operations
imagetype	integer	Type of image. Possible values: 1 - *(default)* icon; 2 - background image. Property behavior: - *constant* - *required* for create operations
image	string	Base64 encoded image. The maximum size of the encoded image is 1 MB. Maximum size can be adjusted by changing `ZBX_MAX_IMAGE_SIZE` constant value. Supported image formats: PNG, JPEG, GIF, and WebP. Property behavior: - *required* for create operations

image.create

Description

`object image.create(object/array images)`

This method allows you to create new images.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Images to create.

The method accepts images with the standard image properties.

Return values

`(object)` Returns an object containing the IDs of the created images under the `imageids` property. The order of the returned IDs matches the order of the passed images.

Examples

Create an icon image

Create a cloud icon.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "image.create",
    "params": {
        "imagetype": 1,
        "name": "Cloud_(24)",
        "image": "iVBORw0KGgoAAAANSUhEUgAAABgAAAANCAYAAACzbK7QAAAABHNCSVQICAgIfAhkiAAAAAlwSFlzAAACmAAAApgB
    },
    "id": 1
}
```


Response:


```
{
    "jsonrpc": "2.0",
    "result": {
        "imageids": [
            "188"
        ]
    },
    "id": 1
}
```


Create a background image

Create a small background image to be used as a pattern.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "image.create",
    "params": {
        "imagetype": 2,
        "name": "background_image",
        "image": "iVBORw0KGgoAAAANSUhEUgAAABgAAAAYCAIAAABvFaqvAAABhWlDQ1BJQ0MgcHJvZmlsZQAAKJF9kb1Lw0AYxp+m
    },
    "id": 1
}
```


Response:


```
{
    "jsonrpc": "2.0",
    "result": {
        "imageids": [
            "189"
        ]
    },
    "id": 1
}
```


Source

CImage::create() in *ui/include/classes/api/services/CImage.php*.

image.delete

Description

`object image.delete(array imageIds)`

This method allows you to delete images.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the images to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted images under the `imageids` property.

Examples

Delete multiple images

Delete two images.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "image.delete",
    "params": [
        "188",
        "192"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "imageids": [
            "188",
            "192"
        ]
    },
    "id": 1
}
```

Source

CImage::delete() in *ui/include/classes/api/services/CImage.php*.

image.get

Description

`integer/array image.get(object parameters)`

The method allows you to retrieve images according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
imageids	ID/array	Return only images with the given IDs.
sysmapids	ID/array	Return images that are used on the given maps.
select_image	flag	Return an `image` property with the Base64 encoded image.

Parameter	Type	Description
sortfield	string/array	Sort the result by the given properties. Possible values: `imageid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve an image

Retrieve all data for image with ID "2".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "image.get",
    "params": {
        "output": "extend",
        "select_image": true,
        "imageids": "2"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "imageid": "2",
            "imagetype": "1",
            "name": "Cloud_(24)",
            "image": "iVBORw0KGgoAAAANSUhEUgAAABgAAAANCAYAAACzbK7QAAAABHNCSVQICAgIfAhkiAAAAAlwSFlzAAACmAAA
        }
    ],
    "id": 1
}
```

Source

CImage::get() in *ui/include/classes/api/services/CImage.php*.

image.update

Description

`object image.update(object/array images)`

This method allows you to update existing images.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Image properties to be updated.

The `imageid` property must be defined for each image, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

The method accepts images with the standard image properties.

Return values

`(object)` Returns an object containing the IDs of the updated images under the `imageids` property.

Examples

Rename image

Rename image to "Cloud icon".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "image.update",
    "params": {
        "imageid": "2",
        "name": "Cloud icon"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "imageids": [
            "2"
        ]
    },
    "id": 1
}
```

Source

CImage::update() in *ui/include/classes/api/services/CImage.php*.

Item

This class is designed to work with items.

Object references:

- Item
  - HTTP header
  - HTTP query field
- Item tag
- Item preprocessing

Available methods:

- item.create - create new items
- item.delete - delete items
- item.get - retrieve items
- item.update - update items

Item object

The following objects are directly related to the `item` API.

Item

Note:
Web items cannot be directly created, updated or deleted via the Zabbix API.

The item object has the following properties.

Property	Type	Description
itemid	ID	ID of the item. Property behavior: - *read-only* - *required* for update operations
delay	string	Update interval of the item. Accepts a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., 30s, 1m, 2h, 1d), or, optionally, one or more custom intervals, all separated by semicolons. Custom intervals can be a mix of flexible and scheduling intervals. Accepts user macros. If used, the value must be a single macro. Multiple macros or macros mixed with text are not supported. Flexible intervals may be written as two macros separated by a forward slash (e.g., `{$FLEX_INTERVAL}/{$FLEX_PERIOD}`). Example: `1h;wd1-5h9-18;{$Macro1}/1-7,00:00-24:00;0/6-7,12:00-24:00;{$Macr` Property behavior: - *required* if `type` is set to ”Zabbix agent” (0), ”Simple check” (3), ”Zabbix internal” (5), ”External check” (10), ”Database monitor” (11), ”IPMI agent” (12), ”SSH agent” (13), ”TELNET agent” (14), ”Calculated” (15), ”JMX agent” (16), ”HTTP agent” (19), ”SNMP agent” (20), ”Script” (21), ”Browser” (22), ”Telemetry query” (24), or if `type` is set to ”Zabbix agent (active)” (7) and `key_` does not contain ”mqtt.get”
hostid	ID	ID of the host or template that the item belongs to. Property behavior: - *constant* - *required* for create operations
interfaceid	ID	ID of the item’s host interface. Property behavior: - *required* if item belongs to host and `type` is set to ”Zabbix agent”, ”IPMI agent”, ”JMX agent”, ”SNMP trap”, or ”SNMP agent” - *supported* if item belongs to host and `type` is set to ”Simple check”, ”External check”, ”SSH agent”, ”TELNET agent”, or ”HTTP agent” - *read-only* for discovered objects
key_	string	Item key. Property behavior: - *required* for create operations - *read-only* for inherited objects or discovered objects

Property	Type	Description
name	string	Name of the item. Supports user macros. Property behavior: - *required* for create operations - *read-only* for inherited objects or discovered objects
name_resolved	string	Name of the item with resolved user macros. Property behavior: - *read-only*
type	integer	Type of the item. Possible values: 0 - Zabbix agent; 2 - Zabbix trapper; 3 - Simple check; 5 - Zabbix internal; 7 - Zabbix agent (active); 9 - Web item; 10 - External check; 11 - Database monitor; 12 - IPMI agent; 13 - SSH agent; 14 - TELNET agent; 15 - Calculated; 16 - JMX agent; 17 - SNMP trap; 18 - Dependent item; 19 - HTTP agent; 20 - SNMP agent; 21 - Script; 22 - Browser; 24 - Telemetry query. Property behavior: - *required* for create operations - *read-only* for inherited objects or discovered objects
url	string	URL string. Supports user macros, {HOST.IP}, {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.NAME}, {HOST.PORT}, {ITEM.ID}, {ITEM.KEY}. Property behavior: - *required* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects
value_type	integer	Type of information of the item. Possible values: 0 - numeric float; 1 - character; 2 - log; 3 - numeric unsigned; 4 - text; 5 - binary; 6 - JSON. Property behavior: - *required* for create operations - *read-only* for inherited objects or discovered objects

Property	Type	Description
allow_traps	integer	Allow the value to be populated similarly to the trapper item. 0 - *(default)* Do not accept incoming data; 1 - Accept incoming data. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for discovered objects
authtype	integer	Authentication method. Possible values if `type` is set to ”SSH agent”: 0 - *(default)* password; 1 - public key. Possible values if `type` is set to ”HTTP agent”: 0 - *(default)* none; 1 - basic; 2 - NTLM; 3 - Kerberos; 4 - Digest. Property behavior: - *supported* if `type` is set to ”SSH agent” or ”HTTP agent” - *read-only* for inherited objects (if `type` is set to ”HTTP agent”) or discovered objects
description	string	Description of the item. Property behavior: - *read-only* for discovered objects
error	string	Error text if there are problems updating the item value. Property behavior: - *read-only*
flags	integer	Origin of the item. Possible values: 0 - a plain item; 4 - an item converted from prototype. Property behavior: - *read-only*
follow_redirects	integer	Follow response redirects while polling data. Possible values: 0 - Do not follow redirects; 1 - *(default)* Follow redirects. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects
headers	array	Array of headers that will be sent when performing an HTTP request. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects
history	string	A time unit of how long the history data should be stored. Also accepts user macro. Default: 31d. Property behavior: - *read-only* for discovered objects

Property	Type	Description
http_proxy	string	HTTP(S) proxy connection string. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects
inventory_link	integer	ID of the host inventory field that is populated by the item. Refer to the host inventory page for a list of supported host inventory fields and their IDs. Default: 0. Property behavior: - *supported* if `value_type` is set to ”numeric float”, ”character”, ”numeric unsigned”, or ”text” - *read-only* for discovered objects
ipmi_sensor	string	IPMI sensor. Property behavior: - *required* if `type` is set to ”IPMI agent” and `key_` is not set to ”ipmi.get” - *supported* if `type` is set to ”IPMI agent” - *read-only* for inherited objects or discovered objects
jmx_endpoint	string	JMX agent custom connection string. Default value: service:jmx:rmi:///jndi/rmi://{HOST.CONN}:{HOST.PORT}/jmxrmi Property behavior: - *supported* if `type` is set to ”JMX agent” - *read-only* for discovered objects
lastclock	timestamp	Time when the item value was last updated. By default, only values that fall within the last 24 hours are displayed. You can extend this time period by changing the value of *Max history display period* parameter in the *Administration > General* menu section. Property behavior: - *read-only*
lastns	integer	Nanoseconds when the item value was last updated. By default, only values that fall within the last 24 hours are displayed. You can extend this time period by changing the value of *Max history display period* parameter in the *Administration > General* menu section. Property behavior: - *read-only*
lastvalue	string	Last value of the item. By default, only values that fall within the last 24 hours are displayed. You can extend this time period by changing the value of *Max history display period* parameter in the *Administration > General* menu section. Property behavior: - *read-only*

Property	Type	Description
logtimefmt	string	Format of the time in log entries. Property behavior: - *supported* if `value_type` is set to ”log” - *read-only* for inherited objects or discovered objects
master_itemid	ID	ID of the master item. Property behavior: - *required* if `type` is set to ”Dependent item” - *read-only* for inherited objects or discovered objects
output_format	integer	Should the response be converted to JSON. 0 - *(default)* Store raw; 1 - Convert to JSON. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects
params	string	Additional parameters depending on the type of the item: - executed script for SSH agent and TELNET agent items; - SQL query for database monitor items; - formula for calculated items; - the script for script and browser items. Property behavior: - *required* if `type` is set to ”Database monitor”, ”SSH agent”, ”TELNET agent”, ”Calculated”, ”Script”, or ”Browser” - *read-only* for inherited objects (if `type` is set to ”Script” or ”Browser”) or discovered objects
parameters	object/array	Additional parameters if `type` is set to ”Script” or ”Browser”. Array of objects with `name` and `value` properties, where `name` must be unique. Property behavior: - *supported* if `type` is set to ”Script” or ”Browser” - *read-only* for inherited objects or discovered objects
password	string	Password for authentication. Property behavior: - *required* if `type` is set to ”JMX agent” and `username` is set - *supported* if `type` is set to ”Simple check”, ”SSH agent”, ”TELNET agent”, ”Database monitor”, or ”HTTP agent” - *read-only* for inherited objects (if `type` is set to ”HTTP agent”) or discovered objects
post_type	integer	Type of post data body stored in `posts` property. Possible values: 0 - *(default)* Raw data; 2 - JSON data; 3 - XML data. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects
posts	string	HTTP(S) request body data. Property behavior: - *required* if `type` is set to ”HTTP agent” and `post_type` is set to ”JSON data” or ”XML data” - *supported* if `type` is set to ”HTTP agent” and `post_type` is set to ”Raw data” - *read-only* for inherited objects or discovered objects

Property	Type	Description
prevvalue	string	Previous value of the item. By default, only values that fall within the last 24 hours are displayed. You can extend this time period by changing the value of *Max history display period* parameter in the *Administration > General* menu section. Property behavior: - *read-only*
privatekey	string	Name of the private key file. Property behavior: - *required* if `type` is set to ”SSH agent” and `authtype` is set to ”public key” - *read-only* for discovered objects
publickey	string	Name of the public key file. Property behavior: - *required* if `type` is set to ”SSH agent” and `authtype` is set to ”public key” - *read-only* for discovered objects
query_fields	array	Array of query fields that will be sent when performing an HTTP request. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects
request_method	integer	Type of request method. Possible values: 0 - *(default)* GET; 1 - POST; 2 - PUT; 3 - HEAD. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects
retrieve_mode	integer	What part of response should be stored. Possible values if `request_method` is set to ”GET”, ”POST”, or ”PUT”: 0 - *(default)* Body; 1 - Headers; 2 - Both body and headers will be stored. Possible values if `request_method` is set to ”HEAD”: 1 - Headers. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects
snmp_oid	string	SNMP OID. Property behavior: - *required* if `type` is set to ”SNMP agent” - *read-only* for inherited objects or discovered objects
ssl_cert_file	string	Public SSL Key file path. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects

Property	Type	Description
ssl_key_file	string	Private SSL Key file path. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects
ssl_key_password	string	Password for SSL Key file. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects
state	integer	State of the item. Possible values: 0 - *(default)* normal; 1 - not supported. Property behavior: - *read-only*
status	integer	Status of the item. Possible values: 0 - *(default)* enabled item; 1 - disabled item.
status_codes	string	Ranges of required HTTP status codes, separated by commas. Also supports user macros as part of comma separated list. Example: 200,200-`{$M}`,`{$M}`,200-400 Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects
templateid	ID	ID of the parent template item. *Hint*: Use the `hostid` property to specify the template that the item belongs to. Property behavior: - *read-only*
timeout	string	Item data polling request timeout. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., `300s`, `1m`). Also accepts user macros. Possible values range: 1-600s. Default: ”” - use proxy/global settings. Property behavior: - *supported* if `type` is set to ”Zabbix agent” (0), ”Simple check” (3) and `key_` does not start with ”`vmware.`” and ”`icmpping`”, ”Zabbix agent (active)” (7), ”External check” (10), ”Database monitor” (11), ”SSH agent” (13), ”TELNET agent” (14), ”HTTP agent” (19), ”SNMP agent” (20) and `snmp_oid` starts with ”`walk[`” or ”`get[`”, ”Script” (21), ”Browser” (22) - *read-only* for inherited and discovered objects, - if the default value is applied and the item belongs to a host monitored by an inaccessible proxy or proxy group, the global timeout is used.

Property	Type	Description
trapper_hosts	string	Allowed hosts. Property behavior: - *readonly* for discovered objects - *supported* if `type` is set to ”Zabbix trapper”, or if `type` is set to ”HTTP agent” and `allow_traps` is set to ”Allow incoming data”. If the property is omitted, it is automatically set to the `{$TRAPPER.ALLOWED_HOSTS}` user macro (default value). By default, `{$TRAPPER.ALLOWED_HOSTS}` is `127.0.0.1,::1`. If the property is set to an empty string, incoming connections will be denied from all hosts.
trends	string	A time unit of how long the trends data should be stored. Also accepts user macro. Default: 365d. Property behavior: - *supported* if `value_type` is set to ”numeric float” or ”numeric unsigned” - *read-only* for discovered objects
units	string	Value units. Property behavior: - *supported* if `value_type` is set to ”numeric float” or ”numeric unsigned” - *read-only* for inherited objects or discovered objects
username	string	Username for authentication. Property behavior: - *required* if `type` is set to ”SSH agent”, ”TELNET agent”, or if `type` is set to ”JMX agent” and `password` is set - *supported* if `type` is set to ”Simple check”, ”Database monitor”, or ”HTTP agent” - *read-only* for inherited objects (if `type` is set to ”HTTP agent”) or discovered objects
uuid	string	Universal unique identifier, used for linking imported item to already existing ones. Auto-generated, if not given. Property behavior: - *supported* if the item belongs to a template
valuemapid	ID	ID of the associated value map. Property behavior: - *supported* if `value_type` is set to ”numeric float”, ”character”, or ”numeric unsigned” - *read-only* for inherited objects or discovered objects
verify_host	integer	Whether to validate that the host name for the connection matches the one in the host’s certificate. Possible values: 0 - *(default)* Do not validate; 1 - Validate. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects or discovered objects

Property	Type	Description
verify_peer	integer	Whether to validate that the host's certificate is authentic. Possible values: 0 - *(default)* Do not validate; 1 - Validate. Property behavior: - *supported* if `type` is set to "HTTP agent" - *read-only* for inherited objects or discovered objects
query	object	Telemetry query configuration. Defines what the item queries from the telemetry data source, how it filters the data, and how it aggregates the result. Property behavior: - *required* if `type` is set to "Telemetry query" - *read-only* for inherited objects or discovered objects
time_shift	string	How far back (before the current time) ends the time range that the item queries. Accepts a number of seconds (e.g., 1, 60, 300) or a value with a time suffix (e.g., `30s`, `1m`, `2h`, `1d`). Also accepts user macros. Possible values range: 0-86400s. Default: 15s. Property behavior: - *supported* if `type` is set to "Telemetry query"
lookback_limit	string	How far back (before the current time) the item queries for data when there is no previous query to continue from (e.g., on the item's first run), or when the previous query is older than `delay` (e.g., because Zabbix server was down). Accepts a number of seconds (e.g., 1, 60, 300) or a value with a time suffix (e.g., `30s`, `1m`, `2h`, `1d`). Also accepts user macros. Possible values range: 1-259200s. Default: 10m. Property behavior: - *supported* if `type` is set to "Telemetry query"
granularity	string	The length of a piece that the query's full time range is split into. Accepts a number of seconds (e.g., 1, 60, 300) or a value with a time suffix (e.g., `30s`, `1m`, `2h`, `1d`). Also accepts user macros. Possible values range: 1-86400s. Default: 15s. Property behavior: - *supported* if `type` is set to "Telemetry query"

HTTP header

The header object has the following properties:

Property	Type	Description
name	string	HTTP header name. Property behavior: - *required*
value	string	Header value. Property behavior: - *required*

HTTP query field

The query field object defines a name and value that is used to specify a URL parameter. It has the following properties:

Property	Type	Description
name	string	Name of the parameter. Property behavior: - *required*
value	string	Parameter value. Property behavior: - *required*

Telemetry query

The telemetry query object has the following properties.

Property	Type	Description
signal_type	integer	Signal type that the item queries. Possible values: 0 - Traces; 1 - Metrics; 2 - Logs. Property behavior: - *required*
metric_point_type	integer	Metric point type that the item queries. Possible values: 0 - Sum; 1 - Gauge; 2 - Histogram; 3 - Exponential histogram. Property behavior: - *required* if `signal_type` is set to ”Metrics”
columns	array	Array of telemetry data columns that the item queries. Can be empty. Property behavior: - *required*
aggregated_columns	array	Array of telemetry data aggregated columns that the item queries, and how it calculates a value from column records. Property behavior: - *required*
filter	object	Telemetry query filter for column records in the query. Property behavior: - *required*

Telemetry data columns

The telemetry data columns object has the following properties.

Property	Type	Description
column	string	Name of the telemetry data column that the item queries. See supported telemetry data columns. For multiple single-value columns, each `column` value must be unique. For multiple columns that hold key-value pairs, each `column` and `attribute_key` pair must be unique.
attribute_key	string	Name of the key that the item queries, within the column that holds key-value pairs. For multiple columns that hold key-value pairs, each `column` and `attribute_key` pair must be unique. Property behavior: - *required* if `column` is set to a column that holds key-value pairs

Telemetry data aggregated columns

The telemetry data aggregated columns object has the following properties.

Property	Type	Description
column	string	Name of the telemetry data column that the item queries. See supported telemetry data columns. For multiple columns, each `column` value must be unique. This property is not supported if `function` is set to ”Count”. Property behavior: - *required* if `function` is set to ”Min”, ”Max”, ”Average”, ”Sum”, or ”Percentile”
function	integer	Function that the item applies to the column records. Possible values: 1 - Min; 2 - Max; 3 - Average; 4 - Count; 5 - Sum; 8 - Percentile. Property behavior: - *required*
parameters	array	Percentile that the item calculates from the column records. Possible value: a number from 0 to 100, given as a string. Property behavior: - *required* if `function` is set to ”Percentile”
alias	string	Unique name for the calculated value, used to identify it in the item value. Each `alias` value must be unique. Property behavior: - *required*

Telemetry data filter

The telemetry data filter object has the following properties.

Property	Type	Description
conditions	array	Telemetry data filter conditions a record in a telemetry data column must meet. Property behavior: - *required*
evaltype	integer	Filter condition evaluation method. Possible values: 0 - And/Or; 1 - And; 2 - Or; 3 - Custom expression. Property behavior: - *required*
eval_formula	string	Generated expression that will be used for evaluating filter conditions. The expression contains IDs that reference specific filter conditions by its `formulaid`. The value of `eval_formula` is equal to the value of `formula` for filters with a custom expression. Property behavior: - *read-only*
formula	string	User-defined expression to be used for evaluating conditions of filters with a custom expression. The expression must contain IDs that reference specific filter conditions by its `formulaid`. The IDs used in the expression must exactly match the ones defined in the filter conditions: no condition can remain unused or omitted. Property behavior: - *required* if `evaltype` is set to ”Custom expression”

Telemetry data filter conditions

The telemetry data filter conditions object has the following properties.

Property	Type	Description
column	string	Name of the telemetry data column that the filter condition applies to. See supported telemetry data columns. Property behavior: - *required*
attribute_key	string	Name of the key that the condition applies to, within the column that holds key-value pairs. Property behavior: - *required* if `column` is set to a column that holds key-value pairs

Property	Type	Description
operator	integer	Filter condition operator. Possible values for single-value columns: 0 - Equals; 1 - Does not equal; 2 - Contains; 3 - Does not contain. Possible values for columns that hold key-value pairs: 0 - Equals; 1 - Does not equal; 12 - Exists. Property behavior: - *required*
value	string	Value to compare with. Property behavior: - *required* if `operator` is set to ”Equals”, ”Does not equal”, ”Contains” or ”Does not contain”
formulaid	string	Arbitrary unique ID that is used to reference the condition from a custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward. Property behavior: - *required* if `evaltype` of Telemetry query filter object is set to ”Custom expression”

Item tag

The item tag object has the following properties.

Property	Type	Description
tag	string	Item tag name. Property behavior: - *required*
value	string	Item tag value. Property behavior: - *read-only*.
object	integer	Type of object the tag was inherited from. Possible values: 0 - template; 1 - host. Property behavior: - *read-only*.
objectid	ID	ID of the object the tag was inherited from. Property behavior: - *read-only*.

Item preprocessing

The item preprocessing object has the following properties.

Property	Type	Description
type	integer	The preprocessing option type. Possible values: 1 - Custom multiplier; 2 - Right trim; 3 - Left trim; 4 - Trim; 5 - Regular expression; 6 - Boolean to decimal; 7 - Octal to decimal; 8 - Hexadecimal to decimal; 9 - Simple change; 10 - Change per second; 11 - XML XPath; 12 - JSONPath; 13 - In range; 14 - Matches regular expression; 15 - Does not match regular expression; 16 - Check for error in JSON; 17 - Check for error in XML; 18 - Check for error using regular expression; 19 - Discard unchanged; 20 - Discard unchanged with heartbeat; 21 - JavaScript; 22 - Prometheus pattern; 23 - Prometheus to JSON; 24 - CSV to JSON; 25 - Replace; 26 - Check unsupported; 27 - XML to JSON; 28 - SNMP walk value; 29 - SNMP walk to JSON; 30 - SNMP get value. Property behavior: - *required*
params	string	Additional parameters used by preprocessing option. Multiple parameters are separated by the newline (\n) character. If `type` is set to ”Check unsupported”, the parameters follow a `<scope>[\n<pattern>]` syntax, where *pattern* is a regular expression, and *scope* is one of: -1 - match any error; 0 - check if error message matches *pattern*; 1 - check if error message does not match *pattern*. Property behavior: - *required* if `type` is set to ”Custom multiplier” (1), ”Right trim” (2), ”Left trim” (3), ”Trim” (4), ”Regular expression” (5), ”XML XPath” (11), ”JSONPath” (12), ”In range” (13), ”Matches regular expression” (14), ”Does not match regular expression” (15), ”Check for error in JSON” (16), ”Check for error in XML” (17), ”Check for error using regular expression” (18), ”Discard unchanged with heartbeat” (20), ”JavaScript” (21), ”Prometheus pattern” (22), ”Prometheus to JSON” (23), ”CSV to JSON” (24), ”Replace” (25), Check unsupported (26), ”SNMP walk value” (28), ”SNMP walk to JSON” (29), or ”SNMP get value” (30)

Property	Type	Description
error_handler	integer	Action type used in case of preprocessing step failure. Possible values: 0 - Error message is set by Zabbix server; 1 - Discard value; 2 - Set custom value; 3 - Set custom error message. Possible values if `type` is set to ”Check unsupported”: 1 - Discard value; 2 - Set custom value; 3 - Set custom error message. Property behavior: - *required* if `type` is set to ”Custom multiplier” (1), ”Regular expression” (5), ”Boolean to decimal” (6), ”Octal to decimal” (7), ”Hexadecimal to decimal” (8), ”Simple change” (9), ”Change per second” (10), ”XML XPath” (11), ”JSONPath” (12), ”In range” (13), ”Matches regular expression” (14), ”Does not match regular expression” (15), ”Check for error in JSON” (16), ”Check for error in XML” (17), ”Check for error using regular expression” (18), ”Prometheus pattern” (22), ”Prometheus to JSON” (23), ”CSV to JSON” (24), ”Check unsupported” (26), ”XML to JSON” (27), ”SNMP walk value” (28), ”SNMP walk to JSON” (29), or ”SNMP get value” (30)
error_handler_params	string	Error handler parameters. Property behavior: - *required* if `error_handler` is set to ”Set custom value” or ”Set custom error message”

The following parameters and error handlers are supported for each preprocessing type.

Preprocessing type	Name	Parameter 1	Parameter 2	Parameter 3	Supported error handlers
1	Custom multiplier	number^1, 6^			0, 1, 2, 3
2	Right trim	list of characters[2]			
3	Left trim	list of characters[2]			
4	Trim	list of characters[2]			
5	Regular expression	pattern[3]	output[2]		0, 1, 2, 3
6	Boolean to decimal				0, 1, 2, 3
7	Octal to decimal				0, 1, 2, 3
8	Hexadecimal to decimal				0, 1, 2, 3

Preprocessing type	Name	Parameter 1	Parameter 2	Parameter 3	Supported error handlers
9	Simple change				0, 1, 2, 3
10	Change per second				0, 1, 2, 3
11	XML XPath	path[4]			0, 1, 2, 3
12	JSONPath	path[4]			0, 1, 2, 3
13	In range	min^1, 6^	max^1, 6^		0, 1, 2, 3
14	Matches regular expression	pattern[3]			0, 1, 2, 3
15	Does not match regular expression	pattern[3]			0, 1, 2, 3
16	Check for error in JSON	path[4]			0, 1, 2, 3
17	Check for error in XML	path[4]			0, 1, 2, 3
18	Check for error using regular expression	pattern[3]	output[2]		0, 1, 2, 3
19	Discard unchanged				
20	Discard unchanged with heartbeat	seconds^5, 6^			
21	JavaScript	script[2]			
22	Prometheus pattern	pattern^6, 7^	`value`, `label`, `function`	output^8, 9^	0, 1, 2, 3
23	Prometheus to JSON	pattern^6, 7^			0, 1, 2, 3

Preprocessing type	Name	Parameter 1	Parameter 2	Parameter 3	Supported error handlers
24	CSV to JSON	character[2]	character[2]	0,1	0, 1, 2, 3
25	Replace	search string[2]	replacement[2]		
26	Check unsupported	scope[1]	pattern^3, 6^		1, 2, 3
27	XML to JSON				0, 1, 2, 3
28	SNMP walk value	OID[2]	Format: 0 - Unchanged 1 - UTF-8 from Hex-STRING 2 - MAC from Hex-STRING 3 - Integer from BITS		0, 1, 2, 3
29	SNMP walk to JSON[10]	Field name[2]	OID prefix[2]	Format: 0 - Unchanged 1 - UTF-8 from Hex-STRING 2 - MAC from Hex-STRING 3 - Integer from BITS	0, 1, 2, 3
30	SNMP get value	Format: 1 - UTF-8 from Hex-STRING 2 - MAC from Hex-STRING 3 - Integer from BITS			0, 1, 2, 3

[1] floating-point number (integers are converted to float values implicitly)
 [2] string
 [3] regular expression
 [4] JSONPath or XML XPath
 [5] positive integer (with support of time suffixes, e.g., 30s, 1m, 2h, 1d)
 [6] user macro
 [7] Prometheus pattern following the syntax: `<metric name>{<label name>="<label value>", ...} == <value>`. Each Prometheus pattern component (metric, label name, label value and metric value) can be user macro.
 [8] Prometheus output following the syntax: `<label name>` (can be a user macro) if `label` is selected as the second parameter.
 [9] One of the aggregation functions: `sum`, `min`, `max`, `avg`, `count` if `function` is selected as the second parameter.
 [10] Supports multiple "Field name,OID prefix,Format records" records delimited by a new line character.

item.create

Description

`object item.create(object/array items)`

This method allows you to create new items.

> **Note:**
> Web items cannot be created via the Zabbix API.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Items to create.

Additionally to the standard item properties, the method accepts the following parameters.

Parameter	Type	Description
preprocessing	array	Item preprocessing options.
tags	array	Item tags.

Return values

(object) Returns an object containing the IDs of the created items under the itemids property. The order of the returned IDs matches the order of the passed items.

Examples

Creating an item

Create a numeric Zabbix agent item with 2 item tags to monitor free disk space on host with ID "30074".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.create",
    "params": {
        "name": "Free disk space on /home/joe/",
        "key_": "vfs.fs.size[/home/joe/,free]",
        "hostid": "30074",
        "type": 0,
        "value_type": 3,
        "interfaceid": "30084",
        "tags": [
            {
                "tag": "component",
                "value": "storage"
            },
            {
                "tag": "equipment",
                "value": "workstation"
            }
        ],
        "delay": "30s"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "24758"
        ]
    },
    "id": 1
}
```

Creating a host inventory item

Create a Zabbix agent item to populate the host's "OS" inventory field.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.create",
    "params": {
        "name": "uname",
```

```
        "key_": "system.uname",
        "hostid": "30021",
        "type": 0,
        "interfaceid": "30007",
        "value_type": 1,
        "delay": "10s",
        "inventory_link": 5
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "24759"
        ]
    },
    "id": 1
}
```

Creating an item with preprocessing

Create an item using custom multiplier.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.create",
    "params": {
        "name": "Device uptime",
        "key_": "sysUpTime",
        "hostid": "10084",
        "type": 20,
        "snmp_oid": "SNMPv2-MIB::sysUpTime.0",
        "value_type": 1,
        "delay": "60s",
        "interfaceid": "83",
        "preprocessing": [
            {
                "type": 1,
                "params": "0.01",
                "error_handler": 1,
                "error_handler_params": ""
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "44210"
        ]
    },
    "id": 1
}
```

Creating dependent item

Create a dependent item for the master item with ID 24759. Only dependencies on the same host are allowed, therefore master and the dependent item should have the same hostid.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.create",
    "params": {
      "hostid": "30074",
      "name": "Dependent test item",
      "key_": "dependent.item",
      "type": 18,
      "master_itemid": "24759",
      "value_type": 2
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "44211"
        ]
    },
    "id": 1
}
```

Create HTTP agent item

Create POST request method item with JSON response preprocessing.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.create",
    "params": {
        "url":"http://127.0.0.1/http.php",
        "query_fields": [
            {
                "name": "mode",
                "value": "json"
            },
            {
                "name": "min",
                "value": "10"
            },
            {
                "name": "max",
                "value": "100"
            }
        ],
        "interfaceid": "1",
        "type": 19,
        "hostid": "10254",
        "delay": "5s",
        "key_": "json",
        "name": "HTTP agent example JSON",
        "value_type": 0,
        "output_format": 1,
        "preprocessing": [
            {
                "type": 12,
```

```
                "params": "$.random",
                "error_handler": 0,
                "error_handler_params": ""
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "23865"
        ]
    },
    "id": 1
}
```

Create script item

Create a simple data collection using a script item.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.create",
    "params": {
        "name": "Script example",
        "key_": "custom.script.item",
        "hostid": "12345",
        "type": 21,
        "value_type": 4,
        "params": "var request = new HttpRequest();\nreturn request.post(\"https://postman-echo.com/post\"
        "parameters": [
            {
                "name": "host",
                "value": "{HOST.CONN}"
            }
        ],
        "timeout": "6s",
        "delay": "30s"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "23865"
        ]
    },
    "id": 1
}
```

Source

CItem::create() in *ui/include/classes/api/services/CItem.php*.

item.delete

Description

`object item.delete(array itemIds)`

This method allows you to delete items.

> **Note:**
> Web items cannot be deleted via the Zabbix API.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the items to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted items under the `itemids` property.

Examples

Deleting multiple items

Delete two items.
 Dependent items and item prototypes are removed automatically if master item is deleted.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.delete",
    "params": [
        "22982",
        "22986"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "22982",
            "22986"
        ]
    },
    "id": 1
}
```

Source

CItem::delete() in *ui/include/classes/api/services/CItem.php*.

item.get

Description

`integer/array item.get(object parameters)`

The method allows you to retrieve items according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
itemids	ID/array	Return only items with the given IDs.
groupids	ID/array	Return only items that belong to the hosts from the given groups.
templateids	ID/array	Return only items that belong to the given templates.
hostids	ID/array	Return only items that belong to the given hosts.
proxyids	ID/array	Return only items that are monitored by the given proxies.
interfaceids	ID/array	Return only items that use the given host interfaces.
graphids	ID/array	Return only items that are used in the given graphs.
triggerids	ID/array	Return only items that are used in the given triggers.
webitems	flag	Include web items in the result.
inherited	boolean	If set to `true` return only items inherited from a template.
inheritedTags	boolean	Return items that have given `tags` also in host/template/linked templates. Possible values: `true` - template/host/linked templates must also have given tags; `false` - *(default)* tags from template/host/linked templates are ignored.
templated	boolean	If set to `true` return only items that belong to templates.
monitored	boolean	If set to `true` return only enabled items that belong to monitored hosts.
group	string	Return only items that belong to a group with the given name.
host	string	Return only items that belong to a host with the given name.
evaltype	integer	Tag evaluation method. Possible values: 0 - *(default)* And/Or; 2 - Or.
tags	array	Return only items with the given tags. Format: `[{"tag": "<tag>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all items. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
with_triggers	boolean	If set to `true` return only items that are used in triggers.
selectHosts	query	Return a `hosts` property with an array of hosts that the item belongs to.
selectInterfaces	query	Return an `interfaces` property with an array of host interfaces used by the item.
selectTriggers	query	Return a `triggers` property with the triggers that the item is used in. Supports `count`.
selectGraphs	query	Return a `graphs` property with the graphs that contain the item. Supports `count`.

Parameter	Type	Description
selectDiscoveryData	query	Return a `discoveryData` property with the item discovery object data. The item discovery object links a discovered item to an item prototype from which it was discovered. It has the following properties: `parent_itemid` - `(string)` ID of the item prototype from which the item has been created; `key_` - `(string)` key of the item prototype; `status` - `(int)` item discovery status: 0 - *(default)* item is discovered, 1 - item is not discovered anymore; `ts_delete` - `(timestamp)` time when an item that is no longer discovered will be deleted; `ts_disable` - `(timestamp)` time when an item that is no longer discovered will be disabled; `disable_source` - `(int)` indicator of whether item was disabled by an LLD rule or manually: 0 - *(default)* disabled automatically, 1 - disabled by an LLD rule.
selectDiscoveryRule	query	Return a `discoveryRule` property with the LLD rule that created the item.
selectInheritedTags	query	Return an `inheritedTags` property with tags that are on template/host/linked templates.
selectPreprocessing	query	Return a `preprocessing` property with item preprocessing options.
selectTags	query	Return the item tags in `tags` property.
selectValueMap	query	Return a `valuemap` property with item value map.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type. Supports additional properties: `host` - technical name of the host that the item belongs to.
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectGraphs` - results will be sorted by `name`; `selectTriggers` - results will be sorted by `description`.
sortfield	string/array	Sort the result by the given properties. Possible values: `itemid`, `name`, `key_`, `delay`, `history`, `trends`, `type`, `status`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	
selectItemDiscovery	query	Return an `itemDiscovery` property with the item discovery object. The item discovery object links the item to an item prototype from which it was created. This query is **deprecated**, please use `selectDiscoveryData` instead.

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Finding items by key

Retrieve all items used in triggers for specific host ID that have word ”system.cpu” in the item key and sort results by name.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.get",
    "params": {
        "output": "extend",
        "hostids": "10084",
        "with_triggers": true,
        "search": {
            "key_": "system.cpu"
        },
        "sortfield": "name"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "itemid": "42269",
            "type": "18",
            "snmp_oid": "",
            "hostid": "10084",
            "name": "CPU utilization",
            "key_": "system.cpu.util",
            "delay": "0",
            "history": "7d",
            "trends": "365d",
            "status": "0",
            "value_type": "0",
            "trapper_hosts": "",
            "units": "%",
            "formula": "",
            "logtimefmt": "",
            "templateid": "42267",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "0",
            "interfaceid": "0",
            "description": "CPU utilization in %.",
            "inventory_link": "0",
            "lifetime": "7d",
            "evaltype": "0",
            "jmx_endpoint": "",
```

```
            "master_itemid": "42264",
            "timeout": "",
            "url": "",
            "query_fields": [],
            "posts": "",
            "status_codes": "200",
            "follow_redirects": "1",
            "post_type": "0",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "0",
            "request_method": "0",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
            "state": "0",
            "error": "",
            "name_resolved": "CPU utilization",
            "parameters": [],
            "lastclock": "0",
            "lastns": "0",
            "lastvalue": "0",
            "prevvalue": "0"
        },
        {
            "itemid": "42259",
            "type": "0",
            "snmp_oid": "",
            "hostid": "10084",
            "name": "Load average (15m avg)",
            "key_": "system.cpu.load[all,avg15]",
            "delay": "1m",
            "history": "7d",
            "trends": "365d",
            "status": "0",
            "value_type": "0",
            "trapper_hosts": "",
            "units": "",
            "formula": "",
            "logtimefmt": "",
            "templateid": "42219",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "0",
            "interfaceid": "1",
            "description": "",
            "inventory_link": "0",
            "lifetime": "7d",
```

```
            "evaltype": "0",
            "jmx_endpoint": "",
            "master_itemid": "0",
            "timeout": "",
            "url": "",
            "query_fields": [],
            "posts": "",
            "status_codes": "200",
            "follow_redirects": "1",
            "post_type": "0",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "0",
            "request_method": "0",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
            "state": "0",
            "error": "",
            "name_resolved": "Load average (15m avg)",
            "parameters": [],
            "lastclock": "0",
            "lastns": "0",
            "lastvalue": "0",
            "prevvalue": "0"
        },
        {
            "itemid": "42249",
            "type": "0",
            "snmp_oid": "",
            "hostid": "10084",
            "name": "Load average (1m avg)",
            "key_": "system.cpu.load[all,avg1]",
            "delay": "1m",
            "history": "7d",
            "trends": "365d",
            "status": "0",
            "value_type": "0",
            "trapper_hosts": "",
            "units": "",
            "formula": "",
            "logtimefmt": "",
            "templateid": "42209",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "0",
            "interfaceid": "1",
            "description": "",
```

```
            "inventory_link": "0",
            "lifetime": "7d",
            "evaltype": "0",
            "jmx_endpoint": "",
            "master_itemid": "0",
            "timeout": "",
            "url": "",
            "query_fields": [],
            "posts": "",
            "status_codes": "200",
            "follow_redirects": "1",
            "post_type": "0",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "0",
            "request_method": "0",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
            "state": "0",
            "error": "",
            "name_resolved": "Load average (1m avg)",
            "parameters": [],
            "lastclock": "0",
            "lastns": "0",
            "lastvalue": "0",
            "prevvalue": "0"
        },
        {
            "itemid": "42257",
            "type": "0",
            "snmp_oid": "",
            "hostid": "10084",
            "name": "Load average (5m avg)",
            "key_": "system.cpu.load[all,avg5]",
            "delay": "1m",
            "history": "7d",
            "trends": "365d",
            "status": "0",
            "value_type": "0",
            "trapper_hosts": "",
            "units": "",
            "formula": "",
            "logtimefmt": "",
            "templateid": "42217",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "0",
```

```
        "interfaceid": "1",
        "description": "",
        "inventory_link": "0",
        "lifetime": "7d",
        "evaltype": "0",
        "jmx_endpoint": "",
        "master_itemid": "0",
        "timeout": "",
        "url": "",
        "query_fields": [],
        "posts": "",
        "status_codes": "200",
        "follow_redirects": "1",
        "post_type": "0",
        "http_proxy": "",
        "headers": [],
        "retrieve_mode": "0",
        "request_method": "0",
        "output_format": "0",
        "ssl_cert_file": "",
        "ssl_key_file": "",
        "ssl_key_password": "",
        "verify_peer": "0",
        "verify_host": "0",
        "allow_traps": "0",
        "uuid": "",
        "lifetime_type": "0",
        "enabled_lifetime_type": "2",
        "enabled_lifetime": "0",
        "state": "0",
        "error": "",
        "name_resolved": "Load average (5m avg)",
        "parameters": [],
        "lastclock": "0",
        "lastns": "0",
        "lastvalue": "0",
        "prevvalue": "0"
    },
    {
        "itemid": "42260",
        "type": "0",
        "snmp_oid": "",
        "hostid": "10084",
        "name": "Number of CPUs",
        "key_": "system.cpu.num",
        "delay": "1m",
        "history": "7d",
        "trends": "365d",
        "status": "0",
        "value_type": "3",
        "trapper_hosts": "",
        "units": "",
        "formula": "",
        "logtimefmt": "",
        "templateid": "42220",
        "valuemapid": "0",
        "params": "",
        "ipmi_sensor": "",
        "authtype": "0",
        "username": "",
        "password": "",
        "publickey": "",
```

```
            "privatekey": "",
            "flags": "0",
            "interfaceid": "1",
            "description": "",
            "inventory_link": "0",
            "lifetime": "7d",
            "evaltype": "0",
            "jmx_endpoint": "",
            "master_itemid": "0",
            "timeout": "",
            "url": "",
            "query_fields": [],
            "posts": "",
            "status_codes": "200",
            "follow_redirects": "1",
            "post_type": "0",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "0",
            "request_method": "0",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
            "state": "0",
            "error": "",
            "name_resolved": "Number of CPUs",
            "parameters": [],
            "lastclock": "0",
            "lastns": "0",
            "lastvalue": "0",
            "prevvalue": "0"
        }
    ],
    "id": 1
}
```


Finding dependent items by key

Retrieve all dependent items from host with ID ”10116” that have the word ”apache” in the key.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.get",
    "params": {
        "output": "extend",
        "hostids": "10116",
        "search": {
            "key_": "apache"
        },
        "filter": {
            "type": 18
        }
    },
    "id": 1
```

```
}
```

Response:


```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "itemid": "25550",
            "type": "18",
            "snmp_oid": "",
            "hostid": "10116",
            "name": "Days",
            "key_": "apache.status.uptime.days",
            "delay": "0",
            "history": "90d",
            "trends": "365d",
            "status": "0",
            "value_type": "3",
            "trapper_hosts": "",
            "units": "",
            "formula": "",
            "logtimefmt": "",
            "templateid": "0",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "0",
            "interfaceid": "0",
            "description": "",
            "inventory_link": "0",
            "lifetime": "7d",
            "evaltype": "0",
            "jmx_endpoint": "",
            "master_itemid": "25545",
            "timeout": "",
            "url": "",
            "query_fields": [],
            "posts": "",
            "status_codes": "200",
            "follow_redirects": "1",
            "post_type": "0",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "0",
            "request_method": "0",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
            "state": "0",
```

```
            "error": "",
            "name_resolved": "Days",
            "parameters": [],
            "lastclock": "0",
            "lastns": "0",
            "lastvalue": "0",
            "prevvalue": "0"
        },
        {
            "itemid": "25555",
            "type": "18",
            "snmp_oid": "",
            "hostid": "10116",
            "name": "Hours",
            "key_": "apache.status.uptime.hours",
            "delay": "0",
            "history": "90d",
            "trends": "365d",
            "status": "0",
            "value_type": "3",
            "trapper_hosts": "",
            "units": "",
            "formula": "",
            "logtimefmt": "",
            "templateid": "0",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "0",
            "interfaceid": "0",
            "description": "",
            "inventory_link": "0",
            "lifetime": "7d",
            "evaltype": "0",
            "jmx_endpoint": "",
            "master_itemid": "25545",
            "timeout": "",
            "url": "",
            "query_fields": [],
            "posts": "",
            "status_codes": "200",
            "follow_redirects": "1",
            "post_type": "0",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "0",
            "request_method": "0",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
```

```
            "enabled_lifetime": "0",
            "state": "0",
            "error": "",
            "name_resolved": "Hours",
            "parameters": [],
            "lastclock": "0",
            "lastns": "0",
            "lastvalue": "0",
            "prevvalue": "0"
        }
    ],
    "id": 1
}
```

Find HTTP agent item

Find HTTP agent item with post body type XML for specific host ID.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.get",
    "params": {
        "hostids": "10255",
        "filter": {
            "type": 19,
            "post_type": 3
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "itemid": "28252",
            "type": "19",
            "snmp_oid": "",
            "hostid": "10255",
            "name": "template item",
            "key_": "ti",
            "delay": "30s",
            "history": "90d",
            "trends": "365d",
            "status": "0",
            "value_type": "3",
            "trapper_hosts": "",
            "units": "",
            "formula": "",
            "logtimefmt": "",
            "templateid": "0",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "0",
            "interfaceid": "0",
```

```
            "description": "",
            "inventory_link": "0",
            "lifetime": "7d",
            "evaltype": "0",
            "jmx_endpoint": "",
            "master_itemid": "0",
            "timeout": "",
            "url": "localhost",
            "query_fields": [
                {
                    "name": "mode",
                    "value": "xml"
                }
            ],
            "posts": "<body>\r\n<![CDATA[{$MACRO}<foo></bar>]]>\r\n</body>",
            "status_codes": "200",
            "follow_redirects": "0",
            "post_type": "3",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "1",
            "request_method": "3",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
            "state": "0",
            "error": "",
            "name_resolved": "template item",
            "parameters": [],
            "lastclock": "0",
            "lastns": "0",
            "lastvalue": "",
            "prevvalue": ""
        }
    ],
    "id": 1
}
```


Retrieving items with preprocessing rules

Retrieve all items and their preprocessing rules for specific host ID.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "item.get",
    "params": {
        "output": ["itemid", "name", "key_"],
        "selectPreprocessing": "extend",
        "hostids": "10254"
    },
    "id": 1
}
```


Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemid": "23865",
        "name": "HTTP agent example JSON",
        "key_": "json",
        "preprocessing": [
            {
                "type": "12",
                "params": "$.random",
                "error_handler": "1",
                "error_handler_params": ""
            }
        ]
    },
    "id": 1
}
```


See also

- Discovery rule
- Graph
- Host
- Host interface
- Trigger

Source

CItem::get() in *ui/include/classes/api/services/CItem.php*.

item.update

Description

`object item.update(object/array items)`

This method allows you to update existing items.

Note:
Web items cannot be updated via the Zabbix API.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Item properties to be updated.

The `itemid` property must be defined for each item, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard item properties, the method accepts the following parameters.

Parameter	Type	Description
preprocessing	array	Item preprocessing options to replace the current preprocessing options. Parameter behavior: - *read-only* for inherited objects or discovered objects
tags	array	Item tags. Parameter behavior: - *read-only* for discovered objects

Return values

(object) Returns an object containing the IDs of the updated items under the itemids property.

Examples

Enabling an item

Enable an item, that is, set its status to "0".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.update",
    "params": {
        "itemid": "10092",
        "status": 0
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "10092"
        ]
    },
    "id": 1
}
```

Update dependent item

Update Dependent item name and Master item ID. Only dependencies on same host are allowed, therefore Master and Dependent item should have same hostid.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.update",
    "params": {
        "name": "Dependent item updated name",
        "master_itemid": "25562",
        "itemid": "189019"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "189019"
        ]
    },
    "id": 1
}
```

Update HTTP agent item

Enable item value trapping.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.update",
    "params": {
        "itemid": "23856",
        "allow_traps": 1
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "23856"
        ]
    },
    "id": 1
}
```

Updating an item with preprocessing

Update an item with item preprocessing rule ”In range”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.update",
    "params": {
        "itemid": "23856",
        "preprocessing": [
            {
                "type": 13,
                "params": "\n100",
                "error_handler": 1,
                "error_handler_params": ""
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "23856"
        ]
    },
    "id": 1
}
```

Updating a script item

Update a script item with a different script and remove unnecessary parameters that were used by previous script.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "item.update",
    "params": {
        "itemid": "23865",
```

```
        "parameters": [],
        "params": "Zabbix.log(3, 'Log test');\nreturn 1;"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "23865"
        ]
    },
    "id": 1
}
```

Source

CItem::update() in *ui/include/classes/api/services/CItem.php*.

Item prototype

This class is designed to work with item prototypes.

Object references:

- Item prototype
  - HTTP header
  - HTTP query field
- Item prototype tag
- Item prototype preprocessing

Available methods:

- itemprototype.create - create new item prototypes
- itemprototype.delete - delete item prototypes
- itemprototype.get - retrieve item prototypes
- itemprototype.update - update item prototypes

Item prototype object

The following objects are directly related to the `itemprototype` API.

Item prototype

The item prototype object has the following properties.

Property	Type	Description
itemid	ID	ID of the item prototype. Property behavior: - *read-only* - *required* for update operations

Property	Type	Description
delay	string	Update interval of the item prototype. Accepts a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., 30s, 1m, 2h, 1d), or, optionally, one or more custom intervals, all separated by semicolons. Custom intervals can be a mix of flexible and scheduling intervals. Accepts user macros and LLD macros. If used, the value must be a single macro. Multiple macros or macros mixed with text are not supported. Flexible intervals may be written as two macros separated by a forward slash (e.g., {$FLEX_INTERVAL}/{$FLEX_PERIOD}). Example: 1h;wd1-5h9-18;{$Macro1}/1-7,00:00-24:00;0/6-7,12:00-24:00;{$Macr Property behavior: - *required* if type is set to "Zabbix agent" (0), "Simple check" (3), "Zabbix internal" (5), "External check" (10), "Database monitor" (11), "IPMI agent" (12), "SSH agent" (13), "TELNET agent" (14), "Calculated" (15), "JMX agent" (16), "HTTP agent" (19), "SNMP agent" (20), "Script" (21), "Browser" (22), "Telemetry query" (24), or if type is set to "Zabbix agent (active)" (7) and key_ does not contain "mqtt.get"
hostid	ID	ID of the host that the item prototype belongs to. Property behavior: - *constant* - *required* for create operations
interfaceid	ID	ID of the item prototype's host interface. Property behavior: - *required* if item prototype belongs to host and type is set to "Zabbix agent", "IPMI agent", "JMX agent", "SNMP trap", or "SNMP agent" - *supported* if item prototype belongs to host and type is set to "Simple check", "External check", "SSH agent", "TELNET agent", or "HTTP agent"
key_	string	Item prototype key. Property behavior: - *required* for create operations - *read-only* for inherited objects
name	string	Name of the item prototype. Supports user macros. Property behavior: - *required* for create operations - *read-only* for inherited objects

Property	Type	Description
type	integer	Type of the item prototype. Possible values: 0 - Zabbix agent; 2 - Zabbix trapper; 3 - Simple check; 5 - Zabbix internal; 7 - Zabbix agent (active); 10 - External check; 11 - Database monitor; 12 - IPMI agent; 13 - SSH agent; 14 - TELNET agent; 15 - Calculated; 16 - JMX agent; 17 - SNMP trap; 18 - Dependent item; 19 - HTTP agent; 20 - SNMP agent; 21 - Script; 22 - Browser; 24 - Telemetry query. Property behavior: - *required* for create operations - *read-only* for inherited objects
url	string	URL string. Supports LLD macros, user macros, {HOST.IP}, {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.NAME}, {HOST.PORT}, {ITEM.ID}, {ITEM.KEY}. Property behavior: - *required* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
value_type	integer	Type of information of the item prototype. Possible values: 0 - numeric float; 1 - character; 2 - log; 3 - numeric unsigned; 4 - text; 5 - binary; 6 - JSON. Property behavior: - *required* for create operations - *read-only* for inherited objects
allow_traps	integer	Allow the value to be populated similarly to the trapper item. 0 - *(default)* Do not accept incoming data; 1 - Accept incoming data. Property behavior: - *supported* if `type` is set to ”HTTP agent”

Property	Type	Description
authtype	integer	Authentication method. Possible values if `type` is set to ”SSH agent”: 0 - *(default)* password; 1 - public key. Possible values if `type` is set to ”HTTP agent”: 0 - *(default)* none; 1 - basic; 2 - NTLM; 3 - Kerberos; 4 - Digest. Property behavior: - *supported* if `type` is set to ”SSH agent” or ”HTTP agent” - *read-only* for inherited objects (if `type` is set to ”HTTP agent”)
description	string	Description of the item prototype.
flags	integer	Origin of the item prototype. Possible values: 2 - an item prototype; 6 - a discovered item prototype Property behavior: - *read-only*
follow_redirects	integer	Follow response redirects while polling data. Possible values: 0 - Do not follow redirects; 1 - *(default)* Follow redirects. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
headers	array	Array of headers that will be sent when performing an HTTP request. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
history	string	A time unit of how long the history data should be stored. Also accepts user macro and LLD macro. Default: 31d.
http_proxy	string	HTTP(S) proxy connection string. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
ipmi_sensor	string	IPMI sensor. Property behavior: - *required* if `type` is set to ”IPMI agent” and `key_` is not set to ”ipmi.get” - *supported* if `type` is set to ”IPMI agent” - *read-only* for inherited objects
jmx_endpoint	string	JMX agent custom connection string. Default: service:jmx:rmi:///jndi/rmi://{HOST.CONN}:{HOST.PORT}/jmxrmi Property behavior: - *supported* if `type` is set to ”JMX agent”

Property	Type	Description
logtimefmt	string	Format of the time in log entries. Property behavior: - *supported* if `value_type` is set to ”log” - *read-only* for inherited objects
master_itemid	ID	ID of the master item. Property behavior: - *required* if `type` is set to ”Dependent item” - *read-only* for inherited objects
output_format	integer	Should the response be converted to JSON. Possible values: 0 - *(default)* Store raw; 1 - Convert to JSON. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
params	string	Additional parameters depending on the type of the item prototype: - executed script for SSH agent and TELNET agent item prototypes; - SQL query for database monitor item prototypes; - formula for calculated item prototypes; - the script for script and browser item prototypes. Property behavior: - *required* if `type` is set to ”Database monitor”, ”SSH agent”, ”TELNET agent”, ”Calculated”, ”Script”, or ”Browser” - *read-only* for inherited objects (if `type` is set to ”Script” or ”Browser”)
parameters	object/array	Additional parameters if `type` is set to ”Script” or ”Browser”. Array of objects with `name` and `value` properties, where `name` must be unique. Property behavior: - *supported* if `type` is set to ”Script” or ”Browser” - *read-only* for inherited objects
password	string	Password for authentication. Property behavior: - *required* if `type` is set to ”JMX agent” and `username` is set - *supported* if `type` is set to ”Simple check”, ”SSH agent”, ”TELNET agent”, ”Database monitor”, or ”HTTP agent” - *read-only* for inherited objects (if `type` is set to ”HTTP agent”)
post_type	integer	Type of post data body stored in `posts` property. Possible values: 0 - *(default)* Raw data. 2 - JSON data. 3 - XML data. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
posts	string	HTTP(S) request body data. Property behavior: - *required* if `type` is set to ”HTTP agent” and `post_type` is set to ”JSON data” or ”XML data” - *supported* if `type` is set to ”HTTP agent” and `post_type` is set to ”Raw data” - *read-only* for inherited objects

Property	Type	Description
privatekey	string	Name of the private key file. Property behavior: - *required* if `type` is set to ”SSH agent” and `authtype` is set to ”public key”
publickey	string	Name of the public key file. Property behavior: - *required* if `type` is set to ”SSH agent” and `authtype` is set to ”public key”
query_fields	array	Array of query fields that will be sent when performing an HTTP request. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
request_method	integer	Type of request method. Possible values: 0 - *(default)* GET; 1 - POST; 2 - PUT; 3 - HEAD. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
retrieve_mode	integer	What part of response should be stored. Possible values if `request_method` is set to ”GET”, ”POST”, or ”PUT”: 0 - *(default)* Body; 1 - Headers; 2 - Both body and headers will be stored. Possible values if `request_method` is set to ”HEAD”: 1 - Headers. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
snmp_oid	string	SNMP OID. Property behavior: - *required* if `type` is set to ”SNMP agent” - *read-only* for inherited objects
ssl_cert_file	string	Public SSL Key file path. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
ssl_key_file	string	Private SSL Key file path. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
ssl_key_password	string	Password for SSL Key file. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects

Property	Type	Description
status	integer	Status of the item prototype. Possible values: 0 - *(default)* enabled item prototype; 1 - disabled item prototype; 3 - unsupported item prototype.
status_codes	string	Ranges of required HTTP status codes, separated by commas. Also supports user macros or LLD macros as part of comma separated list. Example: 200,200-{$M},{$M},200-400 Property behavior: - *supported* if `type` is set to "HTTP agent" - *read-only* for inherited objects
templateid	ID	ID of the parent template item prototype. Property behavior: - *read-only*
timeout	string	Item data polling request timeout. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m). Also accepts user macros and LLD macros. Possible values range: 1-600s. Default: "" - use proxy/global settings. Property behavior: - *supported* if `type` is set to "Zabbix agent" (0), "Simple check" (3) and `key_` does not start with "`vmware.`" and "`icmpping`", "Zabbix agent (active)" (7), "External check" (10), "Database monitor" (11), "SSH agent" (13), "TELNET agent" (14), "HTTP agent" (19), "SNMP agent" (20) and `snmp_oid` starts with "`walk[`" or "`get[`", "Script" (21), "Browser" (22) - *read-only* for inherited objects
trapper_hosts	string	Allowed hosts. Property behavior: - *supported* if `type` is set to "Zabbix trapper", or if `type` is set to "HTTP agent" and `allow_traps` is set to "Allow incoming data". If the property is omitted, it is automatically set to the `{$TRAPPER.ALLOWED_HOSTS}` user macro (default value). By default, `{$TRAPPER.ALLOWED_HOSTS}` is `127.0.0.1,::1`. If the property is set to an empty string, incoming connections will be denied from all hosts.
trends	string	A time unit of how long the trends data should be stored. Also accepts user macro and LLD macro. Default: 365d. Property behavior: - *supported* if `value_type` is set to "numeric float" or "numeric unsigned"
units	string	Value units. Property behavior: - *supported* if `value_type` is set to "numeric float" or "numeric unsigned" - *read-only* for inherited objects

Property	Type	Description
username	string	Username for authentication. Property behavior: - *required* if `type` is set to ”SSH agent” or ”TELNET agent”, or if `type` is set to ”JMX agent” and `password` is set - *supported* if `type` is set to ”Simple check”, ”Database monitor”, or ”HTTP agent” - *read-only* for inherited objects (if `type` is set to ”HTTP agent”)
uuid	string	Universal unique identifier, used for linking imported item prototypes to already existing ones. Auto-generated, if not given. Property behavior: - *supported* if the item prototype belongs to a template
valuemapid	ID	ID of the associated value map. Property behavior: - *supported* if `value_type` is set to ”numeric float”, ”character”, or ”numeric unsigned” - *read-only* for inherited objects
verify_host	integer	Whether to validate that the host name for the connection matches the one in the host’s certificate. Possible values: 0 - *(default)* Do not validate; 1 - Validate. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
verify_peer	integer	Whether to validate that the host’s certificate is authentic. Possible values: 0 - *(default)* Do not validate; 1 - Validate. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
discover	integer	Item prototype discovery status. Possible values: 0 - *(default)* new items will be discovered; 1 - new items will not be discovered and existing items will be marked as lost.
query	object	Telemetry query configuration. Defines what the item queries from the telemetry data source, how it filters the data, and how it aggregates the result. Property behavior: - *required* if `type` is set to ”Telemetry query” - *read-only* for inherited objects or discovered objects

Property	Type	Description
time_shift	string	How far back (before the current time) ends the time range that the item queries. Accepts a number of seconds (e.g., 1, 60, 300) or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Also accepts user macros and LLD macros. Possible values range: 0-86400s. Default: 15s. Property behavior: - *supported* if `type` is set to "Telemetry query"
lookback_limit	string	How far back (before the current time) the item queries for data when there is no previous query to continue from (e.g., on the item's first run), or when the previous query is older than `delay` (e.g., because Zabbix server was down). Accepts a number of seconds (e.g., 1, 60, 300) or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Also accepts user macros and LLD macros. Possible values range: 1-259200s. Default: 10m. Property behavior: - *supported* if `type` is set to "Telemetry query"
granularity	string	The length of a piece that the query's full time range is split into. Accepts a number of seconds (e.g., 1, 60, 300) or a value with a time suffix (e.g., 30s, 1m, 2h, 1d). Also accepts user macros and LLD macros. Possible values range: 1-86400s. Default: 15s. Property behavior: - *supported* if `type` is set to "Telemetry query"

HTTP header

The header object has the following properties:

Property	Type	Description
name	string	HTTP header name. Property behavior: - *required*
value	string	Header value. Property behavior: - *required*

HTTP query field

The query field object defines a name and value that is used to specify a URL parameter. It has the following properties:

Property	Type	Description
name	string	Name of the parameter. Property behavior: - *required*
value	string	Parameter value. Property behavior: - *required*

Telemetry query

The telemetry query object has the following properties.

Property	Type	Description
signal_type	integer	Signal type that the item queries. Possible values: 0 - Traces; 1 - Metrics; 2 - Logs. Property behavior: - *required*
metric_point_type	integer	Metric point type that the item queries. Possible values: 0 - Sum; 1 - Gauge; 2 - Histogram; 3 - Exponential histogram. Property behavior: - *required* if `signal_type` is set to ”Metrics”
columns	array	Array of telemetry data columns that the item queries. Can be empty. Property behavior: - *required*
aggregated_columns	array	Array of telemetry data aggregated columns that the item queries, and how it calculates a value from column records. Property behavior: - *required*
filter	object	Telemetry query filter for column records in the query. Property behavior: - *required*

Telemetry data columns

The telemetry data columns object has the following properties.

Property	Type	Description
column	string	Name of the telemetry data column that the item queries. See supported telemetry data columns. For multiple single-value columns, each `column` value must be unique. For multiple columns that hold key-value pairs, each `column` and `attribute_key` pair must be unique.

Property	Type	Description
attribute_key	string	Name of the key that the item queries, within the column that holds key-value pairs. For multiple columns that hold key-value pairs, each `column` and `attribute_key` pair must be unique. Property behavior: - *required* if `column` is set to a column that holds key-value pairs

Telemetry data aggregated columns

The telemetry data aggregated columns object has the following properties.

Property	Type	Description
column	string	Name of the telemetry data column that the item queries. See supported telemetry data columns. For multiple columns, each `column` value must be unique. This property is not supported if `function` is set to ”Count”. Property behavior: - *required* if `function` is set to ”Min”, ”Max”, ”Average”, ”Sum”, or ”Percentile”
function	integer	Function that the item applies to the column records. Possible values: 1 - Min; 2 - Max; 3 - Average; 4 - Count; 5 - Sum; 8 - Percentile. Property behavior: - *required*
parameters	array	Percentile that the item calculates from the column records. Possible value: a number from 0 to 100, given as a string. Property behavior: - *required* if `function` is set to ”Percentile”
alias	string	Unique name for the calculated value, used to identify it in the item value. Each `alias` value must be unique. Property behavior: - *required*

Telemetry data filter

The telemetry data filter object has the following properties.

Property	Type	Description
conditions	array	Telemetry data filter conditions a record in a telemetry data column must meet. Property behavior: - *required*
evaltype	integer	Filter condition evaluation method. Possible values: 0 - And/Or; 1 - And; 2 - Or; 3 - Custom expression. Property behavior: - *required*
eval_formula	string	Generated expression that will be used for evaluating filter conditions. The expression contains IDs that reference specific filter conditions by its `formulaid`. The value of `eval_formula` is equal to the value of `formula` for filters with a custom expression. Property behavior: - *read-only*
formula	string	User-defined expression to be used for evaluating conditions of filters with a custom expression. The expression must contain IDs that reference specific filter conditions by its `formulaid`. The IDs used in the expression must exactly match the ones defined in the filter conditions: no condition can remain unused or omitted. Property behavior: - *required* if `evaltype` is set to ”Custom expression”

Telemetry data filter conditions

The telemetry data filter conditions object has the following properties.

Property	Type	Description
column	string	Name of the telemetry data column that the filter condition applies to. See supported telemetry data columns. Property behavior: - *required*
attribute_key	string	Name of the key that the condition applies to, within the column that holds key-value pairs. Property behavior: - *required* if `column` is set to a column that holds key-value pairs

Property	Type	Description
operator	integer	Filter condition operator. Possible values for single-value columns: 0 - Equals; 1 - Does not equal; 2 - Contains; 3 - Does not contain. Possible values for columns that hold key-value pairs: 0 - Equals; 1 - Does not equal; 12 - Exists. Property behavior: - *required*
value	string	Value to compare with. Property behavior: - *required* if `operator` is set to ”Equals”, ”Does not equal”, ”Contains” or ”Does not contain”
formulaid	string	Arbitrary unique ID that is used to reference the condition from a custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward. Property behavior: - *required* if `evaltype` of Telemetry query filter object is set to ”Custom expression”

Item prototype tag

The item prototype tag object has the following properties.

Property	Type	Description
tag	string	Item prototype tag name. Property behavior: - *required*
value	string	Item prototype tag value. Property behavior: - *read-only*.
object	integer	Type of object the tag was inherited from. Possible values: 0 - template; 1 - host. Property behavior: - *read-only*.
objectid	ID	ID of the object the tag was inherited from. Property behavior: - *read-only*.

Item prototype preprocessing

The item prototype preprocessing object has the following properties.

Property	Type	Description
type	integer	The preprocessing option type. Possible values: 1 - Custom multiplier; 2 - Right trim; 3 - Left trim; 4 - Trim; 5 - Regular expression; 6 - Boolean to decimal; 7 - Octal to decimal; 8 - Hexadecimal to decimal; 9 - Simple change; 10 - Change per second; 11 - XML XPath; 12 - JSONPath; 13 - In range; 14 - Matches regular expression; 15 - Does not match regular expression; 16 - Check for error in JSON; 17 - Check for error in XML; 18 - Check for error using regular expression; 19 - Discard unchanged; 20 - Discard unchanged with heartbeat; 21 - JavaScript; 22 - Prometheus pattern; 23 - Prometheus to JSON; 24 - CSV to JSON; 25 - Replace; 26 - Check unsupported; 27 - XML to JSON; 28 - SNMP walk value; 29 - SNMP walk to JSON; 30 - SNMP get value. Property behavior: - *required*
params	string	Additional parameters used by preprocessing option. Multiple parameters are separated by the newline (\n) character. If `type` is set to ”Check unsupported”, the parameters follow a `<scope>[\n<pattern>]` syntax, where *pattern* is a regular expression, and *scope* is one of: -1 - match any error; 0 - check if error message matches *pattern*; 1 - check if error message does not match *pattern*. Property behavior: - *required* if `type` is set to ”Custom multiplier” (1), ”Right trim” (2), ”Left trim” (3), ”Trim” (4), ”Regular expression” (5), ”XML XPath” (11), ”JSONPath” (12), ”In range” (13), ”Matches regular expression” (14), ”Does not match regular expression” (15), ”Check for error in JSON” (16), ”Check for error in XML” (17), ”Check for error using regular expression” (18), ”Discard unchanged with heartbeat” (20), ”JavaScript” (21), ”Prometheus pattern” (22), ”Prometheus to JSON” (23), ”CSV to JSON” (24), ”Replace” (25), ”Check unsupported” (26), ”SNMP walk value” (28), ”SNMP walk to JSON” (29), or ”SNMP get value” (30)

Property	Type	Description
error_handler	integer	Action type used in case of preprocessing step failure. Possible values: 0 - Error message is set by Zabbix server; 1 - Discard value; 2 - Set custom value; 3 - Set custom error message. Possible values if `type` is set to ”Check unsupported”: 1 - Discard value; 2 - Set custom value; 3 - Set custom error message. Property behavior: - *required* if `type` is set to ”Custom multiplier” (1), ”Regular expression” (5), ”Boolean to decimal” (6), ”Octal to decimal” (7), ”Hexadecimal to decimal” (8), ”Simple change” (9), ”Change per second” (10), ”XML XPath” (11), ”JSONPath” (12), ”In range” (13), ”Matches regular expression” (14), ”Does not match regular expression” (15), ”Check for error in JSON” (16), ”Check for error in XML” (17), ”Check for error using regular expression” (18), ”Prometheus pattern” (22), ”Prometheus to JSON” (23), ”CSV to JSON” (24), ”Check unsupported” (26), ”XML to JSON” (27), ”SNMP walk value” (28), ”SNMP walk to JSON” (29), or ”SNMP get value” (30)
error_handler_params	string	Error handler parameters. Property behavior: - *required* if `error_handler` is set to ”Set custom value” or ”Set custom error message”

The following parameters and error handlers are supported for each preprocessing type.

Preprocessing type	Name	Parameter 1	Parameter 2	Parameter 3	Supported error handlers
1	Custom multiplier	number^1, 6^			0, 1, 2, 3
2	Right trim	list of characters[2]			
3	Left trim	list of characters[2]			
4	Trim	list of characters[2]			
5	Regular expression	pattern[3]	output[2]		0, 1, 2, 3
6	Boolean to decimal				0, 1, 2, 3
7	Octal to decimal				0, 1, 2, 3
8	Hexadecimal to decimal				0, 1, 2, 3

Preprocessing type	Name	Parameter 1	Parameter 2	Parameter 3	Supported error handlers
9	Simple change				0, 1, 2, 3
10	Change per second				0, 1, 2, 3
11	XML XPath	path[4]			0, 1, 2, 3
12	JSONPath	path[4]			0, 1, 2, 3
13	In range	min^1, 6^	max^1, 6^		0, 1, 2, 3
14	Matches regular expression	pattern[3]			0, 1, 2, 3
15	Does not match regular expression	pattern[3]			0, 1, 2, 3
16	Check for error in JSON	path[4]			0, 1, 2, 3
17	Check for error in XML	path[4]			0, 1, 2, 3
18	Check for error using regular expression	pattern[3]	output[2]		0, 1, 2, 3
19	Discard unchanged				
20	Discard unchanged with heartbeat	seconds^5, 6^			
21	JavaScript	script[2]			
22	Prometheus pattern	pattern^6, 7^	`value`, `label`, `function`	output^8, 9^	0, 1, 2, 3
23	Prometheus to JSON	pattern^6, 7^			0, 1, 2, 3

Preprocessing type	Name	Parameter 1	Parameter 2	Parameter 3	Supported error handlers
24	CSV to JSON	character[2]	character[2]	0,1	0, 1, 2, 3
25	Replace	search string[2]	replacement[2]		
26	Check unsupported	scope[1]	pattern^3, 6^		1, 2, 3
27	XML to JSON				0, 1, 2, 3
28	SNMP walk value	OID[2]	Format: 0 - Unchanged 1 - UTF-8 from Hex-STRING 2 - MAC from Hex-STRING 3 - Integer from BITS		0, 1, 2, 3
29	SNMP walk to JSON[10]	Field name[2]	OID prefix[2]	Format: 0 - Unchanged 1 - UTF-8 from Hex-STRING 2 - MAC from Hex-STRING 3 - Integer from BITS	0, 1, 2, 3
30	SNMP get value	Format: 1 - UTF-8 from Hex-STRING 2 - MAC from Hex-STRING 3 - Integer from BITS			0, 1, 2, 3

[1] integer or floating-point number
 [2] string
 [3] regular expression
 [4] JSONPath or XML XPath
 [5] positive integer (with support of time suffixes, e.g., 30s, 1m, 2h, 1d)
 [6] user macro, LLD macro
 [7] Prometheus pattern following the syntax: `<metric name>{<label name>="<label value>", ...} == <value>`. Each Prometheus pattern component (metric, label name, label value and metric value) can be user macro or LLD macro.
 [8] Prometheus output following the syntax: `<label name>` (can be a user macro or an LLD macro) if `label` is selected as the second parameter.
 [9] One of the aggregation functions: `sum`, `min`, `max`, `avg`, `count` if `function` is selected as the second parameter.
 [10] Supports multiple "Field name,OID prefix,Format records" records delimited by a new line character.

itemprototype.create

Description

`object itemprototype.create(object/array itemPrototypes)`

This method allows you to create new item prototypes.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Item prototype to create.

Additionally to the standard item prototype properties, the method accepts the following parameters.

Parameter	Type	Description
ruleid	ID	ID of the LLD rule that the item belongs to. Parameter behavior: - *required*
preprocessing	array	Item prototype preprocessing options.
tags	array	Item prototype tags.

Return values

`(object)` Returns an object containing the IDs of the created item prototypes under the `itemids` property. The order of the returned IDs matches the order of the passed item prototypes.

Examples

Creating an item prototype

Create an item prototype to monitor free disk space on a discovered file system. Discovered items should be numeric Zabbix agent items updated every 30 seconds.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.create",
    "params": {
        "name": "Free disk space on {#FSNAME}",
        "key_": "vfs.fs.size[{#FSNAME},free]",
        "hostid": "10197",
        "ruleid": "27665",
        "type": 0,
        "value_type": 3,
        "interfaceid": "112",
        "delay": "30s"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "27666"
        ]
    },
    "id": 1
}
```

Creating an item prototype with preprocessing

Create an item using change per second and a custom multiplier as a second step.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.create",
    "params": {
        "name": "Incoming network traffic on {#IFNAME}",
        "key_": "net.if.in[{#IFNAME}]",
        "hostid": "10001",
        "ruleid": "27665",
        "type": 0,
        "value_type": 3,
        "delay": "60s",
```

```
        "units": "bps",
        "interfaceid": "1155",
        "preprocessing": [
            {
                "type": 10,
                "params": "",
                "error_handler": 0,
                "error_handler_params": ""
            },
            {
                "type": 1,
                "params": "8",
                "error_handler": 2,
                "error_handler_params": "10"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "44211"
        ]
    },
    "id": 1
}
```

Creating dependent item prototype

Create Dependent item prototype for Master item prototype with ID 44211. Only dependencies on same host (template/discovery rule) are allowed, therefore Master and Dependent item should have same hostid and ruleid.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.create",
    "params": {
      "hostid": "10001",
      "ruleid": "27665",
      "name": "Dependent test item prototype",
      "key_": "dependent.prototype",
      "type": 18,
      "master_itemid": "44211",
      "value_type": 3
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "44212"
        ]
    },
    "id": 1
}
```

Create HTTP agent item prototype

Create item prototype with URL using user macro, query fields and custom headers.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.create",
    "params": {
        "type": "19",
        "hostid": "10254",
        "ruleid": "28256",
        "interfaceid": "2",
        "name": "api item prototype example",
        "key_": "api_http_item",
        "value_type": 3,
        "url": "{$URL_PROTOTYPE}",
        "query_fields": [
            {
                "name": "min",
                "value": "10"
            },
            {
                "name": "max",
                "value" "100"
            }
        ],
        "headers": [
            {
                "name": "X-Source",
                "value": "api"
            }
        ],
        "delay": "35"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "28305"
        ]
    },
    "id": 1
}
```

Create script item prototype

Create a simple data collection using a script item prototype.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.create",
    "params": {
        "name": "Script example",
        "key_": "custom.script.itemprototype",
        "hostid": "12345",
        "type": 21,
        "value_type": 4,
        "params": "var request = new HttpRequest();\nreturn request.post(\"https://postman-echo.com/post\"
```

```
        "parameters": [
            {
                "name": "host",
                "value": "{HOST.CONN}"
            }
        ],
        "timeout": "6s",
        "delay": "30s"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "23865"
        ]
    },
    "id": 1
}
```

Source

CItemPrototype::create() in *ui/include/classes/api/services/CItemPrototype.php*.

itemprototype.delete

Description

`object itemprototype.delete(array itemPrototypeIds)`

This method allows you to delete item prototypes.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the item prototypes to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted item prototypes under the `prototypeids` property.

Examples

Deleting multiple item prototypes

Delete two item prototypes.
 Dependent item prototypes are removed automatically if master item or item prototype is deleted.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.delete",
    "params": [
        "27352",
        "27356"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "prototypeids": [
            "27352",
            "27356"
        ]
    },
    "id": 1
}
```

Source

CItemPrototype::delete() in *ui/include/classes/api/services/CItemPrototype.php*.

itemprototype.get

Description

`integer/array itemprototype.get(object parameters)`

The method allows you to retrieve item prototypes according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
discoveryids	ID/array	Return only item prototypes that belong to the given LLD rules.
graphids	ID/array	Return only item prototypes that are used in the given graph prototypes.
hostids	ID/array	Return only item prototypes that belong to the given hosts.
inherited	boolean	If set to `true` return only item prototypes inherited from a template.
itemids	ID/array	Return only item prototypes with the given IDs.
monitored	boolean	If set to `true` return only enabled item prototypes that belong to monitored hosts.
templated	boolean	If set to `true` return only item prototypes that belong to templates.
templateids	ID/array	Return only item prototypes that belong to the given templates.
triggerids	ID/array	Return only item prototypes that are used in the given trigger prototypes.

Parameter	Type	Description
selectDiscoveryData	query	Return a `discoveryData` property with the item prototype discovery object data. The item prototype discovery object links a discovered item prototype to an item prototype from which it was discovered. It has the following properties: `parent_itemid` - `(string)` ID of the item prototype from which the item prototype has been created; `key_` - `(string)` key of the item prototype; `status` - `(int)` item prototype discovery status: 0 - *(default)* item prototype is discovered, 1 - item prototype is not discovered anymore; `ts_delete` - `(timestamp)` time when an item prototype that is no longer discovered will be deleted; `ts_disable` - `(timestamp)` time when an item prototype that is no longer discovered will be disabled; `disable_source` - `(int)` indicator of whether item prototype was disabled by an LLD rule or manually: 0 - *(default)* disabled automatically, 1 - disabled by an LLD rule.
selectDiscoveryRule	query	Return a `discoveryRule` property with the low-level discovery rule that the item prototype belongs to.
selectDiscoveryRulePrototype	query	Return a `discoveryRulePrototype` property with the parent LLD rule prototype that the item prototype belongs to.
selectGraphs	query	Return a `graphs` property with graph prototypes that the item prototype is used in. Supports `count`.
selectHosts	query	Return a `hosts` property with an array of hosts that the item prototype belongs to.
selectInheritedTags	query	Return an `inheritedTags` property with tags that are on template/host/linked templates.
selectTags	query	Return the item prototype tags in `tags` property.
selectTriggers	query	Return a `triggers` property with trigger prototypes that the item prototype is used in. Supports `count`.
selectPreprocessing	query	Return a `preprocessing` property with item prototype preprocessing options.
selectValueMap	query	Return a `valuemap` property with item prototype value map.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type. Supports additional properties: `host` - technical name of the host that the item prototype belongs to.
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectGraphs` - results will be sorted by `name`; `selectTriggers` - results will be sorted by `description`.
sortfield	string/array	Sort the result by the given properties. Possible values: `itemid`, `name`, `key_`, `delay`, `type`, `status`, `history`, `trends`, `discovered`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
limit	integer	

Parameter	Type	Description
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving item prototypes from an LLD rule

Retrieve all item prototypes for specific LLD rule ID.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.get",
    "params": {
        "output": "extend",
        "discoveryids": "27426"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "itemid": "23077",
            "type": "0",
            "snmp_oid": "",
            "hostid": "10079",
            "name": "Incoming network traffic on en0",
            "key_": "net.if.in[en0]",
            "delay": "1m",
            "history": "1w",
            "trends": "365d",
            "status": "0",
            "value_type": "3",
            "trapper_hosts": "",
            "units": "bps",
            "formula": "",
            "logtimefmt": "",
            "templateid": "0",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "2",
            "interfaceid": "0",
```

```
        "description": "",
        "inventory_link": "0",
        "lifetime": "7d",
        "evaltype": "0",
        "jmx_endpoint": "",
        "master_itemid": "0",
        "timeout": "",
        "url": "",
        "query_fields": [],
        "posts": "",
        "status_codes": "200",
        "follow_redirects": "1",
        "post_type": "0",
        "http_proxy": "",
        "headers": [],
        "retrieve_mode": "0",
        "request_method": "0",
        "output_format": "0",
        "ssl_cert_file": "",
        "ssl_key_file": "",
        "ssl_key_password": "",
        "verify_peer": "0",
        "verify_host": "0",
        "allow_traps": "0",
        "discover": "0",
        "uuid": "",
        "lifetime_type": "0",
        "enabled_lifetime_type": "2",
        "enabled_lifetime": "0",
        "parameters": []
    },
    {
        "itemid": "10010",
        "type": "0",
        "snmp_oid": "",
        "hostid": "10001",
        "name": "Processor load (1 min average per core)",
        "key_": "system.cpu.load[percpu,avg1]",
        "delay": "1m",
        "history": "1w",
        "trends": "365d",
        "status": "0",
        "value_type": "0",
        "trapper_hosts": "",
        "units": "",
        "formula": "",
        "logtimefmt": "",
        "templateid": "0",
        "valuemapid": "0",
        "params": "",
        "ipmi_sensor": "",
        "authtype": "0",
        "username": "",
        "password": "",
        "publickey": "",
        "privatekey": "",
        "flags": "2",
        "interfaceid": "0",
        "description": "The processor load is calculated as system CPU load divided by number of CPU c
        "inventory_link": "0",
        "lifetime": "7d",
        "evaltype": "0",
```

```
            "jmx_endpoint": "",
            "master_itemid": "0",
            "timeout": "",
            "url": "",
            "query_fields": [],
            "posts": "",
            "status_codes": "200",
            "follow_redirects": "1",
            "post_type": "0",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "0",
            "request_method": "0",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "discover": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
            "parameters": []
        }
    ],
    "id": 1
}
```


Finding dependent item

Find one Dependent item for specific item ID.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.get",
    "params": {
        "output": "extend",
        "filter": {
            "type": 18,
            "master_itemid": "25545"
        },
        "limit": "1"
    },
    "id": 1
}
```

Response:


```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "itemid": "25547",
            "type": "18",
            "snmp_oid": "",
            "hostid": "10116",
            "name": "Seconds",
            "key_": "apache.status.uptime.seconds",
            "delay": "0",
            "history": "90d",
```

```
            "trends": "365d",
            "status": "0",
            "value_type": "3",
            "trapper_hosts": "",
            "units": "",
            "formula": "",
            "logtimefmt": "",
            "templateid": "0",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "0",
            "interfaceid": "0",
            "description": "",
            "inventory_link": "0",
            "lifetime": "7d",
            "evaltype": "0",
            "jmx_endpoint": "",
            "master_itemid": "25545",
            "timeout": "",
            "url": "",
            "query_fields": [],
            "posts": "",
            "status_codes": "200",
            "follow_redirects": "1",
            "post_type": "0",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "0",
            "request_method": "0",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "discover": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
            "parameters": []
        }
    ],
    "id": 1
}
```


Find HTTP agent item prototype

Find HTTP agent item prototype with request method HEAD for specific host ID.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.get",
    "params": {
        "hostids": "10254",
```

```
        "filter": {
            "type": 19,
            "request_method": 3
        }
    },
    "id": 1
}
```

Response:


```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "itemid": "28257",
            "type": "19",
            "snmp_oid": "",
            "hostid": "10254",
            "name": "discovered",
            "key_": "item[{#INAME}]",
            "delay": "{#IUPDATE}",
            "history": "90d",
            "trends": "30d",
            "status": "0",
            "value_type": "3",
            "trapper_hosts": "",
            "units": "",
            "formula": "",
            "logtimefmt": "",
            "templateid": "28255",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "2",
            "interfaceid": "2",
            "description": "",
            "inventory_link": "0",
            "lifetime": "7d",
            "evaltype": "0",
            "jmx_endpoint": "",
            "master_itemid": "0",
            "timeout": "",
            "url": "{#IURL}",
            "query_fields": [],
            "posts": "",
            "status_codes": "",
            "follow_redirects": "0",
            "post_type": "0",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "0",
            "request_method": "3",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
```

```
            "allow_traps": "0",
            "discover": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
            "parameters": []
        }
    ],
    "id": 1
}
```


See also

- Host
- Graph prototype
- Trigger prototype

Source

CItemPrototype::get() in *ui/include/classes/api/services/CItemPrototype.php*.

itemprototype.update

Description

`object itemprototype.update(object/array itemPrototypes)`

This method allows you to update existing item prototypes.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Item prototype properties to be updated.

The `itemid` property must be defined for each item prototype, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard item prototype properties, the method accepts the following parameters.

Parameter	Type	Description
preprocessing	array	Item prototype preprocessing options to replace the current preprocessing options. Parameter behavior: - *read-only* for inherited objects
tags	array	Item prototype tags.

Return values

`(object)` Returns an object containing the IDs of the updated item prototypes under the `itemids` property.

Examples

Changing the interface of an item prototype

Change the host interface that will be used by discovered items.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.update",
    "params": {
```

```
        "itemid": "27428",
        "interfaceid": "132"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "27428"
        ]
    },
    "id": 1
}
```

Update dependent item prototype

Update Dependent item prototype with new Master item prototype ID. Only dependencies on same host (template/discovery rule) are allowed, therefore Master and Dependent item should have same hostid and ruleid.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.update",
    "params": {
        "master_itemid": "25570",
        "itemid": "189030"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "189030"
        ]
    },
    "id": 1
}
```

Update HTTP agent item prototype

Change query fields and remove all custom headers.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.update",
    "params": {
        "itemid":"28305",
        "query_fields": [
            {
                "name": "random",
                "value": "qwertyuiopasdfghjklzxcvbnm"
            }
        ],
        "headers": []
    }
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "28305"
        ]
    },
    "id": 1
}
```

Updating item preprocessing options

Update an item prototype with item preprocessing rule “Custom multiplier”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.update",
    "params": {
        "itemid": "44211",
        "preprocessing": [
            {
                "type": 1,
                "params": "4",
                "error_handler": 2,
                "error_handler_params": "5"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "44211"
        ]
    },
    "id": 1
}
```

Updating a script item prototype

Update a script item prototype with a different script and remove unnecessary parameters that were used by previous script.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "itemprototype.update",
    "params": {
        "itemid": "23865",
        "parameters": [],
        "script": "Zabbix.log(3, 'Log test');\nreturn 1;"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
```

```
        "itemids": [
            "23865"
        ]
    },
    "id": 1
}
```

Source

CItemPrototype::update() in *ui/include/classes/api/services/CItemPrototype.php*.

LLD rule

This class is designed to work with low-level discovery rules.

Object references:

- LLD rule
  - HTTP header
  - HTTP query field
- LLD rule filter
  - LLD rule filter condition
- LLD macro path
- LLD rule preprocessing
- LLD rule overrides
  - LLD rule override filter
    - LLD rule override filter condition
  - LLD rule override operation
    - LLD rule override operation status
    - LLD rule override operation discover
    - LLD rule override operation period
    - LLD rule override operation history
    - LLD rule override operation trends
    - LLD rule override operation severity
    - LLD rule override operation tag
    - LLD rule override operation template
    - LLD rule override operation inventory

Available methods:

- discoveryrule.create - create new LLD rules
- discoveryrule.delete - delete LLD rules
- discoveryrule.get - retrieve LLD rules
- discoveryrule.update - update LLD rules

LLD rule object

The following objects are directly related to the `discoveryrule` API.

LLD rule

The low-level discovery rule object has the following properties.

Property	Type	Description
itemid	ID	ID of the LLD rule. Property behavior: - *read-only* - *required* for update operations

Property	Type	Description
delay	string	Update interval of the LLD rule. Accepts a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., 30s, 1m, 2h, 1d), or, optionally, one or more custom intervals, all separated by semicolons. Custom intervals can be a mix of flexible and scheduling intervals. Accepts user macros. If used, the value must be a single macro. Multiple macros or macros mixed with text are not supported. Flexible intervals may be written as two macros separated by a forward slash (e.g., {$FLEX_INTERVAL}/{$FLEX_PERIOD}). Example: 1h;wd1-5h9-18;{$Macro1}/1-7,00:00-24:00;0/6-7,12:00-24:00;{$Macr Property behavior: - *required* if type is set to ”Zabbix agent” (0), ”Simple check” (3), ”Zabbix internal” (5), ”External check” (10), ”Database monitor” (11), ”IPMI agent” (12), ”SSH agent” (13), ”TELNET agent” (14), ”JMX agent” (16), ”HTTP agent” (19), ”SNMP agent” (20), ”Script” (21), ”Browser” (22), or if type is set to ”Zabbix agent (active)” (7) and key_ does not contain ”mqtt.get”
hostid	ID	ID of the host that the LLD rule belongs to. Property behavior: - *constant* - *required* for create operations
flags	integer	Origin of the discovery rule. Possible values: 1 - a low-level discovery rule; 5 - a low-level discovery rule converted from prototype. Property behavior: - *read-only*
interfaceid	ID	ID of the LLD rule’s host interface. Property behavior: - *required* if LLD rule belongs to host and type is set to ”Zabbix agent”, ”IPMI agent”, ”JMX agent”, or ”SNMP agent” - *supported* if LLD rule belongs to host and type is set to ”Simple check”, ”External check”, ”SSH agent”, ”TELNET agent”, or ”HTTP agent”
key_	string	LLD rule key. Property behavior: - *required* for create operations - *read-only* for inherited objects
name	string	Name of the LLD rule. Property behavior: - *required* for create operations - *read-only* for inherited objects

Property	Type	Description
type	integer	Type of the LLD rule. Possible values: 0 - Zabbix agent; 2 - Zabbix trapper; 3 - Simple check; 5 - Zabbix internal; 7 - Zabbix agent (active); 10 - External check; 11 - Database monitor; 12 - IPMI agent; 13 - SSH agent; 14 - TELNET agent; 16 - JMX agent; 18 - Dependent item; 19 - HTTP agent; 20 - SNMP agent; 21 - Script; 22 - Browser; 23 - Nested. This type is allowed only if the rule's parent is a template (assumed to be linked to a host prototype on discovery) or discovered host (which has an LLD rule at its root). Property behavior: - *required* for create operations - *read-only* for inherited objects
url	string	URL string. Supports user macros, {HOST.IP}, {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.NAME}, {HOST.PORT}, {ITEM.ID}, {ITEM.KEY}. Property behavior: - *required* if `type` is set to "HTTP agent" - *read-only* for inherited objects
allow_traps	integer	Allow the value to be populated similarly to the trapper item. Possible values: 0 - *(default)* Do not accept incoming data; 1 - Accept incoming data. Property behavior: - *supported* if `type` is set to "HTTP agent"
authtype	integer	Authentication method. Possible values if `type` is set to "SSH agent": 0 - *(default)* password; 1 - public key. Possible values if `type` is set to "HTTP agent": 0 - *(default)* none; 1 - basic; 2 - NTLM; 3 - Kerberos; 4 - Digest. Property behavior: - *supported* if `type` is set to "SSH agent" or "HTTP agent" - *read-only* for inherited objects (if `type` is set to "HTTP agent")
description	string	Description of the LLD rule.
error	string	Error text if there are problems updating the LLD rule value. Property behavior: - *read-only*

Property	Type	Description
follow_redirects	integer	Follow response redirects while polling data. Possible values: 0 - Do not follow redirects; 1 - *(default)* Follow redirects. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
headers	array	Array of headers that will be sent when performing an HTTP request. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
http_proxy	string	HTTP(S) proxy connection string. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
ipmi_sensor	string	IPMI sensor. Property behavior: - *required* if `type` is set to ”IPMI agent” and `key_` is not set to ”ipmi.get” - *supported* if `type` is set to ”IPMI agent” - *read-only* for inherited objects
jmx_endpoint	string	JMX agent custom connection string. Default: service:jmx:rmi:///jndi/rmi://{HOST.CONN}:{HOST.PORT}/jmxrmi Property behavior: - *supported* if `type` is set to ”JMX agent”
lifetime	string	Time period after which items that are no longer discovered will be deleted. Accepts a number of seconds (e.g., 3600), or a value with a time suffix (e.g., `3600s`, `180m`, `24h`, `1d`), or a user macro. Default: `7d`.
lifetime_type	integer	Scenario to delete lost LLD resources. Possible values: 0 - *(default)* Delete after lifetime threshold is reached; 1 - Do not delete; 2 - Delete immediately.
enabled_lifetime	string	Time period after which items that are no longer discovered will be disabled. Accepts a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., `30s`, `1m`, `2h`, `1d`), or a user macro. Default: 0.
enabled_lifetime_type	integer	Scenario to disable lost LLD resources. Possible values: 0 - Disable after lifetime threshold is reached; 1 - Do not disable; 2 - *(default)* Disable immediately.
master_itemid	ID	ID of the master item. Discovery rule cannot be master item for another discovery rule. Property behavior: - *required* if `type` is set to ”Dependent item” - *read-only* for inherited objects

Property	Type	Description
output_format	integer	Should the response be converted to JSON. Possible values: 0 - *(default)* Store raw; 1 - Convert to JSON. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
params	string	Additional parameters depending on the type of the LLD rule: - executed script for SSH and Telnet LLD rules; - SQL query for database monitor LLD rules; - formula for calculated LLD rules; - the script for script and browser LLD rules. Property behavior: - *required* if `type` is set to ”Database monitor”, ”SSH agent”, ”TELNET agent”, ”Script” or ”Browser” - *read-only* for inherited objects (if `type` is set to ”Script” or ”Browser”)
parameters	object/array	Additional parameters if `type` is set to ”Script” or ”Browser”. Array of objects with `name` and `value` properties, where `name` must be unique. Property behavior: - *supported* if `type` is set to ”Script” or ”Browser” - *read-only* for inherited objects
password	string	Password for authentication. Property behavior: - *required* if `type` is set to ”JMX agent” and `username` is set - *supported* if `type` is set to ”Simple check”, ”Database monitor”, ”SSH agent”, ”TELNET agent”, or ”HTTP agent” - *read-only* for inherited objects (if `type` is set to ”HTTP agent”)
post_type	integer	Type of post data body stored in `posts` property. Possible values: 0 - *(default)* Raw data; 2 - JSON data; 3 - XML data. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
posts	string	HTTP(S) request body data. Property behavior: - *required* if `type` is set to ”HTTP agent” and `post_type` is set to ”JSON data” or ”XML data” - *supported* if `type` is set to ”HTTP agent” and `post_type` is set to ”Raw data” - *read-only* for inherited objects
privatekey	string	Name of the private key file. Property behavior: - *required* if `type` is set to ”SSH agent” and `authtype` is set to ”public key”
publickey	string	Name of the public key file. Property behavior: - *required* if `type` is set to ”SSH agent” and `authtype` is set to ”public key”

Property	Type	Description
query_fields	array	Array of query fields that will be sent when performing an HTTP request. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
request_method	integer	Type of request method. Possible values: 0 - *(default)* GET; 1 - POST; 2 - PUT; 3 - HEAD. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
retrieve_mode	integer	What part of response should be stored. Possible values if `request_method` is set to ”GET”, ”POST”, or ”PUT”: 0 - *(default)* Body; 1 - Headers; 2 - Both body and headers will be stored. Possible values if `request_method` is set to ”HEAD”: 1 - Headers. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
snmp_oid	string	SNMP OID. Property behavior: - *required* if `type` is set to ”SNMP agent” - *read-only* for inherited objects
ssl_cert_file	string	Public SSL Key file path. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
ssl_key_file	string	Private SSL Key file path. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
ssl_key_password	string	Password for SSL Key file. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
state	integer	State of the LLD rule. Possible values: 0 - *(default)* normal; 1 - not supported. Property behavior: - *read-only*

Property	Type	Description
status	integer	Status of the LLD rule. Possible values: 0 - *(default)* enabled LLD rule; 1 - disabled LLD rule.
status_codes	string	Ranges of required HTTP status codes, separated by commas. Also supports user macros as part of comma separated list. Example: 200,200-{$M},{$M},200-400 Property behavior: - *supported* if `type` is set to "HTTP agent" - *read-only* for inherited objects
templateid	ID	ID of the parent template LLD rule. Property behavior: - *read-only*
timeout	string	Item data polling request timeout. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m). Also accepts user macros. Possible values range: 1-600s. Default: "" - use proxy/global settings. Property behavior: - *supported* if `type` is set to "Zabbix agent" (0), "Simple check" (3) and `key_` does not start with "`vmware.`" and "`icmpping`", "Zabbix agent (active)" (7), "External check" (10), "Database monitor" (11), "SSH agent" (13), "TELNET agent" (14), "HTTP agent" (19), "SNMP agent" (20) and `snmp_oid` starts with "`walk[`" or "`get[`", "Script" (21), "Browser" (22) - *read-only* for inherited objects, - if the default value is applied and the item belongs to a host monitored by an inaccessible proxy or proxy group, the global timeout is used.
trapper_hosts	string	Allowed hosts. Property behavior: - *supported* if `type` is set to "Zabbix trapper", or if `type` is set to "HTTP agent" and `allow_traps` is set to "Allow incoming data". If the property is omitted, it is automatically set to the `{$TRAPPER.ALLOWED_HOSTS}` user macro (default value). By default, `{$TRAPPER.ALLOWED_HOSTS}` is `127.0.0.1,::1`. If the property is set to an empty string, incoming connections will be denied from all hosts.
username	string	Username for authentication. Property behavior: - *required* if `type` is set to "SSH agent", "TELNET agent", or if `type` is set to "JMX agent" and `password` is set - *supported* if `type` is set to "Simple check", "Database monitor", or "HTTP agent" - *read-only* for inherited objects (if `type` is set to "HTTP agent")
uuid	string	Universal unique identifier, used for linking imported LLD rules to already existing ones. Auto-generated, if not given. Property behavior: - *supported* if the LLD rule belongs to a template

Property	Type	Description
verify_host	integer	Whether to validate that the host name for the connection matches the one in the host's certificate. Possible values: 0 - *(default)* Do not validate; 1 - Validate. Property behavior: - *supported* if `type` is set to "HTTP agent" - *read-only* for inherited objects
verify_peer	integer	Whether to validate that the host's certificate is authentic. Possible values: 0 - *(default)* Do not validate; 1 - Validate. Property behavior: - *supported* if `type` is set to "HTTP agent" - *read-only* for inherited objects

HTTP header

The header object has the following properties:

Property	Type	Description
name	string	HTTP header name. Property behavior: - *required*
value	string	Header value. Property behavior: - *required*

HTTP query field

The query field object defines a name and value that is used to specify a URL parameter. It has the following properties:

Property	Type	Description
name	string	Name of the parameter. Property behavior: - *required*
value	string	Parameter value. Property behavior: - *required*

LLD rule filter

The LLD rule filter object defines a set of conditions that can be used to filter discovered objects. It has the following properties:

Property	Type	Description
conditions	object/array	Set of filter conditions to use for filtering results. The conditions will be sorted in the order of their placement in the formula. Property behavior: - *required*

Property	Type	Description
evaltype	integer	Filter condition evaluation method. Possible values: 0 - And/Or; 1 - And; 2 - Or; 3 - Custom expression. Property behavior: - *required*
eval_formula	string	Generated expression that will be used for evaluating filter conditions. The expression contains IDs that reference specific filter conditions by its `formulaid`. The value of `eval_formula` is equal to the value of `formula` for filters with a custom expression. Property behavior: - *read-only*
formula	string	User-defined expression to be used for evaluating conditions of filters with a custom expression. The expression must contain IDs that reference specific filter conditions by its `formulaid`. The IDs used in the expression must exactly match the ones defined in the filter conditions: no condition can remain unused or omitted. Property behavior: - *required* if `evaltype` is set to ”custom expression”

LLD rule filter condition

The LLD rule filter condition object defines a separate check to perform on the value of an LLD macro. It has the following properties:

Property	Type	Description
macro	string	LLD macro to perform the check on. Property behavior: - *required*
value	string	Value to compare with. Property behavior: - *required* if `operator` is set to ”matches regular expression” or ”does not match regular expression”
formulaid	string	Arbitrary unique ID that is used to reference the condition from a custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward. Property behavior: - *required* if `evaltype` of LLD rule filter object is set to ”custom expression”
operator	integer	Condition operator. Possible values: 8 - *(default)* matches regular expression; 9 - does not match regular expression; 12 - exists; 13 - does not exist.

Note:
To better understand how to use filters with various types of expressions, see examples on the discoveryrule.get and discoveryrule.create method pages.

LLD macro path

The LLD macro path has the following properties:

Property	Type	Description
lld_macro	string	LLD macro. Property behavior: - *required*
path	string	Selector for value which will be assigned to corresponding macro. Property behavior: - *required*

LLD rule preprocessing

The LLD rule preprocessing object has the following properties.

Property	Type	Description
type	integer	The preprocessing option type. Possible values: 5 - Regular expression; 11 - XML XPath; 12 - JSONPath; 14 - Matches regular expression; 15 - Does not match regular expression; 16 - Check for error in JSON; 17 - Check for error in XML; 20 - Discard unchanged with heartbeat; 21 - JavaScript; 23 - Prometheus to JSON; 24 - CSV to JSON; 25 - Replace; 27 - XML to JSON; 28 - SNMP walk value; 29 - SNMP walk to JSON; 30 - SNMP get value. Property behavior: - *required*
params	string	Additional parameters used by preprocessing option. Multiple parameters are separated by the newline (\n) character. Property behavior: - *required* if `type` is set to ”Regular expression” (5), ”XML XPath” (11), ”JSONPath” (12), ”Matches regular expression” (14), ”Does not match regular expression” (15), ”Check for error in JSON” (16), ”Check for error in XML” (17), ”Discard unchanged with heartbeat” (20), ”JavaScript” (21), ”Prometheus to JSON” (23), ”CSV to JSON” (24), ”Replace” (25), ”SNMP walk value” (28), ”SNMP walk to JSON” (29), or ”SNMP get value” (30)

Property	Type	Description
error_handler	integer	Action type used in case of preprocessing step failure. Possible values: 0 - Error message is set by Zabbix server; 1 - Discard value; 2 - Set custom value; 3 - Set custom error message. Property behavior: - *required* if `type` is set to ”Regular expression” (5), ”XML XPath” (11), ”JSONPath” (12), ”Matches regular expression” (14), ”Does not match regular expression” (15), ”Check for error in JSON” (16), ”Check for error in XML” (17), ”Prometheus to JSON” (23), ”CSV to JSON” (24), ”XML to JSON” (27), ”SNMP walk value” (28), ”SNMP walk to JSON” (29), or ”SNMP get value” (30)
error_handler_params	string	Error handler parameters. Property behavior: - *required* if `error_handler` is set to ”Set custom value” or ”Set custom error message”

The following parameters and error handlers are supported for each preprocessing type.

Preprocessing type	Name	Parameter 1	Parameter 2	Parameter 3	Supported error handlers
5	Regular expression	pattern[1]	output[2]		0, 1, 2, 3
11	XML XPath	path[3]			0, 1, 2, 3
12	JSONPath	path[3]			0, 1, 2, 3
14	Matches regular expression	pattern[1]			0, 1, 2, 3
15	Does not match regular expression	pattern[1]			0, 1, 2, 3
16	Check for error in JSON	path[3]			0, 1, 2, 3
17	Check for error in XML	path[3]			0, 1, 2, 3

Preprocessing type	Name	Parameter 1	Parameter 2	Parameter 3	Supported error handlers
20	Discard unchanged with heartbeat	seconds^4, 5^			
21	JavaScript	script[2]			
23	Prometheus to JSON	pattern^5, 6^			0, 1, 2, 3
24	CSV to JSON	character[2]	character[2]	0,1	0, 1, 2, 3
25	Replace	search string[2]	replacement[2]		
27	XML to JSON				0, 1, 2, 3
28	SNMP walk value	OID[2]	Format: 0 - Unchanged 1 - UTF-8 from Hex-STRING 2 - MAC from Hex-STRING 3 - Integer from BITS		0, 1, 2, 3
29	SNMP walk to JSON[7]	Field name[2]	OID prefix[2]	Format: 0 - Unchanged 1 - UTF-8 from Hex-STRING 2 - MAC from Hex-STRING 3 - Integer from BITS	0, 1, 2, 3
30	SNMP get value	Format: 1 - UTF-8 from Hex-STRING 2 - MAC from Hex-STRING 3 - Integer from BITS			0, 1, 2, 3

[1] regular expression
 [2] string
 [3] JSONPath or XML XPath
 [4] positive integer (with support of time suffixes, e.g., 30s, 1m, 2h, 1d)
 [5] user macro
 [6] Prometheus pattern following the syntax: `<metric name>{<label name>="<label value>", ...} == <value>`. Each Prometheus pattern component (metric, label name, label value and metric value) can be user macro.
 [7] Supports multiple "Field name,OID prefix,Format records" records delimited by a new line character.

LLD rule overrides

The LLD rule overrides object defines a set of rules (filters, conditions and operations) that are used to override properties of different prototype objects. It has the following properties:

Property	Type	Description
name	string	Unique override name. Property behavior: - *required*
step	integer	Unique order number of the override. Property behavior: - *required*

Property	Type	Description
stop	integer	Stop processing next overrides if matches. Possible values: 0 - *(default)* don't stop processing overrides; 1 - stop processing overrides if filter matches.
filter	object	Override filter.
operations	object/array	Override operations.

LLD rule override filter

The LLD rule override filter object defines a set of conditions that if they match the discovered object the override is applied. It has the following properties:

Property	Type	Description
conditions	object/array	Set of override filter conditions to use for matching the discovered objects. The conditions will be sorted in the order of their placement in the formula. Property behavior: - *required*
evaltype	integer	Override filter condition evaluation method. Possible values: 0 - And/Or; 1 - And; 2 - Or; 3 - Custom expression. Property behavior: - *required*
eval_formula	string	Generated expression that will be used for evaluating override filter conditions. The expression contains IDs that reference specific override filter conditions by its `formulaid`. The value of `eval_formula` is equal to the value of `formula` for filters with a custom expression. Property behavior: - *read-only*
formula	string	User-defined expression to be used for evaluating conditions of override filters with a custom expression. The expression must contain IDs that reference specific override filter conditions by its `formulaid`. The IDs used in the expression must exactly match the ones defined in the override filter conditions: no condition can remain unused or omitted. Property behavior: - *required* if `evaltype` is set to "custom expression"

LLD rule override filter condition

The LLD rule override filter condition object defines a separate check to perform on the value of an LLD macro. It has the following properties:

Property	Type	Description
macro	string	LLD macro to perform the check on. Property behavior: - *required*

Property	Type	Description
value	string	Value to compare with. Property behavior: - *required* if `operator` is set to ”matches regular expression” or ”does not match regular expression”
formulaid	string	Arbitrary unique ID that is used to reference the condition from a custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward. Property behavior: - *required* if `evaltype` of LLD rule override filter object is set to ”custom expression”
operator	integer	Condition operator. Possible values: 8 - *(default)* matches regular expression; 9 - does not match regular expression; 12 - exists; 13 - does not exist.

LLD rule override operation

The LLD rule override operation is combination of conditions and actions to perform on the prototype object. It has the following properties:

Property	Type	Description
operationobject	integer	Type of discovered object to perform the action. Possible values: 0 - Item prototype; 1 - Trigger prototype; 2 - Graph prototype; 3 - Host prototype; 4 - Discovery prototype. Property behavior: - *required*
operator	integer	Override condition operator. Possible values: 0 - *(default)* equals; 1 - does not equal; 2 - contains; 3 - does not contain; 8 - matches; 9 - does not match.
value	string	Pattern to match item, trigger, graph or host prototype name depending on selected object.
opstatus	object	Override operation status object for item, trigger and host prototype objects.
opdiscover	object	Override operation discover status object (all object types).
opperiod	object	Override operation period (update interval) object for item prototype object.
ophistory	object	Override operation history object for item prototype object.
optrends	object	Override operation trends object for item prototype object.
opseverity	object	Override operation severity object for trigger prototype object.
optag	object/array	Override operation tag object for trigger and host prototype objects.
optemplate	object/array	Override operation template object for host prototype object.
opinventory	object	Override operation inventory object for host prototype object.

LLD rule override operation status

LLD rule override operation status that is set to discovered object. It has the following properties:

Property	Type	Description
status	integer	Override the status for selected object. Possible values: 0 - Create enabled; 1 - Create disabled. Property behavior: - *required*

LLD rule override operation discover

LLD rule override operation discover status that is set to discovered object. It has the following properties:

Property	Type	Description
discover	integer	Override the discover status for selected object. Possible values: 0 - Yes, continue discovering the objects; 1 - No, new objects will not be discovered and existing ones will be marked as lost. Property behavior: - *required*

LLD rule override operation period

LLD rule override operation period is an update interval value that is set to discovered item. It has the following properties:

Property	Type	Description
delay	string	Override the update interval of the item prototype. Accepts a number of seconds (e.g., `1`, `60`, `300`), a value with a time suffix (e.g., `30s`, `1m`, `2h`, `1d`), or, optionally, one or more custom intervals, all separated by semicolons. Custom intervals can be a mix of flexible and scheduling intervals. Accepts user macros or LLD macros. If used, the value must be a single macro. Multiple macros or macros mixed with text are not supported. Flexible intervals may be written as two macros separated by a forward slash (e.g., `{$FLEX_INTERVAL}/{$FLEX_PERIOD}`). Example: `1h;wd1-5h9-18;{$Macro1}/1-7,00:00-24:00;0/6-7,12:00-24:00;{$Macr` Property behavior: - *required*

LLD rule override operation history

LLD rule override operation history value that is set to discovered item. It has the following properties:

Property	Type	Description
history	string	Override the history of item prototype which is a time unit of how long the history data should be stored. Also accepts user macro and LLD macro. Property behavior: - *required*

LLD rule override operation trends

LLD rule override operation trends value that is set to discovered item. It has the following properties:

Property	Type	Description
trends	string	Override the trends of item prototype which is a time unit of how long the trends data should be stored. Also accepts user macro and LLD macro. Property behavior: - *required*

LLD rule override operation severity

LLD rule override operation severity value that is set to discovered trigger. It has the following properties:

Property	Type	Description
severity	integer	Override the severity of trigger prototype. Possible values: 0 - *(default)* not classified; 1 - information; 2 - warning; 3 - average; 4 - high; 5 - disaster. Property behavior: - *required*

LLD rule override operation tag

LLD rule override operation tag object contains tag name and value that are set to discovered object. It has the following properties:

Property	Type	Description
tag	string	New tag name. Also accepts user macro and LLD macro. Property behavior: - *required*
value	string	New tag value. Also accepts user macro and LLD macro.

LLD rule override operation template

LLD rule override operation template object that is linked to discovered host. It has the following properties:

Property	Type	Description
templateid	ID	Override the template of host prototype linked templates. Property behavior: - *required*

LLD rule override operation inventory

LLD rule override operation inventory mode value that is set to discovered host. It has the following properties:

Property	Type	Description
inventory_mode	integer	Override the host prototype inventory mode. Possible values: -1 - disabled; 0 - *(default)* manual; 1 - automatic. Property behavior: - *required*

discoveryrule.create

Description

`object discoveryrule.create(object/array lldRules)`

This method allows you to create new LLD rules.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` LLD rules to create.

Additionally to the standard LLD rule properties, the method accepts the following parameters.

Parameter	Type	Description
filter	object	LLD rule filter for the LLD rule.
preprocessing	object/array	LLD rule preprocessing options.
lld_macro_paths	object/array	LLD rule lld_macro_path options.
overrides	object/array	LLD rule overrides options.

Return values

`(object)` Returns an object containing the IDs of the created LLD rules under the `itemids` property. The order of the returned IDs matches the order of the passed LLD rules.

Examples

Creating an LLD rule

Create a Zabbix agent LLD rule to discover mounted file systems. Discovered items will be updated every 30 seconds.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.create",
    "params": {
        "name": "Mounted filesystem discovery",
        "key_": "vfs.fs.discovery",
        "hostid": "10197",
        "type": 0,
        "interfaceid": "112",
        "delay": "30s"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "27665"
        ]
    },
    "id": 1
}
```

Using a filter

Create an LLD rule with a set of conditions to filter the results by. The conditions will be grouped together using the logical "and" operator.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.create",
    "params": {
        "name": "Filtered LLD rule",
        "key_": "lld",
        "hostid": "10116",
        "type": 0,
        "interfaceid": "13",
        "delay": "30s",
        "filter": {
            "evaltype": 1,
            "conditions": [
                {
                    "macro": "{#MACRO1}",
                    "value": "@regex1"
                },
                {
                    "macro": "{#MACRO2}",
                    "value": "@regex2",
                    "operator": "9"
                },
                {
                    "macro": "{#MACRO3}",
                    "value": "",
                    "operator": "12"
                },
                {
                    "macro": "{#MACRO4}",
                    "value": "",
                    "operator": "13"
                }
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "27665"
        ]
    },
```

```
    "id": 1
}
```

Creating an LLD rule with macro paths

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.create",
    "params": {
        "name": "LLD rule with LLD macro paths",
        "key_": "lld",
        "hostid": "10116",
        "type": 0,
        "interfaceid": "13",
        "delay": "30s",
        "lld_macro_paths": [
            {
                "lld_macro": "{#MACRO1}",
                "path": "$.path.1"
            },
            {
                "lld_macro": "{#MACRO2}",
                "path": "$.path.2"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "27665"
        ]
    },
    "id": 1
}
```

Using a custom expression filter

Create an LLD rule with a filter that will use a custom expression to evaluate the conditions. The LLD rule must only discover objects the "{#MACRO1}" macro value of which matches both regular expression "regex1" and "regex2", and the value of "{#MACRO2}" matches either "regex3" or "regex4". The formula IDs "A", "B", "C" and "D" have been chosen arbitrarily.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.create",
    "params": {
        "name": "Filtered LLD rule",
        "key_": "lld",
        "hostid": "10116",
        "type": 0,
        "interfaceid": "13",
        "delay": "30s",
        "filter": {
            "evaltype": 3,
            "formula": "(A and B) and (C or D)",
            "conditions": [
                {
                    "macro": "{#MACRO1}",
```

```
                    "value": "@regex1",
                    "formulaid": "A"
                },
                {
                    "macro": "{#MACRO1}",
                    "value": "@regex2",
                    "formulaid": "B"
                },
                {
                    "macro": "{#MACRO2}",
                    "value": "@regex3",
                    "formulaid": "C"
                },
                {
                    "macro": "{#MACRO2}",
                    "value": "@regex4",
                    "formulaid": "D"
                }
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "27665"
        ]
    },
    "id": 1
}
```

Using custom query fields and headers

Create LLD rule with custom query fields and headers.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.create",
    "params": {
        "hostid": "10257",
        "interfaceid": "5",
        "type": 19,
        "name": "API HTTP agent",
        "key_": "api_discovery_rule",
        "delay": "5s",
        "url": "http://127.0.0.1?discoverer.php",
        "query_fields": [
            {
                "name": "mode",
                "value": "json"
            },
            {
                "name": "elements",
                "value": "2"
            }
        ],
        "headers": [
            {
```

```
                "name": "X-Type",
                "value": "api"
            },
            {
                "name": "Authorization",
                "value": "Bearer mF_A.B5f-2.1JcM"
            }
        ],
        "allow_traps": 1,
        "trapper_hosts": "127.0.0.1"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "28336"
        ]
    },
    "id": 1
}
```

Creating an LLD rule with preprocessing

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.create",
    "params": {
        "name": "Discovery rule with preprocessing",
        "key_": "lld.with.preprocessing",
        "hostid": "10001",
        "ruleid": "27665",
        "type": 0,
        "delay": "60s",
        "interfaceid": "1155",
        "preprocessing": [
            {
                "type": 20,
                "params": "20",
                "error_handler": 0,
                "error_handler_params": ""
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "44211"
        ]
    },
    "id": 1
}
```

Creating an LLD rule with overrides

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.create",
    "params": {
        "name": "Discover database host",
        "key_": "lld.with.overrides",
        "hostid": "10001",
        "type": 0,
        "delay": "60s",
        "interfaceid": "1155",
        "overrides": [
            {
                "name": "Discover MySQL host",
                "step": "1",
                "stop": "1",
                "filter": {
                    "evaltype": "2",
                    "conditions": [
                        {
                            "macro": "{#UNIT.NAME}",
                            "operator": "8",
                            "value": "^mysqld\\.service$"
                        },
                        {
                            "macro": "{#UNIT.NAME}",
                            "operator": "8",
                            "value": "^mariadb\\.service$"
                        }
                    ]
                },
                "operations": [
                    {
                        "operationobject": "3",
                        "operator": "2",
                        "value": "Database host",
                        "opstatus": {
                            "status": "0"
                        },
                        "optemplate": [
                            {
                                "templateid": "10170"
                            }
                        ],
                        "optag": [
                            {
                                "tag": "database",
                                "value": "mysql"
                            }
                        ]
                    }
                ]
            },
            {
                "name": "Discover PostgreSQL host",
                "step": "2",
                "stop": "1",
                "filter": {
                    "evaltype": "0",
                    "conditions": [
                        {
                            "macro": "{#UNIT.NAME}",
```

```
                            "operator": "8",
                            "value": "^postgresql\\.service$"
                        }
                    ]
                },
                "operations": [
                    {
                        "operationobject": "3",
                        "operator": "2",
                        "value": "Database host",
                        "opstatus": {
                            "status": "0"
                        },
                        "optemplate": [
                            {
                                "templateid": "10263"
                            }
                        ],
                        "optag": [
                            {
                                "tag": "database",
                                "value": "postgresql"
                            }
                        ]
                    }
                ]
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "30980"
        ]
    },
    "id": 1
}
```

Create script LLD rule

Create a simple data collection using a script LLD rule.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.create",
    "params": {
        "name": "Script example",
        "key_": "custom.script.lldrule",
        "hostid": "12345",
        "type": 21,
        "params": "var request = new HttpRequest();\nreturn request.post(\"https://postman-echo.com/post\"
        "parameters": [{
            "name": "host",
            "value": "{HOST.CONN}"
        }],
        "timeout": "6s",
        "delay": "30s"
```

```
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "23865"
        ]
    },
    "id": 1
}
```

Create LLD rule with a specified time period for disabling and no deletion

Create an LLD rule with custom time period for disabling entity after it is no longer discovered, with the setting that it will never be deleted.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.create",
    "params": {
        "name": "lld disable after 1h",
        "key_": "lld.disable",
        "hostid": "10001",
        "type": 2,
        "lifetime_type": 1,
        "enabled_lifetime_type": 0,
        "enabled_lifetime": "1h"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "46864"
        ]
    },
    "id": 1
}
```

See also

- LLD rule filter
- LLD macro paths
- LLD rule preprocessing

Source

CDiscoveryRule::create() in *ui/include/classes/api/services/CDiscoveryRule.php*.

discoveryrule.delete

Description

`object discoveryrule.delete(array lldRuleIds)`

This method allows you to delete LLD rules.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.
 The user will not be able to delete a LLD rule that discover entities using inaccessible proxy. See user group configuration (*Proxy access list*) for more information.

Parameters

`(array)` IDs of the LLD rules to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted LLD rules under the `ruleids` property.

Examples

Deleting multiple LLD rules

Delete two LLD rules.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.delete",
    "params": [
        "27665",
        "27668"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "ruleids": [
            "27665",
            "27668"
        ]
    },
    "id": 1
}
```

Source

CDiscoveryRule::delete() in *ui/include/classes/api/services/CDiscoveryRule.php*.

discoveryrule.get

Description

`integer/array discoveryrule.get(object parameters)`

The method allows you to retrieve LLD rules according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
itemids	ID/array	Return only LLD rules with the given IDs.
groupids	ID/array	Return only LLD rules that belong to the hosts from the given groups.

Parameter	Type	Description
hostids	ID/array	Return only LLD rules that belong to the given hosts.
inherited	boolean	If set to `true` return only LLD rules inherited from a template.
interfaceids	ID/array	Return only LLD rules use the given host interfaces.
monitored	boolean	If set to `true` return only enabled LLD rules that belong to monitored hosts.
templated	boolean	If set to `true` return only LLD rules that belong to templates.
templateids	ID/array	Return only LLD rules that belong to the given templates.
selectDiscoveryData	query	Return a `discoveryData` property with the LLD rule discovery object data. The LLD rule discovery object links a discovered LLD rule to an LLD rule prototype from which it was discovered.
selectDiscoveryRulePrototypes	query	Return a `discoveryRulePrototypes` property with LLD rule prototypes that belong to the LLD rule. Supports `count`.
selectFilter	query	Return a `filter` property with data of the filter used by the LLD rule.
selectGraphs	query	Returns a `graphs` property with graph prototypes that belong to the LLD rule. Supports `count`.
selectHostPrototypes	query	Return a `hostPrototypes` property with host prototypes that belong to the LLD rule. Supports `count`.
selectHosts	query	Return a `hosts` property with an array of hosts that the LLD rule belongs to.
selectItems	query	Return an `items` property with item prototypes that belong to the LLD rule. Supports `count`.
selectTriggers	query	Return a `triggers` property with trigger prototypes that belong to the LLD rule. Supports `count`.
selectLLDMacroPaths	query	Return an `lld_macro_paths` property with a list of LLD macros and paths to values assigned to each corresponding macro.
selectPreprocessing	query	Return a `preprocessing` property with LLD rule preprocessing options.
selectOverrides	query	Return an `lld_rule_overrides` property with a list of override filters, conditions and operations that are performed on prototype objects.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type. Supports additional properties: `host` - technical name of the host that the LLD rule belongs to.
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectItems`, `selectGraphs`, `selectTriggers`.
sortfield	string/array	Sort the result by the given properties. Possible values: `itemid`, `name`, `key_`, `delay`, `type`, `status`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
limit	integer	
output	query	

Parameter	Type	Description
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

(integer/array) Returns either:

- an array of objects
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving discovery rules from a host

Retrieve all discovery rules for specific host ID.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.get",
    "params": {
        "output": "extend",
        "hostids": "10202"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "itemid": "27425",
            "type": "0",
            "snmp_oid": "",
            "hostid": "10202",
            "name": "Network interface discovery",
            "key_": "net.if.discovery",
            "delay": "1h",
            "status": "0",
            "trapper_hosts": "",
            "templateid": "22444",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "1",
            "interfaceid": "119",
            "description": "Discovery of network interfaces as defined in global regular expression \"Netw
            "lifetime": "30d",
            "jmx_endpoint": "",
            "master_itemid": "0",
            "timeout": "",
            "url": "",
            "query_fields": [],
```

```
            "posts": "",
            "status_codes": "200",
            "follow_redirects": "1",
            "post_type": "0",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "0",
            "request_method": "0",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
            "state": "0",
            "error": "",
            "parameters": []
        },
        {
            "itemid": "27426",
            "type": "0",
            "snmp_oid": "",
            "hostid": "10202",
            "name": "Mounted filesystem discovery",
            "key_": "vfs.fs.discovery",
            "delay": "1h",
            "status": "0",
            "trapper_hosts": "",
            "templateid": "22450",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "1",
            "interfaceid": "119",
            "description": "Discovery of file systems of different types as defined in global regular expr
            "lifetime": "30d",
            "jmx_endpoint": "",
            "master_itemid": "0",
            "timeout": "",
            "url": "",
            "query_fields": [],
            "posts": "",
            "status_codes": "200",
            "follow_redirects": "1",
            "post_type": "0",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "0",
            "request_method": "0",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
```

```
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
            "state": "0",
            "error": "",
            "parameters": []
        }
    ],
    "id": 1
}
```

Retrieving filter conditions

Retrieve the name of the LLD rule "24681" and its filter conditions. The filter uses the "and" evaluation type, so the `formula` property is empty and `eval_formula` is generated automatically.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.get",
    "params": {
        "output": ["name"],
        "selectFilter": "extend",
        "itemids": ["24681"]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "itemid": "24681",
            "name": "Filtered LLD rule",
            "filter": {
                "evaltype": "1",
                "formula": "",
                "conditions": [
                    {
                        "macro": "{#MACRO1}",
                        "value": "@regex1",
                        "operator": "8",
                        "formulaid": "A"
                    },
                    {
                        "macro": "{#MACRO2}",
                        "value": "@regex2",
                        "operator": "9",
                        "formulaid": "B"
                    },
                    {
                        "macro": "{#MACRO3}",
                        "value": "",
                        "operator": "12",
                        "formulaid": "C"
                    },
                    {
```

```
                        "macro": "{#MACRO4}",
                        "value": "",
                        "operator": "13",
                        "formulaid": "D"
                    }
                ],
                "eval_formula": "A and B and C and D"
            }
        }
    ],
    "id": 1
}
```

Retrieve LLD rule by URL

Retrieve LLD rule for host by rule URL field value. Only exact match of URL string defined for LLD rule is supported.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.get",
    "params": {
        "hostids": "10257",
        "filter": {
            "type": 19,
            "url": "http://127.0.0.1/discoverer.php"
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "itemid": "28336",
            "type": "19",
            "snmp_oid": "",
            "hostid": "10257",
            "name": "API HTTP agent",
            "key_": "api_discovery_rule",
            "delay": "5s",
            "status": "0",
            "trapper_hosts": "",
            "templateid": "0",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "1",
            "interfaceid": "5",
            "description": "",
            "lifetime": "30d",
            "jmx_endpoint": "",
            "master_itemid": "0",
            "timeout": "",
            "url": "http://127.0.0.1/discoverer.php",
            "query_fields": [
```

```
                {
                    "name": "mode",
                    "value": "json"
                },
                {
                    "name": "elements",
                    "value": "2"
                }
            ],
            "posts": "",
            "status_codes": "200",
            "follow_redirects": "1",
            "post_type": "0",
            "http_proxy": "",
            "headers": [
                {
                    "name" : "X-Type",
                    "value": "api"
                },
                {
                    "name": "Authorization",
                    "value": "Bearer mF_A.B5f-2.1JcM"
                }
            ],
            "retrieve_mode": "0",
            "request_method": "1",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
            "state": "0",
            "error": "",
            "parameters": []
        }
    ],
    "id": 1
}
```


Retrieve LLD rule with overrides

Retrieve one LLD rule that has various override settings.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.get",
    "params": {
        "output": ["name"],
        "itemids": "30980",
        "selectOverrides": ["name", "step", "stop", "filter", "operations"]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "name": "Discover database host",
            "overrides": [
                {
                    "name": "Discover MySQL host",
                    "step": "1",
                    "stop": "1",
                    "filter": {
                        "evaltype": "2",
                        "formula": "",
                        "conditions": [
                            {
                                "macro": "{#UNIT.NAME}",
                                "operator": "8",
                                "value": "^mysqld\\.service$",
                                "formulaid": "A"
                            },
                            {
                                "macro": "{#UNIT.NAME}",
                                "operator": "8",
                                "value": "^mariadb\\.service$",
                                "formulaid": "B"
                            }
                        ],
                        "eval_formula": "A or B"
                    },
                    "operations": [
                        {
                            "operationobject": "3",
                            "operator": "2",
                            "value": "Database host",
                            "opstatus": {
                                "status": "0"
                            },
                            "optag": [
                                {
                                    "tag": "database",
                                    "value": "mysql"
                                }
                            ],
                            "optemplate": [
                                {
                                    "templateid": "10170"
                                }
                            ]
                        }
                    ]
                },
                {
                    "name": "Discover PostgreSQL host",
                    "step": "2",
                    "stop": "1",
                    "filter": {
                        "evaltype": "0",
                        "formula": "",
                        "conditions": [
                            {
                                "macro": "{#UNIT.NAME}",
                                "operator": "8",
```

```
                                "value": "^postgresql\\.service$",
                                "formulaid": "A"
                            }
                        ],
                        "eval_formula": "A"
                    },
                    "operations": [
                        {
                            "operationobject": "3",
                            "operator": "2",
                            "value": "Database host",
                            "opstatus": {
                                "status": "0"
                            },
                            "optag": [
                                {
                                    "tag": "database",
                                    "value": "postgresql"
                                }
                            ],
                            "optemplate": [
                                {
                                    "templateid": "10263"
                                }
                            ]
                        }
                    ]
                }
            ]
        }
    ],
    "id": 1
}
```


See also

- Graph prototype
- Host
- Item prototype
- LLD rule filter
- Trigger prototype

Source

CDiscoveryRule::get() in *ui/include/classes/api/services/CDiscoveryRule.php*.

discoveryrule.update

Description

`object discoveryrule.update(object/array lldRules)`

This method allows you to update existing LLD rules.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` LLD rule properties to be updated.

The `itemid` property must be defined for each LLD rule, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard LLD rule properties, the method accepts the following parameters.

Parameter	Type	Description
filter	object	LLD rule filter to replace the existing filter.
preprocessing	object/array	LLD rule preprocessing options to replace the existing preprocessing options. Parameter behavior: - *read-only* for inherited objects
lld_macro_paths	object/array	LLD rule lld_macro_path options to replace the existing lld_macro_path options. Parameter behavior: - *read-only* for inherited objects
overrides	object/array	LLD rule overrides options to replace the existing overrides options. Parameter behavior: - *read-only* for inherited objects

Return values

`(object)` Returns an object containing the IDs of the updated LLD rules under the `itemids` property.

Examples

Adding a filter to an LLD rule

Add a filter so that the contents of the *{#FSTYPE}* macro would match the *@File systems for discovery* regexp.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.update",
    "params": {
        "itemid": "22450",
        "filter": {
            "evaltype": 1,
            "conditions": [
                {
                    "macro": "{#FSTYPE}",
                    "value": "@File systems for discovery"
                }
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "22450"
        ]
    },
    "id": 1
}
```

Adding LLD macro paths

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.update",
    "params": {
```

```
        "itemid": "22450",
        "lld_macro_paths": [
            {
                "lld_macro": "{#MACRO1}",
                "path": "$.json.path"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "22450"
        ]
    },
    "id": 1
}
```

Disable trapping

Disable LLD trapping for discovery rule.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.update",
    "params": {
        "itemid": "28336",
        "allow_traps": 0
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "28336"
        ]
    },
    "id": 1
}
```

Updating LLD rule preprocessing options

Update an LLD rule with preprocessing rule ”JSONPath”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.update",
    "params": {
        "itemid": "44211",
        "preprocessing": [
            {
                "type": 12,
                "params": "$.path.to.json",
                "error_handler": 2,
                "error_handler_params": "5"
```

```
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "44211"
        ]
    },
    "id": 1
}
```

Updating LLD rule script

Update an LLD rule script with a different script and remove unnecessary parameters that were used by previous script.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.update",
    "params": {
        "itemid": "23865",
        "parameters": [],
        "script": "Zabbix.log(3, 'Log test');\nreturn 1;"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "23865"
        ]
    },
    "id": 1
}
```

Updating LLD rule lifetime

Update LLD rule to disable no-longer discovered entity after 12 hours and to delete it after 7 days.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryrule.update",
    "params": {
        "itemid": "46864",
        "lifetime_type": 0,
        "lifetime": "7d",
        "enabled_lifetime_type": 0,
        "enabled_lifetime": "12h"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "46864"
        ]
    },
    "id": 1
}
```

Source

CDiscoveryRule::update() in *ui/include/classes/api/services/CDiscoveryRule.php*.

LLD rule prototype

This class is designed to work with low-level discovery rule prototypes.

Object references:

- LLD rule prototype
  - HTTP header
  - HTTP query field
- LLD rule prototype filter
  - LLD rule prototype filter condition
- LLD prototype macro path
- LLD rule prototype preprocessing
- LLD rule prototype overrides
  - LLD rule prototype override filter
    - LLD rule prototype override filter condition
  - LLD rule prototype override operation
    - LLD rule prototype override operation status
    - LLD rule prototype override operation discover
    - LLD rule prototype override operation period
    - LLD rule prototype override operation history
    - LLD rule prototype override operation trends
    - LLD rule prototype override operation severity
    - LLD rule prototype override operation tag
    - LLD rule prototype override operation template
    - LLD rule prototype override operation inventory

Available methods:

- discoveryruleprototype.create - create new LLD rule prototypes
- discoveryruleprototype.delete - delete LLD rule prototypes
- discoveryruleprototype.get - retrieve LLD rule prototypes
- discoveryruleprototype.update - update LLD rule prototypes

LLD rule prototype object

The following objects are directly related to the `discoveryruleprototype` API.

LLD rule prototype

The low-level discovery rule prototype object has the following properties.

Property	Type	Description
itemid	ID	ID of the LLD rule prototype. Property behavior: - *read-only* - *required* for update operations

Property	Type	Description
ruleid	ID	ID of the parent LLD rule/LLD rule prototype. Property behavior: - *read-only* - *required* for create operations
delay	string	Update interval of the LLD rule prototype. Accepts a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., 30s, 1m, 2h, 1d), or, optionally, one or more custom intervals, all separated by semicolons. Custom intervals can be a mix of flexible and scheduling intervals. Accepts user macros. If used, the value must be a single macro. Multiple macros or macros mixed with text are not supported. Flexible intervals may be written as two macros separated by a forward slash (e.g., {$FLEX_INTERVAL}/{$FLEX_PERIOD}). Example: 1h;wd1-5h9-18;{$Macro1}/1-7,00:00-24:00;0/6-7,12:00-24:00;{$Macr Property behavior: - *required* if type is set to ”Zabbix agent” (0), ”Simple check” (3), ”Zabbix internal” (5), ”External check” (10), ”Database monitor” (11), ”IPMI agent” (12), ”SSH agent” (13), ”TELNET agent” (14), ”JMX agent” (16), ”HTTP agent” (19), ”SNMP agent” (20), ”Script” (21), ”Browser” (22), or if type is set to ”Zabbix agent (active)” (7) and key_ does not contain ”mqtt.get”
hostid	ID	ID of the host that the LLD rule prototype belongs to. Property behavior: - *constant* - *required* for create operations
flags	integer	Origin of the LLD rule prototype. Possible values: 3 - a low-level discovery rule prototype; 7 - a discovered low-level discovery rule prototype. Property behavior: - *read-only*
interfaceid	ID	ID of the LLD rule prototype’s host interface. Property behavior: - *required* if LLD rule prototype belongs to host and type is set to ”Zabbix agent”, ”IPMI agent”, ”JMX agent”, or ”SNMP agent” - *supported* if LLD rule prototype belongs to host and type is set to ”Simple check”, ”External check”, ”SSH agent”, ”TELNET agent”, or ”HTTP agent”
key_	string	LLD rule prototype key. At least one LLD macro is required. Property behavior: - *required* for create operations - *read-only* for inherited objects
name	string	Name of the LLD rule prototype. Property behavior: - *required* for create operations - *read-only* for inherited objects

Property	Type	Description
type	integer	Type of the LLD rule prototype. Possible values: 0 - Zabbix agent; 2 - Zabbix trapper; 3 - Simple check; 5 - Zabbix internal; 7 - Zabbix agent (active); 10 - External check; 11 - Database monitor; 12 - IPMI agent; 13 - SSH agent; 14 - TELNET agent; 16 - JMX agent; 18 - Dependent item; 19 - HTTP agent; 20 - SNMP agent; 21 - Script; 22 - Browser; 23 - Nested. Property behavior: - *required* for create operations - *read-only* for inherited objects
url	string	URL string. Supports user macros, {HOST.IP}, {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.NAME}, {HOST.PORT}, {ITEM.ID}, {ITEM.KEY}. Property behavior: - *required* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
allow_traps	integer	Allow the value to be populated similarly to the trapper item. Possible values: 0 - *(default)* Do not accept incoming data; 1 - Accept incoming data. Property behavior: - *supported* if `type` is set to ”HTTP agent”
authtype	integer	Authentication method. Possible values if `type` is set to ”SSH agent”: 0 - *(default)* password; 1 - public key. Possible values if `type` is set to ”HTTP agent”: 0 - *(default)* none; 1 - basic; 2 - NTLM. Property behavior: - *supported* if `type` is set to ”SSH agent” or ”HTTP agent” - *read-only* for inherited objects (if `type` is set to ”HTTP agent”)
description	string	Description of the LLD rule prototype.
error	string	Error text if there are problems updating the LLD rule prototype value. Property behavior: - *read-only*

Property	Type	Description
follow_redirects	integer	Follow response redirects while polling data. Possible values: 0 - Do not follow redirects; 1 - *(default)* Follow redirects. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
headers	array	Array of headers that will be sent when performing an HTTP request. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
http_proxy	string	HTTP(S) proxy connection string. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
ipmi_sensor	string	IPMI sensor. Property behavior: - *required* if `type` is set to ”IPMI agent” and `key_` is not set to ”ipmi.get” - *supported* if `type` is set to ”IPMI agent” - *read-only* for inherited objects
jmx_endpoint	string	JMX agent custom connection string. Default: service:jmx:rmi:///jndi/rmi://{HOST.CONN}:{HOST.PORT}/jmxrmi Property behavior: - *supported* if `type` is set to ”JMX agent”
lifetime	string	Time period after which items that are no longer discovered will be deleted. Accepts a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., `30s`, `1m`, `2h`, `1d`), or a user macro. Default: `7d`.
lifetime_type	integer	Scenario to delete lost LLD resources. Possible values: 0 - *(default)* Delete after lifetime threshold is reached; 1 - Do not delete; 2 - Delete immediately.
enabled_lifetime	string	Time period after which items that are no longer discovered will be disabled. Accepts a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., `30s`, `1m`, `2h`, `1d`), or a user macro. Default: 0.
enabled_lifetime_type	integer	Scenario to disable lost LLD resources. Possible values: 0 - Disable after lifetime threshold is reached; 1 - Do not disable; 2 - *(default)* Disable immediately.
master_itemid	ID	ID of the master item. Discovery rule cannot be master item for another discovery rule. Property behavior: - *required* if `type` is set to ”Dependent item” - *read-only* for inherited objects

Property	Type	Description
output_format	integer	Should the response be converted to JSON. Possible values: 0 - *(default)* Store raw; 1 - Convert to JSON. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
params	string	Additional parameters depending on the type of the LLD rule prototype: - executed script for SSH and Telnet LLD rules; - SQL query for database monitor LLD rules; - formula for calculated LLD rules; - the script for script and browser LLD rules. Property behavior: - *required* if `type` is set to ”Database monitor”, ”SSH agent”, ”TELNET agent”, ”Script” or ”Browser” - *read-only* for inherited objects (if `type` is set to ”Script” or ”Browser”)
parameters	object/array	Additional parameters if `type` is set to ”Script” or ”Browser”. Array of objects with `name` and `value` properties, where `name` must be unique. Property behavior: - *supported* if `type` is set to ”Script” or ”Browser” - *read-only* for inherited objects
password	string	Password for authentication. Property behavior: - *required* if `type` is set to ”JMX agent” and `username` is set - *supported* if `type` is set to ”Simple check”, ”Database monitor”, ”SSH agent”, ”TELNET agent”, or ”HTTP agent” - *read-only* for inherited objects (if `type` is set to ”HTTP agent”)
post_type	integer	Type of post data body stored in `posts` property. Possible values: 0 - *(default)* Raw data; 2 - JSON data; 3 - XML data. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
posts	string	HTTP(S) request body data. Property behavior: - *required* if `type` is set to ”HTTP agent” and `post_type` is set to ”JSON data” or ”XML data” - *supported* if `type` is set to ”HTTP agent” and `post_type` is set to ”Raw data” - *read-only* for inherited objects
privatekey	string	Name of the private key file. Property behavior: - *required* if `type` is set to ”SSH agent” and `authtype` is set to ”public key”
publickey	string	Name of the public key file. Property behavior: - *required* if `type` is set to ”SSH agent” and `authtype` is set to ”public key”

Property	Type	Description
query_fields	array	Array of query fields that will be sent when performing an HTTP request. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
request_method	integer	Type of request method. Possible values: 0 - *(default)* GET; 1 - POST; 2 - PUT; 3 - HEAD. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
retrieve_mode	integer	What part of response should be stored. Possible values if `request_method` is set to ”GET”, ”POST”, or ”PUT”: 0 - *(default)* Body; 1 - Headers; 2 - Both body and headers will be stored. Possible values if `request_method` is set to ”HEAD”: 1 - Headers. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
snmp_oid	string	SNMP OID. Property behavior: - *required* if `type` is set to ”SNMP agent” - *read-only* for inherited objects
ssl_cert_file	string	Public SSL Key file path. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
ssl_key_file	string	Private SSL Key file path. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
ssl_key_password	string	Password for SSL Key file. Property behavior: - *supported* if `type` is set to ”HTTP agent” - *read-only* for inherited objects
state	integer	State of the LLD rule prototype. Possible values: 0 - *(default)* normal; 1 - not supported. Property behavior: - *read-only*

Property	Type	Description
status	integer	Status of the LLD rule prototype. Possible values: 0 - *(default)* enabled LLD rule prototype; 1 - disabled LLD rule prototype.
status_codes	string	Ranges of required HTTP status codes, separated by commas. Also supports user macros as part of comma separated list. Example: 200,200-{$M},{$M},200-400 Property behavior: - *supported* if `type` is set to "HTTP agent" - *read-only* for inherited objects
templateid	ID	ID of the parent template LLD rule prototype. Property behavior: - *read-only*
timeout	string	Item data polling request timeout. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m). Also accepts user macros. Possible values range: 1-600s. Default: "" - use proxy/global settings. Property behavior: - *supported* if `type` is set to "Zabbix agent" (0), "Simple check" (3) and `key_` does not start with "`vmware.`" and "`icmpping`", "Zabbix agent (active)" (7), "External check" (10), "Database monitor" (11), "SSH agent" (13), "TELNET agent" (14), "HTTP agent" (19), "SNMP agent" (20) and `snmp_oid` starts with "`walk[`" or "`get[`", "Script" (21), "Browser" (22) - *read-only* for inherited objects
trapper_hosts	string	Allowed hosts. Property behavior: - *supported* if `type` is set to "Zabbix trapper", or if `type` is set to "HTTP agent" and `allow_traps` is set to "Allow incoming data". If the property is omitted, it is automatically set to the `{$TRAPPER.ALLOWED_HOSTS}` user macro (default value). By default, `{$TRAPPER.ALLOWED_HOSTS}` is `127.0.0.1,::1`. If the property is set to an empty string, incoming connections will be denied from all hosts.
username	string	Username for authentication. Property behavior: - *required* if `type` is set to "SSH agent", "TELNET agent", or if `type` is set to "JMX agent" and `password` is set - *supported* if `type` is set to "Simple check", "Database monitor", or "HTTP agent" - *read-only* for inherited objects (if `type` is set to "HTTP agent")
uuid	string	Universal unique identifier, used for linking imported LLD rule prototypes to already existing ones. Auto-generated, if not given. Property behavior: - *supported* if the LLD rule prototype belongs to a template

Property	Type	Description
verify_host	integer	Whether to validate that the host name for the connection matches the one in the host's certificate. Possible values: 0 - *(default)* Do not validate; 1 - Validate. Property behavior: - *supported* if `type` is set to "HTTP agent" - *read-only* for inherited objects
verify_peer	integer	Whether to validate that the host's certificate is authentic. Possible values: 0 - *(default)* Do not validate; 1 - Validate. Property behavior: - *supported* if `type` is set to "HTTP agent" - *read-only* for inherited objects

HTTP header

The header object has the following properties:

Property	Type	Description
name	string	HTTP header name. Property behavior: - *required*
value	string	Header value. Property behavior: - *required*

HTTP query field

The query field object defines a name and value that is used to specify a URL parameter. It has the following properties:

Property	Type	Description
name	string	Name of the parameter. Property behavior: - *required*
value	string	Parameter value. Property behavior: - *required*

LLD rule prototype filter

The LLD rule prototype filter object defines a set of conditions that can be used to filter discovered objects. It has the following properties:

Property	Type	Description
conditions	object/array	Set of filter conditions to use for filtering results. The conditions will be sorted in the order of their placement in the formula. Property behavior: - *required*

Property	Type	Description
evaltype	integer	Filter condition evaluation method. Possible values: 0 - And/Or; 1 - And; 2 - Or; 3 - Custom expression. Property behavior: - *required*
eval_formula	string	Generated expression that will be used for evaluating filter conditions. The expression contains IDs that reference specific filter conditions by its `formulaid`. The value of `eval_formula` is equal to the value of `formula` for filters with a custom expression. Property behavior: - *read-only*
formula	string	User-defined expression to be used for evaluating conditions of filters with a custom expression. The expression must contain IDs that reference specific filter conditions by its `formulaid`. The IDs used in the expression must exactly match the ones defined in the filter conditions: no condition can remain unused or omitted. Property behavior: - *required* if `evaltype` is set to ”custom expression”

LLD rule prototype filter condition

The LLD rule prototype filter condition object defines a separate check to perform on the value of an LLD macro. It has the following properties:

Property	Type	Description
macro	string	LLD macro to perform the check on. Property behavior: - *required*
value	string	Value to compare with. Property behavior: - *required* if `operator` is set to ”matches regular expression” or ”does not match regular expression”
formulaid	string	Arbitrary unique ID that is used to reference the condition from a custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward. Property behavior: - *required* if `evaltype` of LLD rule prototype filter object is set to ”custom expression”
operator	integer	Condition operator. Possible values: 8 - *(default)* matches regular expression; 9 - does not match regular expression; 12 - exists; 13 - does not exist.

Note:
To better understand how to use filters with various types of expressions, see examples on the discoveryruleprototype.get and discoveryruleprototype.create method pages.

LLD macro path

The LLD macro path has the following properties:

Property	Type	Description
lld_macro	string	LLD macro. Property behavior: - *required*
path	string	Selector for value which will be assigned to corresponding macro. Property behavior: - *required*

LLD rule prototype preprocessing

The LLD rule prototype preprocessing object has the following properties.

Property	Type	Description
type	integer	The preprocessing option type. Possible values: 5 - Regular expression; 11 - XML XPath; 12 - JSONPath; 14 - Matches regular expression; 15 - Does not match regular expression; 16 - Check for error in JSON; 17 - Check for error in XML; 20 - Discard unchanged with heartbeat; 21 - JavaScript; 23 - Prometheus to JSON; 24 - CSV to JSON; 25 - Replace; 27 - XML to JSON; 28 - SNMP walk value; 29 - SNMP walk to JSON; 30 - SNMP get value. Property behavior: - *required*
params	string	Additional parameters used by preprocessing option. Multiple parameters are separated by the newline (\n) character. Property behavior: - *required* if `type` is set to ”Regular expression” (5), ”XML XPath” (11), ”JSONPath” (12), ”Matches regular expression” (14), ”Does not match regular expression” (15), ”Check for error in JSON” (16), ”Check for error in XML” (17), ”Discard unchanged with heartbeat” (20), ”JavaScript” (21), ”Prometheus to JSON” (23), ”CSV to JSON” (24), ”Replace” (25), ”SNMP walk value” (28), ”SNMP walk to JSON” (29), or ”SNMP get value” (30)

Property	Type	Description
error_handler	integer	Action type used in case of preprocessing step failure. Possible values: 0 - Error message is set by Zabbix server; 1 - Discard value; 2 - Set custom value; 3 - Set custom error message. Property behavior: - *required* if `type` is set to ”Regular expression” (5), ”XML XPath” (11), ”JSONPath” (12), ”Matches regular expression” (14), ”Does not match regular expression” (15), ”Check for error in JSON” (16), ”Check for error in XML” (17), ”Prometheus to JSON” (23), ”CSV to JSON” (24), ”XML to JSON” (27), ”SNMP walk value” (28), ”SNMP walk to JSON” (29), or ”SNMP get value” (30)
error_handler_params	string	Error handler parameters. Property behavior: - *required* if `error_handler` is set to ”Set custom value” or ”Set custom error message”

The following parameters and error handlers are supported for each preprocessing type.

Preprocessing type	Name	Parameter 1	Parameter 2	Parameter 3	Supported error handlers
5	Regular expression	pattern[1]	output[2]		0, 1, 2, 3
11	XML XPath	path[3]			0, 1, 2, 3
12	JSONPath	path[3]			0, 1, 2, 3
14	Matches regular expression	pattern[1]			0, 1, 2, 3
15	Does not match regular expression	pattern[1]			0, 1, 2, 3
16	Check for error in JSON	path[3]			0, 1, 2, 3
17	Check for error in XML	path[3]			0, 1, 2, 3

Preprocessing type	Name	Parameter 1	Parameter 2	Parameter 3	Supported error handlers
20	Discard unchanged with heartbeat	seconds^4, 5^			
21	JavaScript	script[2]			
23	Prometheus to JSON	pattern^5, 6^			0, 1, 2, 3
24	CSV to JSON	character[2]	character[2]	0,1	0, 1, 2, 3
25	Replace	search string[2]	replacement[2]		
27	XML to JSON				0, 1, 2, 3
28	SNMP walk value	OID[2]	Format: 0 - Unchanged 1 - UTF-8 from Hex-STRING 2 - MAC from Hex-STRING 3 - Integer from BITS		0, 1, 2, 3
29	SNMP walk to JSON[7]	Field name[2]	OID prefix[2]	Format: 0 - Unchanged 1 - UTF-8 from Hex-STRING 2 - MAC from Hex-STRING 3 - Integer from BITS	0, 1, 2, 3
30	SNMP get value	Format: 1 - UTF-8 from Hex-STRING 2 - MAC from Hex-STRING 3 - Integer from BITS			0, 1, 2, 3

[1] regular expression
 [2] string
 [3] JSONPath or XML XPath
 [4] positive integer (with support of time suffixes, e.g., 30s, 1m, 2h, 1d)
 [5] user macro
 [6] Prometheus pattern following the syntax: `<metric name>{<label name>="<label value>", ...} == <value>`. Each Prometheus pattern component (metric, label name, label value and metric value) can be user macro.
 [7] Supports multiple "Field name,OID prefix,Format records" records delimited by a new line character.

LLD rule prototype overrides

The LLD rule prototype overrides object defines a set of rules (filters, conditions and operations) that are used to override properties of different prototype objects. It has the following properties:

Property	Type	Description
name	string	Unique override name. Property behavior: - *required*
step	integer	Unique order number of the override. Property behavior: - *required*

Property	Type	Description
stop	integer	Stop processing next overrides if matches. Possible values: 0 - *(default)* don't stop processing overrides; 1 - stop processing overrides if filter matches.
filter	object	Override filter.
operations	object/array	Override operations.

LLD rule prototype override filter

The LLD rule prototype override filter object defines a set of conditions that if they match the discovered object the override is applied. It has the following properties:

Property	Type	Description
conditions	object/array	Set of override filter conditions to use for matching the discovered objects. The conditions will be sorted in the order of their placement in the formula. Property behavior: - *required*
evaltype	integer	Override filter condition evaluation method. Possible values: 0 - And/Or; 1 - And; 2 - Or; 3 - Custom expression. Property behavior: - *required*
eval_formula	string	Generated expression that will be used for evaluating override filter conditions. The expression contains IDs that reference specific override filter conditions by its `formulaid`. The value of `eval_formula` is equal to the value of `formula` for filters with a custom expression. Property behavior: - *read-only*
formula	string	User-defined expression to be used for evaluating conditions of override filters with a custom expression. The expression must contain IDs that reference specific override filter conditions by its `formulaid`. The IDs used in the expression must exactly match the ones defined in the override filter conditions: no condition can remain unused or omitted. Property behavior: - *required* if `evaltype` is set to "custom expression"

LLD rule prototype override filter condition

The LLD rule prototype override filter condition object defines a separate check to perform on the value of an LLD macro. It has the following properties:

Property	Type	Description
macro	string	LLD macro to perform the check on. Property behavior: - *required*

Property	Type	Description
value	string	Value to compare with. Property behavior: - *required* if `operator` is set to ”matches regular expression” or ”does not match regular expression”
formulaid	string	Arbitrary unique ID that is used to reference the condition from a custom expression. Can only contain capital-case letters. The ID must be defined by the user when modifying filter conditions, but will be generated anew when requesting them afterward. Property behavior: - *required* if `evaltype` of LLD rule prototype override filter object is set to ”custom expression”
operator	integer	Condition operator. Possible values: 8 - *(default)* matches regular expression; 9 - does not match regular expression; 12 - exists; 13 - does not exist.

LLD rule prototype override operation

The LLD rule prototype override operation is combination of conditions and actions to perform on the prototype object. It has the following properties:

Property	Type	Description
operationobject	integer	Type of discovered object to perform the action. Possible values: 0 - Item prototype; 1 - Trigger prototype; 2 - Graph prototype; 3 - Host prototype. Property behavior: - *required*
operator	integer	Override condition operator. Possible values: 0 - *(default)* equals; 1 - does not equal; 2 - contains; 3 - does not contain; 8 - matches; 9 - does not match.
value	string	Pattern to match item, trigger, graph or host prototype name depending on selected object.
opstatus	object	Override operation status object for item, trigger and host prototype objects.
opdiscover	object	Override operation discover status object (all object types).
opperiod	object	Override operation period (update interval) object for item prototype object.
ophistory	object	Override operation history object for item prototype object.
optrends	object	Override operation trends object for item prototype object.
opseverity	object	Override operation severity object for trigger prototype object.
optag	object/array	Override operation tag object for trigger and host prototype objects.
optemplate	object/array	Override operation template object for host prototype object.
opinventory	object	Override operation inventory object for host prototype object.

LLD rule prototype override operation status

LLD rule prototype override operation status that is set to discovered object. It has the following properties:

Property	Type	Description
status	integer	Override the status for selected object. Possible values: 0 - Create enabled; 1 - Create disabled. Property behavior: - *required*

LLD rule prototype override operation discover

LLD rule prototype override operation discover status that is set to discovered object. It has the following properties:

Property	Type	Description
discover	integer	Override the discover status for selected object. Possible values: 0 - Yes, continue discovering the objects; 1 - No, new objects will not be discovered and existing ones will be marked as lost. Property behavior: - *required*

LLD rule prototype override operation period

LLD rule prototype override operation period is an update interval value that is set to discovered item. It has the following properties:

Property	Type	Description
delay	string	Override the update interval of the item prototype. Accepts a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., 30s, 1m, 2h, 1d), or, optionally, one or more custom intervals, all separated by semicolons. Custom intervals can be a mix of flexible and scheduling intervals. Accepts user macros or LLD macros. If used, the value must be a single macro. Multiple macros or macros mixed with text are not supported. Flexible intervals may be written as two macros separated by a forward slash (e.g., {$FLEX_INTERVAL}/{$FLEX_PERIOD}). Example: 1h;wd1-5h9-18;{$Macro1}/1-7,00:00-24:00;0/6-7,12:00-24:00;{$Macr Property behavior: - *required*

LLD rule prototype override operation history

LLD rule prototype override operation history value that is set to discovered item. It has the following properties:

Property	Type	Description
history	string	Override the history of item prototype which is a time unit of how long the history data should be stored. Also accepts user macro and LLD macro. Property behavior: - *required*

LLD rule prototype override operation trends

LLD rule prototype override operation trends value that is set to discovered item. It has the following properties:

Property	Type	Description
trends	string	Override the trends of item prototype which is a time unit of how long the trends data should be stored. Also accepts user macro and LLD macro. Property behavior: - *required*

LLD rule prototype override operation severity

LLD rule prototype override operation severity value that is set to discovered trigger. It has the following properties:

Property	Type	Description
severity	integer	Override the severity of trigger prototype. Possible values: 0 - *(default)* not classified; 1 - information; 2 - warning; 3 - average; 4 - high; 5 - disaster. Property behavior: - *required*

LLD rule prototype override operation tag

LLD rule prototype override operation tag object contains tag name and value that are set to discovered object. It has the following properties:

Property	Type	Description
tag	string	New tag name. Property behavior: - *required*
value	string	New tag value.

LLD rule prototype override operation template

LLD rule prototype override operation template object that is linked to discovered host. It has the following properties:

Property	Type	Description
templateid	ID	Override the template of host prototype linked templates. Property behavior: - *required*

LLD rule prototype override operation inventory

LLD rule prototype override operation inventory mode value that is set to discovered host. It has the following properties:

Property	Type	Description
inventory_mode	integer	Override the host prototype inventory mode. Possible values: -1 - disabled; 0 - *(default)* manual; 1 - automatic. Property behavior: - *required*

discoveryruleprototype.create

Description

`object discoveryruleprototype.create(object/array lldRules)`

This method allows you to create new LLD rule prototypes.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` LLD rule prototypes to create.

Additionally to the standard LLD rule prototype properties, the method accepts the following parameters.

Parameter	Type	Description
filter	object	LLD rule prototype filter for the LLD rule.
preprocessing	object/array	LLD rule prototype preprocessing options.
lld_macro_paths	object/array	LLD rule prototype lld_macro_path options.
overrides	object/array	LLD rule prototype overrides options.

Return values

`(object)` Returns an object containing the IDs of the created LLD rule prototypes under the `itemids` property. The order of the returned IDs matches the order of the passed LLD rule prototypes.

Examples

Creating an LLD rule prototype

Create an LLD rule prototype (Type: Nested) to discover tablespaces in database instance.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryruleprototype.create",
    "params": {
        "name": "Discover tablespaces for {#DB}",
        "key_": "db.tablespace.discovery[{#DB}]",
        "hostid": "10084",
        "ruleid": "47251",
        "type": 23
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "47252"
        ]
    },
    "id": 1
}
```

See also

- LLD rule prototype filter
- LLD macro paths
- LLD rule prototype preprocessing

Source

CDiscoveryRulePrototype::create() in *ui/include/classes/api/services/CDiscoveryRulePrototype.php*.

discoveryruleprototype.delete

Description

`object discoveryruleprototype.delete(array lldRuleIds)`

This method allows you to delete LLD rule prototypes.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the LLD rule prototypes to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted LLD rule prototypes under the `ruleids` property.

Examples

Deleting multiple LLD rule prototypes

Delete two LLD rule prototypes.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryruleprototype.delete",
    "params": [
        "47252",
        "47253"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "ruleids": [
            "47252",
            "47253"
        ]
    },
```

```
    "id": 1
}
```

Source

CDiscoveryRulePrototype::delete() in *ui/include/classes/api/services/CDiscoveryRulePrototype.php*.

discoveryruleprototype.get

Description

`integer/array discoveryruleprototype.get(object parameters)`

The method allows you to retrieve LLD rule prototypes according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
discoveryids	ID/array	Return only LLD rule prototypes that belong to the given LLD rules or LLD rule prototypes.
itemids	ID/array	Return only LLD rule prototypes with the given IDs.
groupids	ID/array	Return only LLD rule prototypes that belong to the hosts from the given groups.
hostids	ID/array	Return only LLD rule prototypes that belong to the given hosts.
inherited	boolean	If set to `true` return only LLD rule prototypes inherited from a template.
interfaceids	ID/array	Return only LLD rule prototypes use the given host interfaces.
monitored	boolean	If set to `true` return only enabled LLD rule prototypes that belong to monitored hosts.
templated	boolean	If set to `true` return only LLD rule prototypes that belong to templates.
templateids	ID/array	Return only LLD rule prototypes that belong to the given templates.
selectDiscoveryRule	query	Return a `discoveryRule` property with the parent LLD rule of the LLD rule prototype. It will return an empty response if the parent is a discovery rule prototype.
selectDiscoveryRulePrototype	query	Return a `discoveryRulePrototype` property with the parent LLD rule prototype of the LLD rule prototype. It will return an empty response if the parent is a discovery rule (not prototype).
selectDiscoveryRulePrototypes	query	Return a `discoveryRulePrototypes` property with child LLD rule prototypes that belong to the LLD rule prototype. Supports `count`.
selectFilter	query	Return a `filter` property with data of the filter used by the LLD rule prototype.
selectGraphs	query	Returns a `graphs` property with graph prototypes that belong to the LLD rule prototype. Supports `count`.
selectHostPrototypes	query	Return a `hostPrototypes` property with host prototypes that belong to the LLD rule prototype. Supports `count`.
selectHosts	query	Return a `hosts` property with an array of hosts that the LLD rule prototype belongs to.

Parameter	Type	Description
selectItems	query	Return an `items` property with item prototypes that belong to the LLD rule prototype. Supports `count`.
selectTriggers	query	Return a `triggers` property with trigger prototypes that belong to the LLD rule prototype. Supports `count`.
selectLLDMacroPaths	query	Return an `lld_macro_paths` property with a list of LLD macros and paths to values assigned to each corresponding macro.
selectPreprocessing	query	Return a `preprocessing` property with LLD rule prototype preprocessing options.
selectOverrides	query	Return an `lld_rule_overrides` property with a list of override filters, conditions and operations that are performed on prototype objects.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type. Supports additional properties: `host` - technical name of the host that the LLD rule prototype belongs to.
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectItems`, `selectGraphs`, `selectTriggers`.
sortfield	string/array	Sort the result by the given properties. Possible values: `itemid`, `name`, `key_`, `delay`, `type`, `status`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving discovery rule prototypes from a host

Retrieve all discovery rule prototypes for specific host ID.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryruleprototype.get",
    "params": {
        "hostids": "10084"
```

```
    },
    "id": 1
}
```

Response:


```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "itemid": "47253",
            "type": "23",
            "snmp_oid": "",
            "hostid": "10084",
            "name": "Discover tablespaces for {#DB}",
            "key_": "db.tablespace.discovery[{#DB}]",
            "delay": "0",
            "history": "31d",
            "trends": "365d",
            "status": "0",
            "value_type": "4",
            "trapper_hosts": "",
            "units": "",
            "logtimefmt": "",
            "templateid": "0",
            "valuemapid": "0",
            "params": "",
            "ipmi_sensor": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "flags": "3",
            "interfaceid": "0",
            "description": "",
            "inventory_link": "0",
            "lifetime": "7d",
            "jmx_endpoint": "",
            "master_itemid": "0",
            "timeout": "",
            "url": "",
            "query_fields": [],
            "posts": "",
            "status_codes": "200",
            "follow_redirects": "1",
            "post_type": "0",
            "http_proxy": "",
            "headers": [],
            "retrieve_mode": "0",
            "request_method": "0",
            "output_format": "0",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
            "verify_host": "0",
            "allow_traps": "0",
            "discover": "0",
            "uuid": "",
            "lifetime_type": "0",
            "enabled_lifetime_type": "2",
            "enabled_lifetime": "0",
```

```
            "parameters": []
        }
    ],
    "id": 1
}
```

See also

- Graph prototype
- Host
- Item prototype
- LLD rule prototype filter
- Trigger prototype

Source

CDiscoveryRulePrototype::get() in *ui/include/classes/api/services/CDiscoveryRulePrototype.php*.

discoveryruleprototype.update

Description

`object discoveryruleprototype.update(object/array lldRules)`

This method allows you to update existing LLD rule prototypes.

Note that the updating of already discovered prototypes is limited.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` LLD rule prototype properties to be updated.

The `itemid` property must be defined for each LLD rule prototype, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard LLD rule prototype properties, the method accepts the following parameters.

Parameter	Type	Description
filter	object	LLD rule prototype filter to replace the existing filter.
preprocessing	object/array	LLD rule prototype preprocessing options to replace the existing preprocessing options. Parameter behavior: - *read-only* for inherited objects
lld_macro_paths	object/array	LLD rule prototype lld_macro_path options to replace the existing lld_macro_path options. Parameter behavior: - *read-only* for inherited objects
overrides	object/array	LLD rule prototype overrides options to replace the existing overrides options. Parameter behavior: - *read-only* for inherited objects

Return values

`(object)` Returns an object containing the IDs of the updated LLD rule prototypes under the `itemids` property.

Examples

Updating LLD rule prototype preprocessing options

Update an LLD rule prototype with a JSONPath preprocessing rule. Request:

```
{
    "jsonrpc": "2.0",
    "method": "discoveryruleprototype.update",
    "params": {
        "itemid": "47253",
        "preprocessing": [
            {
                "type": 12,
                "params": "$.tablespaces",
                "error_handler": 1
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "itemids": [
            "47253"
        ]
    },
    "id": 1
}
```

Source

CDiscoveryRulePrototype::update() in *ui/include/classes/api/services/CDiscoveryRulePrototype.php*.

Maintenance

This class is designed to work with maintenances.

Object references:

- Maintenance
- Time period
- Problem tag

Available methods:

- maintenance.create - create new maintenances
- maintenance.delete - delete maintenances
- maintenance.get - retrieve maintenances
- maintenance.update - update maintenances

Maintenance object

The following objects are directly related to the `maintenance` API.

Maintenance

The maintenance object has the following properties.

Property	Type	Description
maintenanceid	ID	ID of the maintenance. Property behavior: - *read-only* - *required* for update operations

Property	Type	Description
name	string	Name of the maintenance. Property behavior: - *required* for create operations
active_since	timestamp	Time when the maintenance becomes active (inclusive). The given value will be rounded down to minutes. Property behavior: - *required* for create operations
active_till	timestamp	Time when the maintenance stops being active (exclusive). The given value will be rounded down to minutes. Property behavior: - *required* for create operations
description	string	Description of the maintenance.
maintenance_type	integer	Type of maintenance. Possible values: 0 - *(default)* with data collection; 1 - without data collection.
tags_evaltype	integer	Problem tag evaluation method. Possible values: 0 - *(default)* And/Or; 2 - Or.

Time period

The time period object is used to define periods when the maintenance must come into effect. It has the following properties.

Property	Type	Description
period	integer	Duration of the maintenance period in seconds. The given value will be rounded down to minutes. Possible values range: 300 - 86399940. Default: 3600.
timeperiod_type	integer	Type of time period. Possible values: 0 - *(default)* one time only; 2 - daily; 3 - weekly; 4 - monthly.
start_date	timestamp	Date when the maintenance period must come into effect. The given value will be rounded down to minutes. Default: current date. Property behavior: - *supported* if `timeperiod_type` is set to ”one time only”

Property	Type	Description
start_time	integer	Time of day when the maintenance starts in seconds. The given value will be rounded down to minutes. Default: 0. Property behavior: - *supported* if `timeperiod_type` is set to ”daily”, ”weekly”, or ”monthly”
every	integer	For daily and weekly periods `every` defines the day or week intervals at which the maintenance must come into effect. Default value if `timeperiod_type` is set to ”daily” or ”weekly”: 1. For monthly periods when `day` is set, the `every` property defines the day of the month when the maintenance must come into effect. Default value if `timeperiod_type` is set to ”monthly” and `day` is set: 1. For monthly periods when `dayofweek` is set, the `every` property defines the week of the month when the maintenance must come into effect. Possible values if `timeperiod_type` is set to ”monthly” and `dayofweek` is set: 1 - *(default)* first week; 2 - second week; 3 - third week; 4 - fourth week; 5 - last week. Property behavior: - *supported* if `timeperiod_type` is set to ”daily”, ”weekly”, or ”monthly”
dayofweek	integer	Days of the week when the maintenance must come into effect. Possible bitmap values: 1 - Monday; 2 - Tuesday; 4 - Wednesday; 8 - Thursday; 16 - Friday; 32 - Saturday; 64 - Sunday. This is a bitmask field; any sum of possible bitmap values is acceptable (for example, 21 for Monday, Wednesday, and Friday). Property behavior: - *required* if `timeperiod_type` is set to ”weekly” or if `timeperiod_type` is set to ”monthly” and `day` is not set
day	integer	Day of the month when the maintenance must come into effect. Property behavior: - *required* if `timeperiod_type` is set to ”monthly” and `dayofweek` is not set

Property	Type	Description
month	integer	Months when the maintenance must come into effect. Possible bitmap values: 1 - January; 2 - February; 4 - March; 8 - April; 16 - May; 32 - June; 64 - July; 128 - August; 256 - September; 512 - October; 1024 - November; 2048 - December. This is a bitmask field; any sum of possible bitmap values is acceptable (for example, 585 for January, April, July, and October). Property behavior: - *required* if `timeperiod_type` is set to ”monthly”

Event name

The event name object is used to define which problems must be suppressed when the maintenance comes into effect. Event names can only be specified if `maintenance_type` of Maintenance object is set to ”with data collection”. It has the following properties.

Property	Type	Description
operator	integer	Condition operator. Possible values: 2 - *(default)* Contains; 3 - Does not contain.
value	string	Event name value. Property behavior: - *required*

Problem tag

The problem tag object is used to define which problems must be suppressed when the maintenance comes into effect. Tags can only be specified if `maintenance_type` of Maintenance object is set to ”with data collection”. It has the following properties.

Property	Type	Description
tag	string	Problem tag name. Property behavior: - *required*
operator	integer	Condition operator. Possible values: 0 - Equals; 1 - Does not equal; 2 - *(default)* Contains; 3 - Does not contain.
value	string	Problem tag value.

maintenance.create

Description

`object maintenance.create(object/array maintenances)`

This method allows you to create new maintenances.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Maintenances to create.

Additionally to the standard maintenance properties, the method accepts the following parameters.

Parameter	Type	Description
groups	object/array	Host groups to be put into maintenance. The host groups must have only the `groupid` property defined. Parameter behavior: - *required* if `hosts` is not set
hosts	object/array	Hosts to be put into maintenance. The hosts must have only the `hostid` property defined. Parameter behavior: - *required* if `groups` is not set
triggers	object/array	Triggers to be put into maintenance. The triggers must have only the `triggerid` property defined. Parameter behavior: - *required* if `groups` and `hosts` are not set - *supported* if maintenance_type=0 (with data collection)
event_names	object/array	Event names to define problems to be put into by maintenance. Parameter behavior: - *supported* if maintenance_type=0 (with data collection)
tags	object/array	Problem tags to define problems to be put into maintenance. Parameter behavior: - *supported* if `maintenance_type` of Maintenance object is set to "with data collection"
timeperiods	object/array	Maintenance time periods. Parameter behavior: - *required*

Return values

`(object)` Returns an object containing the IDs of the created maintenances under the `maintenanceids` property. The order of the returned IDs matches the order of the passed maintenances.

Examples

Creating a maintenance

Create a maintenance with data collection for host group with ID "20", and with problem tags **service:mysqld** and **error**. It must be active from 17.03.2026 till 17.03.2027, come in effect each Sunday at 18:00 and last for one hour.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "maintenance.create",
    "params": {
        "name": "Sunday maintenance",
        "active_since": 1773720240,
        "active_till": 1805256240,
        "tags_evaltype": 0,
        "groups": [
            {"groupid": "20"}
        ],
        "timeperiods": [
            {
                "period": 3600,
                "timeperiod_type": 3,
                "start_time": 64800,
                "every": 1,
                "dayofweek": 64
            }
        ],
        "tags": [
            {
                "tag": "service",
                "operator": "0",
                "value": "mysqld"
            },
            {
                "tag": "error",
                "operator": "2",
                "value": ""
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "maintenanceids": [
            "3"
        ]
    },
    "id": 1
}
```

Creating ad-hoc maintenance

Create a one hour ad-hoc maintenance with data collection for host group with ID ”4”, trigger with ID ”12345”, and event name **Printer: ink low**.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "maintenance.create",
    "params": {
        "name": "Ad-hoc maintenance",
        "active_since": 1792057740,
        "active_till": 1792061340,
        "tags_evaltype": 0,
        "groups": [
            {"groupid": "4"}
```

```
        ],
        "triggers": [
            {"triggerid": "12345"}
        ],
        "timeperiods": [
            {
                "period": 3600,
                "timeperiod_type": 0,
                "start_date": 1792057740
            }
        ],
        "event_names": [
            {
                "operator": "2",
                "value": "Printer: ink low"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "maintenanceids": [
            "4"
        ]
    },
    "id": 1
}
```

See also

- Host
- Host group
- Trigger
- Event name
- Problem tag
- Time period

Source

CMaintenance::create() in *ui/include/classes/api/services/CMaintenance.php*.

maintenance.delete

Description

`object maintenance.delete(array maintenanceIds)`

This method allows you to delete maintenance periods.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the maintenance periods to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted maintenance periods under the `maintenanceids` property.

Examples

Deleting multiple maintenance periods

Delete two maintenance periods.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "maintenance.delete",
    "params": [
        "3",
        "1"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "maintenanceids": [
            "3",
            "1"
        ]
    },
    "id": 1
}
```

Source

CMaintenance::delete() in *ui/include/classes/api/services/CMaintenance.php*.

maintenance.get

Description

`integer/array maintenance.get(object parameters)`

The method allows you to retrieve maintenances according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
groupids	ID/array	Return only maintenances that are assigned to the given host groups.
hostids	ID/array	Return only maintenances that are assigned to the given hosts.
triggerids	ID/array	Return only maintenances that are assigned to the given triggers.
maintenanceids	ID/array	Return only maintenances with the given IDs.
selectHostGroups	query	Return a `hostgroups` property with host groups assigned to the maintenance.
selectHosts	query	Return a `hosts` property with hosts assigned to the maintenance.
selectTriggers	query	Return a `triggers` property with triggers assigned to the maintenance.
selectEventNames	query	Return an `event_names` proper,ty with event names assigned to the maintenance.
selectTags	query	Return a `tags` property with problem tags of the maintenance.
selectTimeperiods	query	Return a `timeperiods` property with time periods of the maintenance.

Parameter	Type	Description
sortfield	string/array	Sort the result by the given properties. Possible values: `maintenanceid`, `name`, `maintenance_type`, `active_since`, `active_till`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving maintenances

Retrieve all configured maintenances, the data about the assigned maintenance scope (host group, host, trigger, event name, tags), and defined time periods.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "maintenance.get",
    "params": {
        "output": "extend",
        "selectHostGroups": "extend",
        "selectHosts": "extend",
        "selectTriggers": "extend",
        "selectEventNames": "extend",
        "selectTags": "extend",
        "selectTimeperiods": "extend"
    },
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "maintenanceid": "1",
      "name": "Regular DB maintenance",
      "maintenance_type": "0",
      "description": "",
      "active_since": "1791936000",
      "active_till": "2107555200",
      "tags_evaltype": "0",
      "hostgroups": [
        {
          "groupid": "20",
          "name": "Databases",
```

```
          "flags": "0",
          "uuid": "748ad4d098d447d492bb935c907f652f"
        }
      ],
      "triggers": [],
      "event_names": [],
      "hosts": [],
      "tags": [],
      "timeperiods": [
        {
          "timeperiod_type": "4",
          "every": "1",
          "month": "4095",
          "dayofweek": "0",
          "day": "1",
          "start_time": "64800",
          "period": "3600",
          "start_date": "0"
        }
      ]
    },
    {
      "maintenanceid": "2",
      "name": "Ad-hoc: Printer: ink OUT",
      "maintenance_type": "0",
      "description": "",
      "active_since": "1792057740",
      "active_till": "1792061340",
      "tags_evaltype": "0",
      "hostgroups": [
        {
          "groupid": "4",
          "name": "Zabbix servers",
          "flags": "0",
          "uuid": "6f6799aa69e844b4b3918f779f2abf08"
        }
      ],
      "triggers": [],
      "event_names": [
        {
          "operator": "2",
          "value": "Printer: ink OUT"
        }
      ],
      "hosts": [],
      "tags": [],
      "timeperiods": [
        {
          "timeperiod_type": "0",
          "every": "1",
          "month": "0",
          "dayofweek": "0",
          "day": "0",
          "start_time": "0",
          "period": "3600",
          "start_date": "1792057740"
        }
      ]
    }
  ],
  "id": 1
}
```

See also

- Host
- Host group
- Trigger
- Event name
- Problem tag
- Time period

Source

CMaintenance::get() in *ui/include/classes/api/services/CMaintenance.php*.

maintenance.update

Description

`object maintenance.update(object/array maintenances)`

This method allows you to update existing maintenances.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Maintenance properties to be updated.

The `maintenanceid` property must be defined for each maintenance, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard maintenance properties, the method accepts the following parameters.

Parameter	Type	Description
groups	object/array	Host groups to replace the current groups. The host groups must have only the `groupid` property defined. Parameter behavior: - *required* if `hosts` is not set
hosts	object/array	Hosts to replace the current hosts. The hosts must have only the `hostid` property defined. Parameter behavior: - *required* if `groups` is not set
triggers	object/array	Triggers to replace the current triggers. The triggers must have only the `triggerid` property defined. Parameter behavior: - *required* if `groups` and `hosts` are not set - *supported* if maintenance_type=0 (with data collection)
event_names	object/array	Event names to replace the current event names. Parameter behavior: - *supported* if maintenance_type=0 (with data collection)
tags	object/array	Problem tags to replace the current tags. Parameter behavior: - *supported* if `maintenance_type` of Maintenance object is set to "with data collection"
timeperiods	object/array	Maintenance time periods to replace the current periods.

Return values

(object) Returns an object containing the IDs of the updated maintenances under the maintenanceids property.

Examples

Assigning different hosts

Replace the hosts currently assigned to maintenance with two different ones.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "maintenance.update",
    "params": {
        "maintenanceid": "3",
        "hosts": [
            {"hostid": "10085"},
            {"hostid": "10084"}
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "maintenanceids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Time period

Source

CMaintenance::update() in *ui/include/classes/api/services/CMaintenance.php*.

Map

This class is designed to work with maps.

Object references:

- Map
- Map element
  - Map element Host
  - Map element Host group
  - Map element Map
  - Map element Trigger
  - Map element Tag
  - Map element URL
- Map link
  - Map link trigger
  - Map link indicators
- Map URL
- Map user
- Map user group
- Map shapes
- Map lines

Available methods:

- map.create - create new maps
- map.delete - delete maps
- map.get - retrieve maps
- map.update - update maps

Map object

The following objects are directly related to the `map` API.

Map

The map object has the following properties.

Property	Type	Description
sysmapid	ID	ID of the map. Property behavior: - *read-only* - *required* for update operations
height	integer	Height of the map in pixels. Property behavior: - *required* for create operations
name	string	Name of the map. Property behavior: - *required* for create operations
width	integer	Width of the map in pixels. Property behavior: - *required* for create operations
backgroundid	ID	ID of the image used as the background for the map.
background_scale	integer	Whether to enable background image proportional scaling. Possible values: 0 - disable scaling; 1 - *(default)* enable image scaling.
expand_macros	integer	Whether to expand macros in labels when configuring the map. Possible values: 0 - *(default)* do not expand macros; 1 - expand macros.
expandproblem	integer	Whether the problem trigger will be displayed for elements with a single problem. Possible values: 0 - always display the number of problems; 1 - *(default)* display the problem trigger if there's only one problem.
grid_align	integer	Whether to enable grid aligning. Possible values: 0 - disable grid aligning; 1 - *(default)* enable grid aligning.
grid_show	integer	Whether to show the grid on the map. Possible values: 0 - do not show the grid; 1 - *(default)* show the grid.
grid_size	integer	Size of the map grid in pixels. Supported values: 20, 40, 50, 75 and 100. Default: 50.

Property	Type	Description
highlight	integer	Whether icon highlighting is enabled. Possible values: 0 - highlighting disabled; 1 - *(default)* highlighting enabled.
iconmapid	ID	ID of the icon map used on the map.
label_format	integer	Whether to enable advanced labels. Possible values: 0 - *(default)* disable advanced labels; 1 - enable advanced labels.
label_location	integer	Location of the map element label. Possible values: 0 - *(default)* bottom; 1 - left; 2 - right; 3 - top.
label_string_host	string	Custom label for host elements. Property behavior: - *required* if `label_type_host` is set to ”custom”
label_string_hostgroup	string	Custom label for host group elements. Property behavior: - *required* if `label_type_hostgroup` is set to ”custom”
label_string_image	string	Custom label for image elements. Property behavior: - *required* if `label_type_image` is set to ”custom”
label_string_map	string	Custom label for map elements. Property behavior: - *required* if `label_type_map` is set to ”custom”
label_string_trigger	string	Custom label for trigger elements. Property behavior: - *required* if `label_type_trigger` is set to ”custom”
label_type	integer	Map element label type. Possible values: 0 - label; 1 - IP address; 2 - *(default)* element name; 3 - status only; 4 - nothing.
label_type_host	integer	Label type for host elements. Possible values: 0 - label; 1 - IP address; 2 - *(default)* element name; 3 - status only; 4 - nothing; 5 - custom.

Property	Type	Description
label_type_hostgroup	integer	Label type for host group elements. Possible values: 0 - label; 2 - *(default)* element name; 3 - status only; 4 - nothing; 5 - custom.
label_type_image	integer	Label type for host group elements. Possible values: 0 - label; 2 - *(default)* element name; 4 - nothing; 5 - custom.
label_type_map	integer	Label type for map elements. Possible values: 0 - label; 2 - *(default)* element name; 3 - status only; 4 - nothing; 5 - custom.
label_type_trigger	integer	Label type for trigger elements. Possible values: 0 - label; 2 - *(default)* element name; 3 - status only; 4 - nothing; 5 - custom.
markelements	integer	Whether to highlight map elements that have recently changed their status. Possible values: 0 - *(default)* do not highlight elements; 1 - highlight elements.
severity_min	integer	Minimum severity of the triggers that will be displayed on the map. Refer to the trigger `severity` property for a list of supported trigger severities.
show_element_label	integer	How to display element labels by default. Possible values: 0 - always display; 1 - *(default)* auto-hide.
show_link_label	integer	How to display link labels by default. Possible values: 0 - always display; 1 - *(default)* auto-hide.
show_unack	integer	How problems should be displayed. Possible values: 0 - *(default)* display the count of all problems; 1 - display only the count of unacknowledged problems; 2 - display the count of acknowledged and unacknowledged problems separately.
userid	ID	ID of the user that is the owner of the map.

Property	Type	Description
private	integer	Type of map sharing. Possible values: 0 - public map; 1 - *(default)* private map.
show_suppressed	integer	Whether suppressed problems are shown. Possible values: 0 - *(default)* hide suppressed problems; 1 - show suppressed problems.

Map element

The map element object defines an object displayed on a map. It has the following properties.

Property	Type	Description
selementid	ID	ID of the map element. Property behavior: - *read-only*
elements	array	Element data object. Property behavior: - *required* if `elementtype` is set to ”host”, ”map”, ”trigger” or ”host group”
elementtype	integer	Type of map element. Possible values: 0 - host; 1 - map; 2 - trigger; 3 - host group; 4 - image. Property behavior: - *required*
iconid_off	ID	ID of the image used to display the element in default state. Property behavior: - *required*
areatype	integer	How separate host group hosts should be displayed. Possible values: 0 - *(default)* the host group element will take up the whole map; 1 - the host group element will have a fixed size.
elementsubtype	integer	How a host group element should be displayed on a map. Possible values: 0 - *(default)* display the host group as a single element; 1 - display each host in the group separately.
evaltype	integer	Map element tag filtering condition evaluation method. Possible values: 0 - *(default)* And/Or; 2 - Or.
height	integer	Height of the fixed size host group element in pixels. Default: 200.

Property	Type	Description
iconid_disabled	ID	ID of the image used to display disabled map elements. Property behavior: - *supported* if `elementtype` is set to ”host”, ”map”, ”trigger”, or ”host group”
iconid_maintenance	ID	ID of the image used to display map elements in maintenance. Property behavior: - *supported* if `elementtype` is set to ”host”, ”map”, ”trigger”, or ”host group”
iconid_on	ID	ID of the image used to display map elements with problems. Property behavior: - *supported* if `elementtype` is set to ”host”, ”map”, ”trigger”, or ”host group”
label	string	Label of the element.
label_location	integer	Location of the map element label. Possible values: -1 - *(default)* default location; 0 - bottom; 1 - left; 2 - right; 3 - top.
permission	integer	Type of permission level. Possible values: -1 - none; 2 - read only; 3 - read-write.
show_label	integer	How to display the element label. Possible values: -1 - *(default)* map default; 0 - always display; 1 - auto-hide.
sysmapid	ID	ID of the map that the element belongs to. Property behavior: - *read-only*
urls	array	Map element URLs. The map element URL object is described in detail below.
use_iconmap	integer	Whether icon mapping must be used for host elements. Possible values: 0 - do not use icon mapping; 1 - *(default)* use icon mapping.
viewtype	integer	Host group element placing algorithm. Possible values: 0 - *(default)* grid.
width	integer	Width of the fixed size host group element in pixels. Default: 200.
x	integer	X-coordinates of the element in pixels. Default: 0.
y	integer	Y-coordinates of the element in pixels. Default: 0.

Property	Type	Description
zindex	integer	Value used to order map elements (z-index). Default: 0.

Map element Host

The map element Host object defines one host element.

Property	Type	Description
hostid	ID	ID of the host.

Map element Host group

The map element Host group object defines one host group element.

Property	Type	Description
groupid	ID	ID of the host group.

Map element Map

The map element Map object defines one map element.

Property	Type	Description
sysmapid	ID	ID of the map.

Map element Trigger

The map element Trigger object defines one or more trigger elements.

Property	Type	Description
triggerid	ID	ID of the trigger.

Map element Tag

The map element Tag object has the following properties.

Property	Type	Description
tag	string	Map element tag name. Property behavior: - *required*
operator	integer	Map element tag condition operator. Possible values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
value	string	Map element tag value.

Map element URL

The map element URL object defines a clickable link that will be available for a specific map element. It has the following properties:

Property	Type	Description
sysmapelementurlid	ID	ID of the map element URL. Property behavior: - *read-only*
name	string	Link caption. Property behavior: - *required*
url	string	Link URL. Property behavior: - *required*
selementid	ID	ID of the map element that the URL belongs to.

Map link

The map link object defines a link between two map elements. It has the following properties.

Property	Type	Description
linkid	ID	ID of the map link. Property behavior: - *read-only*
sysmapid	ID	ID of the map the link belongs to.
selementid1	ID	ID of the first map element linked on one end. Property behavior: - *required*
selementid2	ID	ID of the first map element linked on the other end. Property behavior: - *required*
drawtype	integer	Link line draw style. Possible values: 0 - *(default)* line; 2 - bold line; 3 - dotted line; 4 - dashed line.
color	string	Line color as a hexadecimal color code. Default: 000000.
label	string	Link label.
show_label	integer	How to display the link label. Possible values: -1 - *(default)* map default; 0 - always display; 1 - auto-hide.
indicator_type	integer	Select the link indicator type. Possible values: 0 - *(default)* static link; 1 - trigger; 2 - item value.

Property	Type	Description
linktriggers	array	Map link triggers to use as link status indicators. The map link trigger object is described in detail below. Property behavior: - *required* if `indicator_type` is set to ”trigger” (1).
itemid	ID	ID of the item. Property behavior: - *required* if `indicator_type` is set to ”item value” (2).
highlights	array	Map link highlights to use as link status indicators. The map link indicators object is described in detail below. Property behavior: - *required* if `indicator_type` is set to ”item value” (2).
thresholds	array	Map link thresholds to use as link status indicators. The map link indicators object is described in detail below. Property behavior: - *required* if `indicator_type` is set to ”item value” (2).
permission	integer	Type of permission level. Possible values: -1 - none; 2 - read only; 3 - read-write.

Map link trigger

The map link trigger object defines a map link status indicator based on the state of a trigger. It has the following properties:

Property	Type	Description
triggerid	ID	ID of the trigger used as a link indicator. Property behavior: - *required*
color	string	Indicator color as a hexadecimal color code. Default: `DD0000`.
drawtype	integer	Indicator draw style. Possible values: 0 - *(default)* line; 2 - bold line; 3 - dotted line; 4 - dashed line.

Map link indicators

The map link indicators object defines a map link status indicator based on the item value. It has the following properties:

Property	Type	Description
pattern	string	Regular expression to compare against. Available only for highlights.
sortorder	integer	Used for defining the sort order of highlights. Available only for highlights.
threshold	string	Numeric value to compare against. Available only for thresholds.

Property	Type	Description
drawtype	integer	Indicator draw style. Possible values: 0 - *(default)* line; 2 - bold line; 3 - dotted line; 4 - dashed line.
color	string	Indicator color as a hexadecimal color code. Default: DD0000. Property behavior: - *required*

Map URL

The map URL object defines a clickable link that will be available for all elements of a specific type on the map. It has the following properties:

Property	Type	Description
sysmapurlid	ID	ID of the map URL. Property behavior: - *read-only*
name	string	Link caption. Property behavior: - *required*
url	string	Link URL. Property behavior: - *required*
elementtype	integer	Type of map element for which the URL will be available. Refer to the map element `type` property for a list of supported types. Default: 0.
sysmapid	ID	ID of the map that the URL belongs to.

Map user

List of map permissions based on users. It has the following properties:

Property	Type	Description
sysmapuserid	ID	ID of the map user. Property behavior: - *read-only*
userid	ID	ID of the user. Property behavior: - *required*

Property	Type	Description
permission	integer	Type of permission level. Possible values: 2 - read only; 3 - read-write. Property behavior: - *required*

Map user group

List of map permissions based on user groups. It has the following properties:

Property	Type	Description
sysmapusrgrpid	ID	ID of the map user group. Property behavior: - *read-only*
usrgrpid	ID	ID of the user group. Property behavior: - *required*
permission	integer	Type of permission level. Possible values: 2 - read only; 3 - read-write. Property behavior: - *required*

Map shapes

The map shape object defines a geometric shape (with or without text) displayed on a map. It has the following properties:

Property	Type	Description
sysmap_shapeid	ID	ID of the map shape element. Property behavior: - *read-only*
type	integer	Type of map shape element. Possible values: 0 - rectangle; 1 - ellipse. Property is required when new shapes are created. Property behavior: - *required*
x	integer	X-coordinates of the shape in pixels. Default: 0.
y	integer	Y-coordinates of the shape in pixels. Default: 0.
width	integer	Width of the shape in pixels. Default: 200.

Property	Type	Description
height	integer	Height of the shape in pixels. Default: 200.
text	string	Text of the shape.
font	integer	Font of the text within shape. Possible values: 0 - Georgia, serif 1 - ”Palatino Linotype”, ”Book Antiqua”, Palatino, serif 2 - ”Times New Roman”, Times, serif 3 - Arial, Helvetica, sans-serif 4 - ”Arial Black”, Gadget, sans-serif 5 - ”Comic Sans MS”, cursive, sans-serif 6 - Impact, Charcoal, sans-serif 7 - ”Lucida Sans Unicode”, ”Lucida Grande”, sans-serif 8 - Tahoma, Geneva, sans-serif 9 - ”Trebuchet MS”, Helvetica, sans-serif 10 - Verdana, Geneva, sans-serif 11 - ”Courier New”, Courier, monospace 12 - ”Lucida Console”, Monaco, monospace Default: 9.
font_size	integer	Font size in pixels. Default: 11.
font_color	string	Font color. Default: 000000.
text_halign	integer	Horizontal alignment of text. Possible values: 0 - center; 1 - left; 2 - right. Default: 0.
text_valign	integer	Vertical alignment of text. Possible values: 0 - middle; 1 - top; 2 - bottom. Default: 0.
border_type	integer	Type of the border. Possible values: 0 - none; 1 - ——————; 2 - ····; 3 - – – –. Default: 0.
border_width	integer	Width of the border in pixels. Default: 0.
border_color	string	Border color. Default: 000000.
background_color	string	Background color (fill color). Default: (empty).

Property	Type	Description
zindex	integer	Value used to order all shapes and lines (z-index). Default: 0.

Map lines

The map line object defines a line displayed on a map. It has the following properties:

Property	Type	Description
sysmap_shapeid	ID	ID of the map shape element. Property behavior: - *read-only*
x1	integer	X-coordinates of the line point 1 in pixels. Default: 0.
y1	integer	Y-coordinates of the line point 1 in pixels. Default: 0.
x2	integer	X-coordinates of the line point 2 in pixels. Default: 200.
y2	integer	Y-coordinates of the line point 2 in pixels. Default: 200.
line_type	integer	Type of the lines. Possible values: 0 - none; 1 - ————; 2 - ┅; 3 - − − −. Default: 0.
line_width	integer	Width of the lines in pixels. Default: 0.
line_color	string	Line color. Default: 000000.
zindex	integer	Value used to order all shapes and lines (z-index). Default: 0.

map.create

Description

`object map.create(object/array maps)`

This method allows you to create new maps.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Maps to create.

Additionally to the standard map properties, the method accepts the following parameters.

Parameter	Type	Description
links	array	Map links to be created on the map.
selements	array	Map elements to be created on the map.
urls	array	Map URLs to be created on the map.
users	array	Map user shares to be created on the map.
userGroups	array	Map user group shares to be created on the map.
shapes	array	Map shapes to be created on the map.
lines	array	Map lines to be created on the map.

Note:
To create map links you'll need to set a map element `selementid` to an arbitrary value and then use this value to reference this element in the links `selementid1` or `selementid2` properties. When the element is created, this value will be replaced with the correct ID generated by Zabbix. See example.

Return values

`(object)` Returns an object containing the IDs of the created maps under the `sysmapids` property. The order of the returned IDs matches the order of the passed maps.

Examples

Create an empty map

Create a map with no elements.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "map.create",
    "params": {
        "name": "Map",
        "width": 600,
        "height": 600
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "sysmapids": [
            "8"
        ]
    },
    "id": 1
}
```

Create a host map

Create a map with two host elements and a link between them. Note the use of temporary "selementid1" and "selementid2" values in the map link object to refer to map elements.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "map.create",
    "params": {
        "name": "Host map",
        "width": 600,
        "height": 600,
        "selements": [
            {
```

```
                "selementid": "1",
                "elements": [
                    {"hostid": "1033"}
                ],
                "elementtype": 0,
                "iconid_off": "2"
            },

            {
                "selementid": "2",
                "elements": [
                    {"hostid": "1037"}
                ],
                "elementtype": 0,
                "iconid_off": "2"
            }
        ],
        "links": [
            {
                "selementid1": "1",
                "selementid2": "2"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "sysmapids": [
            "9"
        ]
    },
    "id": 1
}
```

Create a trigger map

Create a map with trigger element, which contains two triggers.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "map.create",
    "params": {
        "name": "Trigger map",
        "width": 600,
        "height": 600,
        "selements": [
            {
                "elements": [
                    {"triggerid": "12345"},
                    {"triggerid": "67890"}
                ],
                "elementtype": 2,
                "iconid_off": "2"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "sysmapids": [
            "10"
        ]
    },
    "id": 1
}
```

Map sharing

Create a map with two types of sharing (user and user group).

Request:

```
{
    "jsonrpc": "2.0",
    "method": "map.create",
    "params": {
        "name": "Map sharing",
        "width": 600,
        "height": 600,
        "users": [
            {
                "userid": "4",
                "permission": "3"
            }
        ],
        "userGroups": [
            {
                "usrgrpid": "7",
                "permission": "2"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "sysmapids": [
            "9"
        ]
    },
    "id": 1
}
```

Map shapes

Create a map with map name title.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "map.create",
    "params": {
        "name": "Host map",
        "width": 600,
        "height": 600,
        "shapes": [
            {
```

```
                "type": 0,
                "x": 0,
                "y": 0,
                "width": 600,
                "height": 11,
                "text": "{MAP.NAME}"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "sysmapids": [
            "10"
        ]
    },
    "id": 1
}
```

Map lines

Create a map line.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "map.create",
    "params": {
        "name": "Map API lines",
        "width": 500,
        "height": 500,
        "lines": [
            {
                "x1": 30,
                "y1": 10,
                "x2": 100,
                "y2": 50,
                "line_type": 1,
                "line_width": 10,
                "line_color": "009900"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "sysmapids": [
            "11"
        ]
    },
    "id": 1
}
```

See also

- Map element

- Map link
- Map URL
- Map user
- Map user group
- Map shape
- Map line

Source

CMap::create() in *ui/include/classes/api/services/CMap.php*.

map.delete

Description

`object map.delete(array mapIds)`

This method allows you to delete maps.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the maps to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted maps under the `sysmapids` property.

Examples

Delete multiple maps

Delete two maps.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "map.delete",
    "params": [
        "12",
        "34"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "sysmapids": [
            "12",
            "34"
        ]
    },
    "id": 1
}
```

Source

CMap::delete() in *ui/include/classes/api/services/CMap.php*.

map.get

Description

`integer/array map.get(object parameters)`

The method allows you to retrieve maps according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
sysmapids	ID/array	Returns only maps with the given IDs.
userids	ID/array	Returns only maps that belong to the given user IDs.
expandUrls	flag	Adds global map URLs to the corresponding map elements and expands macros in all map element URLs.
selectIconMap	query	Returns an `iconmap` property with the icon map used on the map.
selectLinks	query	Returns a `links` property with the map links between elements.
selectSelements	query	Returns a `selements` property with the map elements.
selectUrls	query	Returns a `urls` property with the map URLs.
selectUsers	query	Returns a `users` property with users that the map is shared with.
selectUserGroups	query	Returns a `userGroups` property with user groups that the map is shared with.
selectShapes	query	Returns a `shapes` property with the map shapes.
selectLines	query	Returns a `lines` property with the map lines.
sortfield	string/array	Sort the result by the given properties. Possible values: `name`, `width`, `height`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve a map

Retrieve all data about map "3".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "map.get",
    "params": {
        "output": "extend",
        "selectSelements": "extend",
        "selectLinks": "extend",
```

```
        "selectUsers": "extend",
        "selectUserGroups": "extend",
        "selectShapes": "extend",
        "selectLines": "extend",
        "sysmapids": "3"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "selements": [
                {
                    "selementid": "10",
                    "sysmapid": "3",
                    "elementtype": "4",
                    "evaltype": "0",
                    "iconid_off": "1",
                    "iconid_on": "0",
                    "label": "Zabbix server",
                    "label_location": "3",
                    "x": "11",
                    "y": "141",
                    "iconid_disabled": "0",
                    "iconid_maintenance": "0",
                    "elementsubtype": "0",
                    "areatype": "0",
                    "width": "200",
                    "height": "200",
                    "viewtype": "0",
                    "use_iconmap": "1",
                    "show_label": "-1",
                    "zindex":"0",
                    "elements": [],
                    "urls": [],
                    "tags": [
                        {
                            "tag": "service",
                            "value": "mysqld",
                            "operator": "0"
                        }
                    ]
                },
                {
                    "selementid": "11",
                    "sysmapid": "3",
                    "elementtype": "4",
                    "evaltype": "0",
                    "iconid_off": "1",
                    "iconid_on": "0",
                    "label": "Web server",
                    "label_location": "3",
                    "x": "211",
                    "y": "191",
                    "iconid_disabled": "0",
                    "iconid_maintenance": "0",
                    "elementsubtype": "0",
                    "areatype": "0",
                    "width": "200",
```

```
                "height": "200",
                "viewtype": "0",
                "use_iconmap": "1",
                "show_label": "0",
                "zindex":"0",
                "elements": [],
                "urls": [],
                "tags": []
            },
            {
                "selementid": "12",
                "sysmapid": "3",
                "elementtype": "0",
                "evaltype": "0",
                "iconid_off": "185",
                "iconid_on": "0",
                "label": "{HOST.NAME}\r\n{HOST.CONN}",
                "label_location": "0",
                "x": "111",
                "y": "61",
                "iconid_disabled": "0",
                "iconid_maintenance": "0",
                "elementsubtype": "0",
                "areatype": "0",
                "width": "200",
                "height": "200",
                "viewtype": "0",
                "use_iconmap": "0",
                "show_label": "1",
                "zindex":"0",
                "elements": [
                    {
                        "hostid": "10084"
                    }
                ],
                "urls": [],
                "tags": []
            }
        ],
        "links": [
            {
                "linkid": "23",
                "sysmapid": "3",
                "selementid1": "10",
                "selementid2": "11",
                "drawtype": "0",
                "color": "00CC00",
                "label": "",
                "show_label": "1",
                "indicator_type": "0",
                "itemid": "0",
                "linktriggers": [],
                "thresholds": [],
                "highlights": []
            }
        ],
        "users": [
            {
                "sysmapuserid": "1",
                "userid": "2",
                "permission": "2"
            }
```

```
        ],
        "userGroups": [
            {
                "sysmapusrgrpid": "1",
                "usrgrpid": "7",
                "permission": "2"
            }
        ],
        "shapes":[
            {
                "sysmap_shapeid":"1",
                "type":"0",
                "x":"0",
                "y":"0",
                "width":"680",
                "height":"15",
                "text":"{MAP.NAME}",
                "font":"9",
                "font_size":"11",
                "font_color":"000000",
                "text_halign":"0",
                "text_valign":"0",
                "border_type":"0",
                "border_width":"0",
                "border_color":"000000",
                "background_color":"",
                "zindex":"0"
            }
        ],
        "lines":[
            {
                "sysmap_shapeid":"2",
                "x1": 30,
                "y1": 10,
                "x2": 100,
                "y2": 50,
                "line_type": 1,
                "line_width": 10,
                "line_color": "009900",
                "zindex":"1"
            }
        ],
        "sysmapid": "3",
        "name": "Local network",
        "width": "400",
        "height": "400",
        "backgroundid": "0",
        "background_scale": "1",
        "label_type": "2",
        "label_location": "3",
        "show_element_label": "0",
        "show_link_label": "1",
        "highlight": "1",
        "expandproblem": "1",
        "markelements": "0",
        "show_unack": "0",
        "grid_size": "50",
        "grid_show": "1",
        "grid_align": "1",
        "label_format": "0",
        "label_type_host": "2",
        "label_type_hostgroup": "2",
```

```
            "label_type_trigger": "2",
            "label_type_map": "2",
            "label_type_image": "2",
            "label_string_host": "",
            "label_string_hostgroup": "",
            "label_string_trigger": "",
            "label_string_map": "",
            "label_string_image": "",
            "iconmapid": "0",
            "expand_macros": "0",
            "severity_min": "0",
            "userid": "1",
            "private": "1",
            "show_suppressed": "1"
        }
    ],
    "id": 1
}
```


See also

- Icon map
- Map element
- Map link
- Map URL
- Map user
- Map user group
- Map shapes
- Map lines

Source

CMap::get() in *ui/include/classes/api/services/CMap.php*.

map.update

Description

`object map.update(object/array maps)`

This method allows you to update existing maps.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Map properties to be updated.

The `mapid` property must be defined for each map, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard map properties, the method accepts the following parameters.

Parameter	Type	Description
links	array	Map links to replace the existing links.
selements	array	Map elements to replace the existing elements.
urls	array	Map URLs to replace the existing URLs.
users	array	Map user shares to replace the existing elements.
userGroups	array	Map user group shares to replace the existing elements.
shapes	array	Map shapes to replace the existing shapes.
lines	array	Map lines to replace the existing lines.

Note:
To create map links between new map elements you'll need to set an element's `selementid` to an arbitrary value and then use this value to reference this element in the links `selementid1` or `selementid2` properties. When the element is created, this value will be replaced with the correct ID generated by Zabbix. See example for map.create.

Return values

`(object)` Returns an object containing the IDs of the updated maps under the `sysmapids` property.

Examples

Resize a map

Change the size of the map to 1200x1200 pixels.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "map.update",
    "params": {
        "sysmapid": "8",
        "width": 1200,
        "height": 1200
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "sysmapids": [
            "8"
        ]
    },
    "id": 1
}
```

Change map owner

Available only for admins and super admins.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "map.update",
    "params": {
        "sysmapid": "9",
        "userid": "1"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "sysmapids": [
            "9"
        ]
    },
    "id": 1
}
```

See also

- Map element
- Map link
- Map URL
- Map user
- Map user group
- Map shapes
- Map lines

Source

CMap::update() in *ui/include/classes/api/services/CMap.php*.

Media type

This class is designed to work with media types.

Object references:

- Media type
  - Webhook parameters
  - Script parameters
- Message template

Available methods:

- mediatype.create - create new media types
- mediatype.delete - delete media types
- mediatype.get - retrieve media types
- mediatype.update - update media types

Media type object

The following objects are directly related to the `mediatype` API.

Media type

The media type object has the following properties.

Property	Type	Description
mediatypeid	ID	ID of the media type. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the media type. Property behavior: - *required* for create operations
type	integer	Transport used by the media type. Possible values: 0 - Email; 1 - Script; 2 - SMS; 4 - Webhook; 5 - Push. Property behavior: - *required* for create operations
exec_path	string	Name of the script file (e.g., notification.sh) that is located in the directory specified in the AlertScriptsPath server configuration parameter. Property behavior: - *required* if `type` is set to ”Script”

Property	Type	Description
gsm_modem	string	Serial device name of the GSM modem. Property behavior: - *required* if `type` is set to ”SMS”
passwd	string	Authentication password. Property behavior: - *supported* if `smtp_authentication` is set to ”Normal password”
provider	integer	Email provider. Possible values: 0 - *(default)* Generic SMTP; 1 - Gmail; 2 - Gmail relay; 3 - Office365; 4 - Office365 relay.
smtp_email	string	Email address from which notifications will be sent. Property behavior: - *required* if `type` is set to ”Email”
smtp_helo	string	SMTP HELO. Property behavior: - *supported* if `type` is set to ”Email”
smtp_server	string	SMTP server. Property behavior: - *required* if `type` is set to ”Email”
smtp_port	integer	SMTP server port to connect to. Default: 25. Property behavior: - *supported* if `type` is set to ”Email”
smtp_security	integer	SMTP connection security level to use. Possible values: 0 - *(default)* None; 1 - STARTTLS; 2 - SSL/TLS. Property behavior: - *supported* if `type` is set to ”Email”
smtp_verify_host	integer	SSL verify host for SMTP. Possible values: 0 - *(default)* No; 1 - Yes. Property behavior: - *supported* if `smtp_security` is set to ”STARTTLS” or ”SSL/TLS”
smtp_verify_peer	integer	SSL verify peer for SMTP. Possible values: 0 - *(default)* No; 1 - Yes. Property behavior: - *supported* if `smtp_security` is set to ”STARTTLS” or ”SSL/TLS”

Property	Type	Description
smtp_authentication	integer	SMTP authentication method to use. Possible values: 0 - *(default)* None; 1 - Normal password; 2 - OAuth token. OAuth authentication is not allowed for *Office365 relay* email provider. Property behavior: - *supported* if `type` is set to ”Email”
redirection_url	string	Zabbix frontend URL to redirect back OAuth authorization. Default: Value of API settings property `url` with part `zabbix.php?action=oauth.authorize` Property behavior: - *required* if `smtp_authentication` is set to ”OAuth token”
client_id	string	The client identifier registered within the OAuth authorization server. Property behavior: - *required* if `smtp_authentication` is set to ”OAuth token”
client_secret	string	The client secret registered within the OAuth authorization server. Accessible only for user of type Super Admin. Property behavior: - *required* if `smtp_authentication` is set to ”OAuth token”
authorization_url	string	OAuth URL, with parameters, to get access and refresh tokens. Property behavior: - *required* if `smtp_authentication` is set to ”OAuth token”
token_url	string	OAuth URL to exchange authorization token to access and refresh tokens. This URL also is used by server to refresh invalid access token. Property behavior: - *required* if `smtp_authentication` is set to ”OAuth token”
tokens_status	integer	Bit mask on tokens’ status. Possible values: 0 - (default) Both tokens contain invalid value 1 - Access token contains valid value 2 - Refresh token contains valid value 3 - Both tokens contain valid value. Property behavior: - *supported* if `smtp_authentication` is set to ”OAuth token”
access_token	string	OAuth access token value. Property behavior: - *required* if `smtp_authentication` is set to ”OAuth token”
access_token_updated	timestamp	Timestamp of last modification of `access_token` done by server when refreshing with `refresh_token` or API on token changes. Property behavior: - *supported* if `smtp_authentication` is set to ”OAuth token”

Property	Type	Description
access_expires_in	integer	Time in seconds when `access_token` will become outdated and will require to make request to `refresh_url`. Is set by Zabbix server on `access_token` refresh or by API on token changes. Timestamp is calculated by adding value of `access_token_updated`. Property behavior: - *supported* if `smtp_authentication` is set to ”OAuth token”
refresh_token	string	OAuth refresh token value. Property behavior: - *required* if `smtp_authentication` is set to ”OAuth token”
status	integer	Whether the media type is enabled. Possible values: 0 - *(default)* Enabled; 1 - Disabled.
username	string	User name. Property behavior: - *supported* if `smtp_authentication` is set to ”Normal password”
maxsessions	integer	The maximum number of alerts that can be processed in parallel. Possible values if `type` is set to ”SMS”: 1. Possible values if `type` is set to ”Email”, ”Script”, or ”Webhook”: 0-100. Default: 1.
maxattempts	integer	The maximum number of attempts to send an alert. Possible values: 1-100. Default: 3.
attempt_interval	string	The interval between retry attempts. Accepts a number of seconds (e.g., 1, 30, 60), or a value with a time suffix (e.g., 30s, 1m, 1h). Possible values: 0-1h. Default: 10s.
message_format	integer	Message format. Possible values: 0 - Plain text; 1 - *(default)* HTML. Property behavior: - *supported* if `type` is set to ”Email”
script	text	Webhook script body (JavaScript). Property behavior: - *required* if `type` is set to ”Webhook”

Property	Type	Description
timeout	string	Webhook script timeout. Accepts a number of seconds (e.g., 1, 60), or a value with a time suffix (e.g., 30s, 1m) Possible values: 1-60s. Default: 30s. Property behavior: - *supported* if `type` is set to ”Webhook”
process_tags	integer	Process JSON property values in Webhook script response as tags. These tags are added to any existing problem tags. Possible values: 0 - *(default)* Ignore webhook script response; 1 - Process webhook script response as tags. Property behavior: - *supported* if `type` is set to ”Webhook”
show_event_menu	integer	Include an entry in the event menu that links to a custom URL. Also adds the `urls` property to the output of problem.get and event.get. Possible values: 0 - *(default)* Do not include event menu entry or `urls` property; 1 - Include event menu entry and `urls` property. Property behavior: - *supported* if `type` is set to ”Webhook”
event_menu_url	string	URL used in the event menu entry and in the `urls` property returned by problem.get and event.get. Property behavior: - *supported* if `show_event_menu` is set to ”Include event menu entry and `urls` property”
event_menu_name	string	Name used for the event menu entry and in the `urls` property returned by problem.get and event.get. Property behavior: - *supported* if `show_event_menu` is set to ”Include event menu entry and `urls` property”
parameters	array	Webhook or script parameters. Property behavior: - *supported* if `type` is set to ”Webhook” or ”Script”
description	text	Media type description.

Webhook parameters

Webhook parameters have the following properties.

Property	Type	Description
name	string	Parameter name. Property behavior: - *required*
value	string	Parameter value, supports macros. Supported macros are described on the Supported macros page.

Script parameters

Script parameters have the following properties.

Property	Type	Description
sortorder	integer	The order in which parameter values will be passed to the script as command-line arguments, starting with 0 as the first one. Property behavior: - *required*
value	string	Parameter value, supports macros. Supported macros are described on the Supported macros page.

Message template

The message template object defines a template that will be used as a default message for action operations to send a notification. It has the following properties.

Property	Type	Description
eventsource	integer	Event source. Possible values: 0 - Triggers; 1 - Discovery; 2 - Autoregistration; 3 - Internal; 4 - Services. Property behavior: - *required*
recovery	integer	Operation mode. Possible values: 0 - Operations; 1 - Recovery operations; 2 - Update operations. Property behavior: - *required*
subject	string	Message subject.
message	string	Message text.

mediatype.create

Description

`object mediatype.create(object/array mediaTypes)`

This method allows you to create new media types.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.
 If delivery methods are specified in the `$ZBX_FEATURE_FLAGS['media_type_denylist']` feature flag in the frontend configuration file, the user is not allowed to call the method for the corresponding media types.

Parameters

`(object/array)` Media types to create.

Additionally to the standard media type properties, the method accepts the following parameters.

Parameter	Type	Description
message_templates	array	Message templates to be created for the media type.

Return values

(object) Returns an object containing the IDs of the created media types under the mediatypeids property. The order of the returned IDs matches the order of the passed media types.

Examples

Creating an email media type

Create a new email media type with a custom SMTP port and message templates.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "mediatype.create",
    "params": {
        "type": "0",
        "name": "Email",
        "smtp_server": "mail.example.com",
        "smtp_helo": "example.com",
        "smtp_email": "zabbix@example.com",
        "smtp_port": "587",
        "message_format": "1",
        "message_templates": [
            {
                "eventsource": "0",
                "recovery": "0",
                "subject": "Problem: {EVENT.NAME}",
                "message": "Problem \"{EVENT.NAME}\" on host \"{HOST.NAME}\" started at {EVENT.TIME}."
            },
            {
                "eventsource": "0",
                "recovery": "1",
                "subject": "Resolved in {EVENT.DURATION}: {EVENT.NAME}",
                "message": "Problem \"{EVENT.NAME}\" on host \"{HOST.NAME}\" has been resolved at {EVENT.R
            },
            {
                "eventsource": "0",
                "recovery": "2",
                "subject": "Updated problem in {EVENT.AGE}: {EVENT.NAME}",
                "message": "{USER.FULLNAME} {EVENT.UPDATE.ACTION} problem \"{EVENT.NAME}\" on host \"{HOST
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "mediatypeids": [
            "7"
        ]
    },
    "id": 1
}
```

Creating a script media type

Create a new script media type with a custom value for the number of attempts and the interval between them.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "mediatype.create",
    "params": {
        "type": "1",
        "name": "Push notifications",
        "exec_path": "push-notification.sh",
        "maxattempts": "5",
        "attempt_interval": "11s",
        "parameters": [
            {
                "sortorder": "0",
                "value": "{ALERT.SENDTO}"
            },
            {
                "sortorder": "1",
                "value": "{ALERT.SUBJECT}"
            },
            {
                "sortorder": "2",
                "value": "{ALERT.MESSAGE}"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "mediatypeids": [
            "8"
        ]
    },
    "id": 1
}
```

Creating a webhook media type

Create a new webhook media type.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "mediatype.create",
    "params": {
        "type": "4",
        "name": "Webhook",
        "script": "var Webhook = {\r\n    token: null,\r\n    to: null,\r\n    subject: null,\r\n    messa
        "parameters": [
            {
                "name": "Message",
                "value": "{ALERT.MESSAGE}"
            },
            {
                "name": "Subject",
                "value": "{ALERT.SUBJECT}"
            },
            {
                "name": "To",
                "value": "{ALERT.SENDTO}"
```

```
            },
            {
                "name": "Token",
                "value": "<Token>"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "mediatypeids": [
            "9"
        ]
    },
    "id": 1
}
```

Source

CMediaType::create() in *ui/include/classes/api/services/CMediaType.php*.

mediatype.delete

Description

`object mediatype.delete(array mediaTypeIds)`

This method allows you to delete media types.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.
 If delivery methods are specified in the `$ZBX_FEATURE_FLAGS['media_type_denylist']` feature flag in the frontend configuration file, the user is not allowed to call the method for the corresponding media types.

Parameters

`(array)` IDs of the media types to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted media types under the `mediatypeids` property.

Examples

Deleting multiple media types

Delete two media types.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "mediatype.delete",
    "params": [
        "3",
        "5"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "mediatypeids": [
            "3",
            "5"
        ]
    },
    "id": 1
}
```

Source

CMediaType::delete() in *ui/include/classes/api/services/CMediaType.php*.

mediatype.get

Description

`integer/array mediatype.get(object parameters)`

The method allows you to retrieve media types according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.
 If delivery methods are specified in the `$ZBX_FEATURE_FLAGS['media_type_denylist']` feature flag in the frontend configuration file, the user is not allowed to call the method for the corresponding media types.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Note:
When requesting user-related information of media types, *Admin* and *User* type users may retrieve data only about their own user. For an example, see Retrieving media types as *Admin*.

Parameter	Type	Description
mediatypeids	ID/array	Return only media types with the given IDs.
mediaids	ID/array	Return only media types used by the given media.
userids	ID/array	Return only media types used by the given users.
selectActions	query	Return an `actions` property with the actions that use the media type.
selectMessageTemplates	query	Return a `message_templates` property with an array of media type messages. Parameter behavior: - *supported* for *Super admin* type users
selectUsers	query	Return a `users` property with the users that use the media type. See `user.get` for restrictions based on user type.
sortfield	string/array	Sort the result by the given properties. Possible values: `mediatypeid`.

Parameter	Type	Description
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type. Possible Media type object properties for *Admin* and *User* type users: `mediatypeid`, `name`, `type`, `status`, `maxattempts`.
output	query	Media type object properties to be returned. *Admin* and *User* type users may retrieve only the following properties: `mediatypeid`, `name`, `type`, `status`, `maxattempts`, `description`. For an example, see Retrieving media types as *Admin*. Default: `extend`.
search	object	Return results that match the given pattern (case-insensitive). Accepts an object, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a `LIKE "%…%"` search. Supports only properties of `string` and `text` data type. Possible Media type object properties for *Admin* and *User* type users: `name`, `description`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
limit	integer	
preservekeys	boolean	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving media types

Retrieve all configured media types. The following example returns two media types:

- email media type;
- SMS media type.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "mediatype.get",
    "params": {
        "output": "extend",
        "selectMessageTemplates": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "mediatypeid": "1",
            "type": "0",
            "name": "Email",
            "smtp_server": "mail.example.com",
            "smtp_helo": "example.com",
            "smtp_email": "zabbix@example.com",
            "exec_path": "",
            "gsm_modem": "",
            "username": "",
            "passwd": "",
            "status": "0",
            "smtp_port": "25",
            "smtp_security": "0",
            "smtp_verify_peer": "0",
            "smtp_verify_host": "0",
            "smtp_authentication": "0",
            "maxsessions": "1",
            "maxattempts": "3",
            "attempt_interval": "10s",
            "message_format": "0",
            "script": "",
            "timeout": "30s",
            "process_tags": "0",
            "show_event_menu": "1",
            "event_menu_url": "",
            "event_menu_name": "",
            "description": "",
            "provider": "0",
            "message_templates": [
                {
                    "eventsource": "0",
                    "recovery": "0",
                    "subject": "Problem: {EVENT.NAME}",
                    "message": "Problem started at {EVENT.TIME} on {EVENT.DATE}\r\nProblem name: {EVENT.NA
                },
                {
                    "eventsource": "0",
                    "recovery": "1",
                    "subject": "Resolved: {EVENT.NAME}",
                    "message": "Problem has been resolved at {EVENT.RECOVERY.TIME} on {EVENT.RECOVERY.DATE
                },
                {
                    "eventsource": "0",
                    "recovery": "2",
                    "subject": "Updated problem: {EVENT.NAME}",
                    "message": "{USER.FULLNAME} {EVENT.UPDATE.ACTION} problem at {EVENT.UPDATE.DATE} {EVEN
                },
                {
                    "eventsource": "1",
                    "recovery": "0",
                    "subject": "Discovery: {DISCOVERY.DEVICE.STATUS} {DISCOVERY.DEVICE.IPADDRESS}",
                    "message": "Discovery rule: {DISCOVERY.RULE.NAME}\r\n\r\nDevice IP: {DISCOVERY.DEVICE.
                },
                {
                    "eventsource": "2",
                    "recovery": "0",
                    "subject": "Autoregistration: {HOST.HOST}",
                    "message": "Host name: {HOST.HOST}\r\nHost IP: {HOST.IP}\r\nAgent port: {HOST.PORT}"
```

```
            }
        ],
        "parameters": []
    },
    {
        "mediatypeid": "3",
        "type": "2",
        "name": "SMS",
        "smtp_server": "",
        "smtp_helo": "",
        "smtp_email": "",
        "exec_path": "",
        "gsm_modem": "/dev/ttyS0",
        "username": "",
        "passwd": "",
        "status": "0",
        "smtp_port": "25",
        "smtp_security": "0",
        "smtp_verify_peer": "0",
        "smtp_verify_host": "0",
        "smtp_authentication": "0",
        "maxsessions": "1",
        "maxattempts": "3",
        "attempt_interval": "10s",
        "message_format": "1",
        "script": "",
        "timeout": "30s",
        "process_tags": "0",
        "show_event_menu": "1",
        "event_menu_url": "",
        "event_menu_name": "",
        "description": "",
        "provider": "0",
        "message_templates": [
            {
                "eventsource": "0",
                "recovery": "0",
                "subject": "",
                "message": "{EVENT.SEVERITY}: {EVENT.NAME}\r\nHost: {HOST.NAME}\r\n{EVENT.DATE} {EVENT
            },
            {
                "eventsource": "0",
                "recovery": "1",
                "subject": "",
                "message": "RESOLVED: {EVENT.NAME}\r\nHost: {HOST.NAME}\r\n{EVENT.DATE} {EVENT.TIME}"
            },
            {
                "eventsource": "0",
                "recovery": "2",
                "subject": "",
                "message": "{USER.FULLNAME} {EVENT.UPDATE.ACTION} problem at {EVENT.UPDATE.DATE} {EVEN
            },
            {
                "eventsource": "1",
                "recovery": "0",
                "subject": "",
                "message": "Discovery: {DISCOVERY.DEVICE.STATUS} {DISCOVERY.DEVICE.IPADDRESS}"
            },
            {
                "eventsource": "2",
                "recovery": "0",
                "subject": "",
```

```
                    "message": "Autoregistration: {HOST.HOST}\r\nHost IP: {HOST.IP}\r\nAgent port: {HOST.P
                }
            ],
            "parameters": []
        }
    ],
    "id": 1
}
```

Retrieving media types as *Admin*

As an *Admin* type user, retrieve all media types that are enabled, with users that use these media types. The following example returns two media types:

- email media type with one user (only *Admin* type user's own user);
- SMS media type with no users.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "mediatype.get",
    "params": {
        "output": "extend",
        "filter": {
            "status": 0
        },
        "selectUsers": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "mediatypeid": "1",
            "type": "0",
            "name": "Email",
            "status": "0",
            "description": "",
            "maxattempts": "3",
            "users": [
                {
                    "userid": "3",
                    "username": "database-admin",
                    "name": "John",
                    "surname": "Doe",
                    "url": "",
                    "autologin": "0",
                    "autologout": "0",
                    "lang": "default",
                    "refresh": "30s",
                    "theme": "default",
                    "attempt_failed": "0",
                    "attempt_ip": "",
                    "attempt_clock": "0",
                    "rows_per_page": "50",
                    "timezone": "default",
                    "roleid": "2",
                    "provisioned": "0"
                }
            ]
        },
```

```
        {
            "mediatypeid": "3",
            "type": "2",
            "name": "SMS",
            "status": "0",
            "description": "",
            "maxattempts": "3",
            "users": []
        }
    ],
    "id": 1
}
```

Retrieve script and webhook media types

The following example returns three media types:

- script media type with parameters;
- script media type without parameters;
- webhook media type with parameters.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "mediatype.get",
    "params": {
        "output": ["mediatypeid", "name", "parameters"],
        "filter": {
            "type": [1, 4]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "mediatypeid": "10",
            "name": "Script with parameters",
            "parameters": [
                {
                    "sortorder": "0",
                    "value": "{ALERT.SENDTO}"
                },
                {
                    "sortorder": "1",
                    "value": "{EVENT.NAME}"
                },
                {
                    "sortorder": "2",
                    "value": "{ALERT.MESSAGE}"
                },
                {
                    "sortorder": "3",
                    "value": "Zabbix alert"
                }
            ]
        },
        {
            "mediatypeid": "13",
            "name": "Script without parameters",
            "parameters": []
```

```
        },
        {
            "mediatypeid": "11",
            "name": "Webhook",
            "parameters": [
                {
                    "name": "alert_message",
                    "value": "{ALERT.MESSAGE}"
                },
                {
                    "name": "event_update_message",
                    "value": "{EVENT.UPDATE.MESSAGE}"
                },
                {
                    "name": "host_name",
                    "value": "{HOST.NAME}"
                },
                {
                    "name": "trigger_description",
                    "value": "{TRIGGER.DESCRIPTION}"
                },
                {
                    "name": "trigger_id",
                    "value": "{TRIGGER.ID}"
                },
                {
                    "name": "alert_source",
                    "value": "Zabbix"
                }
            ]
        }
    ],
    "id": 1
}
```

See also

- User

Source

CMediaType::get() in *ui/include/classes/api/services/CMediaType.php*.

mediatype.update

Description

`object mediatype.update(object/array mediaTypes)`

This method allows you to update existing media types.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.
 If delivery methods are specified in the `$ZBX_FEATURE_FLAGS['media_type_denylist']` feature flag in the frontend configuration file, the user is not allowed to call the method for the corresponding media types.

Parameters

`(object/array)` Media type properties to be updated.

The `mediatypeid` property must be defined for each media type, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard media type properties, the method accepts the following parameters.

Parameter	Type	Description
message_templates	array	Message templates to replace the current message templates.

Return values

(object) Returns an object containing the IDs of the updated media types under the mediatypeids property.

Examples

Enabling a media type

Enable a media type, that is, set its status to "0".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "mediatype.update",
    "params": {
        "mediatypeid": "6",
        "status": "0"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "mediatypeids": [
            "6"
        ]
    },
    "id": 1
}
```

Source

CMediaType::update() in *ui/include/classes/api/services/CMediaType.php*.

MFA

This class is designed to work with MFA (Multi-Factor Authentication) methods.

Object references:

- MFA

Available methods:

- mfa.create - create new MFA methods
- mfa.delete - delete MFA methods
- mfa.get - retrieve MFA methods
- mfa.update - update MFA methods

MFA object

The following objects are directly related to the mfa API.

MFA

The MFA (Multi-Factor Authentication) object has the following properties.

Property	Type	Description
mfaid	ID	ID of the MFA method. Property behavior: - *read-only* - *required* for update operations
type	integer	Type of the MFA method. Possible values: 1 - TOTP (Time-based One-Time Passwords); 2 - Duo Universal Prompt.
name	string	Unique name of the MFA method. Property behavior: - *required* for create operations
hash_function	integer	Type of the hash function for generating TOTP codes. Possible values: 1 - SHA-1; 2 - SHA-256; 3 - SHA-512. Property behavior: - *required* if `type` is set to ”TOTP”
code_length	integer	Verification code length. Possible values: 6 - 6-digit long; 8 - 8-digit long. Property behavior: - *required* if `type` is set to ”TOTP”
api_hostname	string	API hostname provided by the Duo authentication service. Property behavior: - *required* if `type` is set to ”Duo Universal Prompt”
clientid	string	Client ID provided by the Duo authentication service. Property behavior: - *required* if `type` is set to ”Duo Universal Prompt”
client_secret	string	Client secret provided by the Duo authentication service. Property behavior: - *write-only* - *required* if `type` is set to ”Duo Universal Prompt”

mfa.create

Description

`object mfa.create(object/array MFA methods)`

This method allows you to create new MFA methods.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` MFA methods to create.

The method accepts MFA methods with the standard MFA method properties.

Return values

(object) Returns an object containing the IDs of the created MFA methods under the mfaids property. The order of the returned IDs matches the order of the passed items.

Examples

Creating MFA methods

Create a "Zabbix TOTP" MFA method utilizing time-based one-time passwords (TOTP), with the hash function for generating TOTP codes set to SHA-1 and the verification code length set to 6 digits.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "mfa.create",
    "params": {
        "type": 1,
        "name": "Zabbix TOTP",
        "hash_function": 1,
        "code_length": 6
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "mfaids": [
            "1"
        ]
    },
    "id": 1
}
```

Source

CMfa::create() in *ui/include/classes/api/services/CMfa.php*.

mfa.delete

Description

`object mfa.delete(array mfaids)`

This method allows you to delete MFA methods.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

(array) IDs of the MFA methods to delete.

Return values

(object) Returns an object containing the IDs of the deleted MFA methods under the mfaids property.

Examples

Deleting MFA methods

Delete an MFA method.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "mfa.delete",
    "params": [
        "2"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "mfaids": [
            "2"
        ]
    },
    "id": 1
}
```

Source

CMfa::delete() in *ui/include/classes/api/services/CMfa.php*.

mfa.get

Description

`integer/array mfa.get(object parameters)`

The method allows you to retrieve MFA methods according to the given parameters.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
mfaids	ID/array	Return only MFA methods with the given IDs.
selectUsrgrps	query	Return a `usrgrps` property with user groups associated with MFA methods. Supports `count`.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Supports properties: `mfaid` - ID of the MFA method; `type` - Type of the MFA method.
sortfield	string/array	Sort the result by the given properties. Possible values: `name`.

Parameter	Type	Description
search	object	Return results that match the given pattern (case-insensitive). Accepts an object, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a `LIKE "%…%"` search. Supported properties: `name`.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
limit	integer	
output	query	
preservekeys	boolean	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Finding MFA methods by name

Retrieve all MFA methods with ”Zabbix” in their name.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "mfa.get",
    "params": {
        "output": "extend",
        "search": {
            "name": "Zabbix"
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "mfaid": "1",
            "type": "1",
            "name": "Zabbix TOTP 1",
            "hash_function": "1",
            "code_length": "6",
            "api_hostname": "",
            "clientid": ""
        },
        {
            "mfaid": "2",
            "type": "1",
            "name": "Zabbix TOTP 2",
            "hash_function": "3",
            "code_length": "8",
            "api_hostname": "",
```

```
            "clientid": ""
        }
    ],
    "id": 1
}
```

Source

CMfa::get() in *ui/include/classes/api/services/CMfa.php*.

mfa.update

Description

`object mfa.update(object/array MFA methods)`

This method allows you to update existing MFA methods.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` MFA method properties to be updated.

The `mfaid` property must be defined for each item, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

The method accepts MFA methods with the standard MFA method properties.

Return values

`(object)` Returns an object containing the IDs of the updated MFA methods under the `mfaids` property.

Examples

Updating method properties

Update the hash function for generating TOTP codes and the verification code length for the "Zabbix TOTP" MFA method utilizing time-based one-time passwords (TOTP).

Request:

```
{
    "jsonrpc": "2.0",
    "method": "mfa.update",
    "params": {
        "mfaid": "1",
        "hash_function": 3,
        "code_length": 8
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "mfaids": [
            "1"
        ]
    },
    "id": 1
}
```

Source

CMfa::update() in *ui/include/classes/api/services/CMfa.php*.

Module

This class is designed to work with frontend modules.

Object references:

- Module

Available methods:

- module.create - install new modules
- module.delete - uninstall modules
- module.get - retrieve modules
- module.update - update modules

Module object

The following objects are directly related to the `module` API.

Module

The module object has the following properties.

Property	Type	Description
moduleid	ID	ID of the module as stored in the database. Property behavior: - *read-only* - *required* for update operations
id	string	Unique module ID as defined by a developer in the `manifest.json` file of the module. Possible values for built-in modules: see property "type" description in Dashboard widget. Property behavior: - *required* for create operations
relative_path	string	Path to the directory of the module relative to the directory of the Zabbix frontend. Possible values: `widgets/*` - for built-in widget modules; `modules/*` - for third-party modules. Property behavior: - *required* for create operations
status	integer	Whether the module is enabled or disabled. Possible values: 0 - *(default)* Disabled; 1 - Enabled.
config	object	Module configuration.

module.create

Description

`object module.create(object/array modules)`

This method allows you to install new frontend modules.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.
 If `$ZBX_FEATURE_FLAGS['modules_config_enabled']` is set to `false` in the frontend configuration file, the user will not be allowed to call this method.

Attention:
Module files must be unpacked manually in the correct subdirectories, matching the `relative_path` property of the modules.

Parameters

`(object/array)` Modules to create.

The method accepts modules with the standard module properties.

Return values

`(object)` Returns an object containing the IDs of the installed modules under the `moduleids` property. The order of the returned IDs matches the order of the passed modules.

Examples

Installing a module

Install a module with the status "Enabled".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "module.create",
    "params": {
        "id": "example_module",
        "relative_path": "modules/example_module",
        "status": 1
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "moduleids": [
            "25"
        ]
    },
    "id": 1
}
```

See also

- Module
- Frontend modules

Source

CModule::create() in *ui/include/classes/api/services/CModule.php*.

module.delete

Description

`object module.delete(array moduleids)`

This method allows you to uninstall modules.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.
 If `$ZBX_FEATURE_FLAGS['modules_config_enabled']` is set to `false` in the frontend configuration file, the user will not be allowed to call this method.

Attention:
Module files must be removed manually.

Parameters

`(array)` IDs of the modules to uninstall.

Return values

`(object)` Returns an object containing the IDs of the uninstalled modules under the `moduleids` property.

Examples

Uninstalling multiple modules

Uninstall modules "27" and "28".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "module.delete",
    "params": [
        "27",
        "28"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "moduleids": [
            "27",
            "28"
        ]
    },
    "id": 1
}
```

Source

CModule::delete() in *ui/include/classes/api/services/CModule.php*.

module.get

Description

`integer/array module.get(object parameters)`

The method allows you to retrieve modules according to the given parameters.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.
 If `$ZBX_FEATURE_FLAGS['modules_config_enabled']` is set to `false` in the frontend configuration file, the user will not be allowed to call this method.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
moduleids	ID/array	Return only modules with the given IDs.

Parameter	Type	Description
sortfield	string/array	Sort the result by the given properties. Possible values: `moduleid`, `relative_path`.
countOutput	boolean	These parameters are described in the Reference commentary.
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving a module by ID

Retrieve all data about modules "1", "2", and "25".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "module.get",
    "params": {
        "output": "extend",
        "moduleids": [
            "1",
            "2",
            "25"
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "moduleid": "1",
            "id": "actionlog",
            "relative_path": "widgets/actionlog",
            "status": "1",
            "config": []
        },
        {
            "moduleid": "2",
            "id": "clock",
            "relative_path": "widgets/clock",
            "status": "1",
            "config": []
        },
        {
            "moduleid": "25",
```

```
            "id": "example",
            "relative_path": "modules/example_module",
            "status": "1",
            "config": []
        }
    ],
    "id": 1
}
```

See also

- Module
- Dashboard widget
- Frontend modules

Source

CModule::get() in *ui/include/classes/api/services/CModule.php*.

module.update

Description

`object module.update(object/array modules)`

This method allows you to update existing modules.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.
 If `$ZBX_FEATURE_FLAGS['modules_config_enabled']` is set to `false` in the frontend configuration file, the user will not be allowed to call this method.

Parameters

`(object/array)` Module properties to be updated.

The `moduleid` property must be defined for each module, all other properties are optional. Only the specified properties will be updated.

The method accepts modules with the standard module properties.

Return values

`(object)` Returns an object containing the IDs of the updated modules under the `moduleids` property.

Examples

Disabling a module

Disable module "25".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "module.update",
    "params": {
        "moduleid": "25",
        "status": 0
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "moduleids": [
            "25"
```

```
        ]
    },
    "id": 1
}
```

See also

- Module
- Frontend modules

Source

CModule::update() in *ui/include/classes/api/services/CModule.php*.

Problem

This class is designed to work with problems.

Object references:

- Problem
  - Media type URL
- Problem tag

Available methods:

- problem.get - retrieve problems

Problem object

The following objects are directly related to the `problem` API.

Problem

Note:
Problems are created by the Zabbix server and cannot be modified via the API.

The problem object has the following properties.

Property	Type	Description
eventid	ID	ID of the problem event.
source	integer	Type of the problem event. Possible values: 0 - event created by a trigger; 3 - internal event; 4 - event created on service status update.
object	integer	Type of object that is related to the problem event. Possible values if `source` is set to ”event created by a trigger”: 0 - trigger. Possible values if `source` is set to ”internal event”: 0 - trigger; 4 - item; 5 - LLD rule. Possible values if `source` is set to ”event created on service status update”: 6 - service.
objectid	ID	ID of the related object.
clock	timestamp	Time when the problem event was created.
ns	integer	Nanoseconds when the problem event was created.
r_eventid	ID	ID of the recovery event.
r_clock	timestamp	Time when the recovery event was created.

Property	Type	Description
r_ns	integer	Nanoseconds when the recovery event was created.
cause_eventid	ID	ID of the cause event.
correlationid	ID	ID of the correlation rule if this event was recovered by a global correlation rule.
userid	ID	ID of the user that closed the problem (if the problem was closed manually).
name	string	Resolved problem name.
acknowledged	integer	Acknowledge state for problem. Possible values: 0 - not acknowledged; 1 - acknowledged.
severity	integer	Problem current severity. Possible values: 0 - not classified; 1 - information; 2 - warning; 3 - average; 4 - high; 5 - disaster.
suppressed	integer	Whether the problem is suppressed. Possible values: 0 - problem is in normal state; 1 - problem is suppressed.
opdata	string	Operational data with expanded macros.
urls	array	Active media type URLs.

Media type URL

The media type URL object has the following properties.

Property	Type	Description
name	string	Media type defined URL name.
url	string	Media type defined URL value.

Results will contain entries only for active media types with enabled event menu entry. Macro used in properties will be expanded, but if one of the properties contains an unexpanded macro, both properties will be excluded from results. For supported macros, see *Supported macros*.

Problem tag

The problem tag object has the following properties.

Property	Type	Description
tag	string	Problem tag name.
value	string	Problem tag value.

problem.get

Description

`integer/array problem.get(object parameters)`

The method allows you to retrieve problems according to the given parameters.

This method is for retrieving unresolved problems. It is also possible, if specified, to additionally retrieve recently resolved problems. The period that determines how old is "recently" is defined in *Administration* > *General*. Problems that were resolved prior to that period are not kept in the problem table. To retrieve problems that were resolved further back in the past, use the `event.get` method.

Attention:
This method may return problems of a deleted entity if these problems have not been removed by the housekeeper yet.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
eventids	ID/array	Return only problems with the given IDs.
groupids	ID/array	Return only problems created by objects that belong to the given host groups.
hostids	ID/array	Return only problems created by objects that belong to the given hosts.
objectids	ID/array	Return only problems created by the given objects.
source	integer	Return only problems with the given type. Refer to the problem event object page for a list of supported event types. Default: 0 - problem created by a trigger.
object	integer	Return only problems created by objects of the given type. Refer to the problem event object page for a list of supported object types. Default: 0 - trigger.
acknowledged	boolean	If set to `true` return only acknowledged problems.
action	integer	Return only problems for which the given event update actions have been performed. For multiple actions, use a sum of any acceptable bitmap values as bitmask (for example, 34 for acknowledge and suppress event).
action_userids	ID/array	Return only problems with the given IDs of users who performed the problem event update actions.
suppressed	boolean	If set to `true` return only suppressed problems.
symptom	boolean	If set to `true` return only symptom problem events.
severities	integer/array	Return only problems with given event severities. Applies only if object is trigger.
evaltype	integer	Tag evaluation method. Possible values: 0 - *(default)* And/Or; 2 - Or.
tags	array	Return only problems with the given tags. Format: `[{"tag": "<tag>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all problems. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
recent	boolean	If set to `true`, return recent problems, including those resolved within the `ok_period` and any unresolved problems. If set to `false`, return only unresolved problems.

Parameter	Type	Description
eventid_from	string	Return only problems with IDs greater or equal to the given ID.
eventid_till	string	Return only problems with IDs less or equal to the given ID.
time_from	timestamp	Return only problems that have been created after or at the given time.
time_till	timestamp	Return only problems that have been created before or at the given time.
selectAcknowledges	query	Return an `acknowledges` property with the problem updates. Problem updates are sorted in reverse chronological order. The problem update object has the following properties: `acknowledgeid` - `(ID)` ID of the update; `userid` - `(ID)` ID of the user that updated the event; `eventid` - `(ID)` ID of the updated event; `clock` - `(timestamp)` time when the event was updated; `message` - `(string)` text of the message; `action` - `(integer)`type of update action (see `event.acknowledge`); `old_severity` - `(integer)` event severity before this update action; `new_severity` - `(integer)` event severity after this update action; `suppress_until` - `(timestamp)` time till event will be suppressed; `taskid` - `(ID)` ID of the task if the current event is undergoing a rank change; `maintenanceid` - `(ID)` ID of the maintenance. Supports `count`.
selectTags	query	Return a `tags` property with the problem tags. Output format: `[{"tag": "<tag>", "value": "<value>"}, ...]`.
selectSuppressionData	query	Return a `suppression_data` property with the list of active maintenances and manual suppressions: `maintenanceid` - `(ID)` ID of the maintenance; `userid` - `(ID)` ID of user who suppressed the problem; `suppress_until` - `(integer)` time until the problem is suppressed.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type.
sortfield	string/array	Sort the result by the given properties. Possible values: `eventid`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving trigger problem events

Retrieve recent events from trigger ”15112.”

Request:

```
{
    "jsonrpc": "2.0",
    "method": "problem.get",
    "params": {
        "output": "extend",
        "selectAcknowledges": "extend",
        "selectTags": "extend",
        "selectSuppressionData": "extend",
        "objectids": "15112",
        "recent": true,
        "sortfield": ["eventid"],
        "sortorder": "DESC"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "eventid": "1245463",
            "source": "0",
            "object": "0",
            "objectid": "15112",
            "clock": "1472457242",
            "ns": "209442442",
            "r_eventid": "1245468",
            "r_clock": "1472457285",
            "r_ns": "125644870",
            "correlationid": "0",
            "userid": "1",
            "name": "Zabbix agent on localhost is unreachable for 5 minutes",
            "acknowledged": "1",
            "severity": "3",
            "cause_eventid": "0",
            "suppressed": "1",
            "acknowledges": [
                {
                    "acknowledgeid": "14443",
                    "userid": "1",
                    "eventid": "1245463",
                    "clock": "1472457281",
                    "message": "problem solved",
                    "action": "6",
                    "old_severity": "0",
                    "new_severity": "0",
                    "suppress_until": "1472511600",
                    "taskid": "0",
                    "maintenanceid": "0"
                }
            ],
            "opdata": "",
            "suppression_data": [
                {
                    "maintenanceid": "15",
                    "suppress_until": "1472511600",
                    "userid": "0"
                }
            ],
```

```
            "tags": [
                {
                    "tag": "test-tag",
                    "value": "test-value"
                }
            ],
            "urls": []
        }
    ],
    "id": 1
}
```

Retrieving problems acknowledged by specified user

Retrieving problems acknowledged by user with ID=10

Request:

```
{
    "jsonrpc": "2.0",
    "method": "problem.get",
    "params": {
        "output": "extend",
        "action": 2,
        "action_userids": [10],
        "selectAcknowledges": ["userid", "action"],
        "sortfield": ["eventid"],
        "sortorder": "DESC"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "eventid": "1248566",
            "source": "0",
            "object": "0",
            "objectid": "15142",
            "clock": "1472457242",
            "ns": "209442442",
            "r_eventid": "1245468",
            "r_clock": "1472457285",
            "r_ns": "125644870",
            "correlationid": "0",
            "userid": "10",
            "name": "Zabbix agent on localhost is unreachable for 5 minutes",
            "acknowledged": "1",
            "severity": "3",
            "cause_eventid": "0",
            "suppressed": "0",
            "acknowledges": [
                {
                    "acknowledgeid": "9",
                    "userid": "10",
                    "action": "2"
                }
            ],
            "opdata": "",
            "urls": []
        }
    ],
```

```
    "id": 1
}
```

See also

- Alert
- Item
- Host
- LLD rule
- Trigger

Source

CEvent::get() in *ui/include/classes/api/services/CProblem.php*.

Proxy

This class is designed to work with proxies.

Object references:

- Proxy

Available methods:

- proxy.create - create new proxies
- proxy.delete - delete proxies
- proxy.get - retrieve proxies
- proxy.update - update proxies

Proxy object

The following objects are directly related to the `proxy` API.

Proxy

The proxy object has the following properties.

Property	Type	Description
proxyid	ID	ID of the proxy. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the proxy. Property behavior: - *required* for create operations
proxy_groupid	ID	ID of the proxy group. 0 - proxy is not assigned to any proxy group.
local_address	string	Address for active agents. IP address or DNS name to connect to. Property behavior: - *required* if `proxy_groupid` is not 0
local_port	string	Local proxy port number to connect to. User macros are supported. Default: 10051. Property behavior: - *supported* if `proxy_groupid` is not 0

Property	Type	Description
operating_mode	integer	Type of proxy. Possible values: 0 - active proxy; 1 - passive proxy. Property behavior: - *required* for create operations
description	text	Description of the proxy.
lastaccess	timestamp	Time when the proxy last connected to the server. Property behavior: - *read-only*
address	string	IP address or DNS name to connect to. User macros are supported. Property behavior: - *required* if the Zabbix proxy operating mode is passive
port	string	Port number to connect to. User macros are supported. Default: 10051. Property behavior: - *supported* if the Zabbix proxy operating mode is passive
allowed_addresses	string	Comma-delimited IP addresses or DNS names of active Zabbix proxy.
tls_connect	integer	Connections to proxy. Possible values: 1 - *(default)* No encryption; 2 - PSK; 4 - certificate.
tls_accept	integer	Connections from proxy. Possible bitmap values: 1 - *(default)* No encryption; 2 - PSK; 4 - certificate. This is a bitmask field; any sum of possible bitmap values is acceptable (for example, 6 for PSK and certificate).
tls_issuer	string	Certificate issuer.
tls_subject	string	Certificate subject.
tls_psk_identity	string	PSK identity; must be paired with only one PSK (across autoregistration, hosts, and proxies). Do not include sensitive information in the PSK identity, as it is sent unencrypted over the network to inform the receiver which PSK to use. Property behavior: - *write-only* - *required* if `tls_connect` is set to "PSK", or `tls_accept` contains the "PSK" bit
tls_psk	string	Pre-shared key (PSK); must be at least 32 hex digits. Property behavior: - *write-only* - *required* if `tls_connect` is set to "PSK", or `tls_accept` contains the "PSK" bit

Property	Type	Description
custom_timeouts	integer	Whether to override global item timeouts on the proxy level. Possible values: 0 - *(default)* use global settings; 1 - override timeouts.
timeout_zabbix_agent	string	Spend no more than `timeout_*` seconds on processing. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m). Also accepts user macros. Used only if `custom_timeouts` is set to ”override timeouts”; otherwise, the value is ignored. Possible values range: 1-600s. Property behavior: - *required* if `custom_timeouts` is set to ”override timeouts”
timeout_simple_check		
timeout_snmp_agent		
timeout_external_check		
timeout_db_monitor		
timeout_http_agent		
timeout_ssh_agent		
timeout_telnet_agent		
timeout_script		
timeout_browser		
timeout_telemetry_query		
version	integer	Version of proxy. Three-part Zabbix version number, where two decimal digits are used for each part, e.g., 60401 for version 6.4.1, 70002 for version 7.0.2, etc. 0 - Unknown proxy version. Property behavior: - *read-only*
compatibility	integer	Version of proxy compared to Zabbix server version. Possible values: 0 - Undefined; 1 - Current version (proxy and server have the same major version); 2 - Outdated version (proxy version is older than server version, but is partially supported); 3 - Unsupported version (proxy version is older than server previous LTS release version or server major version is older than proxy major version). Property behavior: - *read-only*
state	integer	State of the proxy. Possible values: 0 - Unknown; 1 - Offline; 2 - Online. Property behavior: - *read-only*

Application performance monitoring (APM)

The APM object has the following properties.

Property	Type	Description
data_collection_status	integer	Whether telemetry data collection is enabled on the proxy. Possible values: 0 - *(default)* Disabled; 1 - Enabled.
max_messages_per_second	integer	Maximum number of messages per second the proxy accepts. Possible values range: 0 (unlimited) – 2147483647. Default: 0. Property behavior: - *supported* if `data_collection_status` is set to ”Enabled”
additional_resource_attributes	array	Array of resource attributes to add to incoming telemetry data. If incoming data already contains a resource attribute with the same name, the incoming attribute is kept instead of the one specified here. Property behavior: - *supported* if `data_collection_status` is set to ”Enabled”

Resource attributes

The resource attributes object has the following properties.

Property	Type	Description
key	string	Resource attribute key. Must be unique for the specified `signal_type`. Property behavior: - *required*
value	string	Resource attribute value. Can be empty.
signal_type	integer	The category of telemetry data the attribute applies to. Possible values: 0 - Traces; 1 - Metrics; 2 - Logs. Property behavior: - *required*

proxy.create

Description

`object proxy.create(object/array proxies)`

This method allows you to create new proxies.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Proxies to create.

Additionally to the standard proxy properties, the method accepts the following parameters.

Parameter	Type	Description
hosts	array	Hosts to be monitored by the proxy. If a host is already monitored by a different proxy, it will be reassigned to the current proxy. The hosts must have only the `hostid` property defined.
apm	object	Application performance monitoring (APM) settings of the proxy.

Return values

`(object)` Returns an object containing the IDs of the created proxies under the `proxyids` property. The order of the returned IDs matches the order of the passed proxies.

Examples

Create an active proxy

Create an action proxy "Active proxy" and assign a host to be monitored by it.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxy.create",
    "params": {
        "name": "Active proxy",
        "operating_mode": "0",
        "hosts": [
            {
                "hostid": "10279"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "proxyids": [
            "10280"
        ]
    },
    "id": 1
}
```

Create a passive proxy

Create a passive proxy "Passive proxy" and assign two hosts to be monitored by it.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxy.create",
    "params": {
        "name": "Passive proxy",
        "operating_mode": "1",
        "address": "127.0.0.1",
        "port": "10051",
        "hosts": [
            {
                "hostid": "10192"
            },
            {
```

```
                "hostid": "10139"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "proxyids": [
            "10284"
        ]
    },
    "id": 1
}
```

Creating a proxy and adding it to a proxy group

Create an active proxy ”Active proxy” and add it to proxy group with ID ”1”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxy.create",
    "params": {
        "name": "Active proxy",
        "proxy_groupid": "1",
        "operating_mode": "0"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "proxyids": [
            "5"
        ]
    },
    "id": 1
}
```

Creating a proxy with certificate-based encryption

Create an active proxy with certificate-based encryption. The proxy is configured to monitor host ”10785”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxy.create",
    "params": {
        "name": "Active proxy",
        "operating_mode": "0",
        "hosts": [
            {
                "hostid": "10785"
            }
        ],
        "tls_accept": "4",
        "tls_issuer": "CN=Signing CA,OU=Development group,O=Zabbix SIA,DC=zabbix,DC=com",
        "tls_subject": "CN=Zabbix server,OU=Development group,O=Zabbix SIA,DC=zabbix,DC=com"
```

```
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "proxyids": [
            "2"
        ]
    },
    "id": 1
}
```

Creating a proxy with APM enabled

Create an active proxy "Active proxy" with APM enabled. The proxy accepts up to 1000 messages per second, adds the `deployment.environment` resource attribute to traces, and adds the `cloud.region` resource attribute to metrics.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxy.create",
    "params": {
        "name": "Active proxy",
        "operating_mode": "0",
        "apm": {
            "data_collection_status": "1",
            "max_messages_per_second": "1000",
            "additional_resource_attributes": [
                {
                    "key": "deployment.environment",
                    "value": "production",
                    "signal_type": "0"
                },
                {
                    "key": "cloud.region",
                    "value": "eu-central-1",
                    "signal_type": "1"
                }
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "proxyids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Host
- Proxy group

Source

CProxy::create() in *ui/include/classes/api/services/CProxy.php*.

proxy.delete

Description

`object proxy.delete(array proxies)`

This method allows you to delete proxies.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of proxies to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted proxies under the `proxyids` property.

Examples

Delete multiple proxies

Delete two proxies.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxy.delete",
    "params": [
        "10286",
        "10285"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "proxyids": [
            "10286",
            "10285"
        ]
    },
    "id": 1
}
```

Source

CProxy::delete() in *ui/include/classes/api/services/CProxy.php*.

proxy.get

Description

`integer/array proxy.get(object parameters)`

The method allows you to retrieve proxies according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. SeeUser roles for more information.
 Only proxies allowed by the *Proxy access list* of the user's user groups are returned.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
proxyids	ID/array	Return only proxies with the given IDs.
proxy_groupids	ID/array	Return only proxies that belong to the given proxy groups.
selectApm	query	Return an `apm` property with the application performance monitoring (APM) settings of the proxy.
selectAssignedHosts	query	Return an `assignedHosts` property with the hosts assigned to the proxy. Supports `count`.
selectHosts	query	Return a `hosts` property with the hosts monitored by the proxy. Supports `count`.
selectProxyGroup	query	Return a `proxyGroup` property with the proxy group object.
sortfield	string/array	Sort the result by the given properties. Possible values: `proxyid`, `name`, `operating_mode`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve all proxies

Retrieve all configured proxies and their interfaces.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxy.get",
    "params": {
        "output": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "proxyid": "11",
            "name": "Active proxy",
```

```
            "proxy_groupid": "0",
            "local_address": "",
            "local_port": "10051",
            "operating_mode": "0",
            "description": "",
            "allowed_addresses": "",
            "address": "127.0.0.1",
            "port": "10051",
            "tls_connect": "1",
            "tls_accept": "1",
            "tls_issuer": "",
            "tls_subject": "",
            "custom_timeouts": "0",
            "timeout_zabbix_agent": "",
            "timeout_simple_check": "",
            "timeout_snmp_agent": "",
            "timeout_external_check": "",
            "timeout_db_monitor": "",
            "timeout_http_agent": "",
            "timeout_ssh_agent": "",
            "timeout_telnet_agent": "",
            "timeout_script": "",
            "last_access": "1693391880",
            "version": "70000",
            "compatibility": "1",
            "state": "1"
        },
        {
            "proxyid": "12",
            "name": "Passive proxy",
            "proxy_groupid": "1",
            "local_address": "127.0.0.1",
            "local_port": "10051",
            "operating_mode": "1",
            "description": "",
            "allowed_addresses": "",
            "address": "127.0.0.1",
            "port": "10051",
            "tls_connect": "1",
            "tls_accept": "1",
            "tls_issuer": "",
            "tls_subject": "",
            "custom_timeouts": "1",
            "timeout_zabbix_agent": "5s",
            "timeout_simple_check": "5s",
            "timeout_snmp_agent": "5s",
            "timeout_external_check": "5s",
            "timeout_db_monitor": "5s",
            "timeout_http_agent": "5s",
            "timeout_ssh_agent": "5s",
            "timeout_telnet_agent": "5s",
            "timeout_script": "5s",
            "lastaccess": "1693391875",
            "version": "60400",
            "compatibility": "2",
            "state": "2"
        }
    ],
    "id": 1
}
```


See also

- Host
- Proxy group

Source

CProxy::get() in *ui/include/classes/api/services/CProxy.php*.

proxy.update

Description

`object proxy.update(object/array proxies)`

This method allows you to update existing proxies.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Proxy properties to be updated.

The `proxyid` property must be defined for each proxy, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard proxy properties, the method accepts the following parameters.

Parameter	Type	Description
hosts	array	Hosts to be monitored by the proxy. If a host is already monitored by a different proxy, it will be reassigned to the current proxy. The hosts must have only the `hostid` property defined.
apm	object	Application performance monitoring (APM) settings of the proxy.

Return values

`(object)` Returns an object containing the IDs of the updated proxies under the `proxyids` property.

Examples

Change hosts monitored by a proxy

Update the proxy to monitor the two given hosts.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxy.update",
    "params": {
        "proxyid": "10293",
        "hosts": [
            {
                "hostid": "10294"
            },
            {
                "hostid": "10295"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "proxyids": [
            "10293"
        ]
    },
    "id": 1
}
```

Change proxy status

Change the proxy to an active proxy and rename it to ”Active proxy”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxy.update",
    "params": {
        "proxyid": "10293",
        "name": "Active proxy",
        "operating_mode": "0"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "proxyids": [
            "10293"
        ]
    },
    "id": 1
}
```

Add a proxy to a proxy group

Update proxy with ID ”5” and add it to proxy group with ID ”1”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxy.create",
    "params": {
        "proxyid": "5",
        "proxy_groupid": "1",
        "local_address": "127.0.0.1"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "proxyids": [
            "5"
        ]
    },
    "id": 1
}
```

See also

- Host
- Proxy group

Source

CProxy::update() in *ui/include/classes/api/services/CProxy.php*.

Proxy group

This class is designed to work with proxy groups.

Object references:

- Proxy group

Available methods:

- proxygroup.create - create new proxy groups
- proxygroup.delete - delete proxy groups
- proxygroup.get - retrieve proxy groups
- proxygroup.update - update proxy groups

Proxy group object

The following objects are directly related to the `proxygroup` API.

Proxy group

The proxy group object has the following properties.

Property	Type	Description
proxy_groupid	ID	ID of the proxy group. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the proxy group. Property behavior: - *required* for create operations
description	text	Description of the proxy group.
failover_delay	string	Period during which a proxy in the proxy group must communicate with Zabbix server to be considered online. Time suffixes are supported (e.g., `30s`, `1m`). User macros are supported. Possible values: 10s-15m. Default: 1m.
min_online	string	Minimum number of online proxies required to keep the proxy group online. User macros are supported. Possible values range: 1-1000. Default: 1.

Property	Type	Description
state	integer	State of the proxy group. Possible values: 0 - Unknown; 1 - Offline; 2 - Recovering; 3 - Online; 4 - Degrading. Property behavior: - *read-only*

proxygroup.create

Description

`object proxygroup.create(object/array proxyGroups)`

This method allows you to create new proxy groups.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Proxy groups to create.

The method accepts proxy groups with the standard proxy group properties.

Return values

`(object)` Returns an object containing the IDs of the created proxy groups under the `proxy_groupids` property. The order of the returned IDs matches the order of the passed proxy groups.

Examples

Create a proxy group

Create a proxy group with custom settings.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxygroup.create",
    "params": {
        "name": "Proxy group",
        "failover_delay": "5m",
        "min_online": "10"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "proxy_groupids": [
            "5"
        ]
    },
    "id": 1
}
```

Source

CProxyGroup::create() in *ui/include/classes/api/services/CProxyGroup.php*.

proxygroup.delete

Description

`object proxygroup.delete(array proxyGroupIds)`

This method allows you to delete proxy groups.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of proxy groups to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted proxy groups under the `proxy_groupids` property.

Examples

Delete multiple proxy groups

Delete two proxy groups.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxygroup.delete",
    "params": [
        "5",
        "10"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "proxy_groupids": [
            "5",
            "10"
        ]
    },
    "id": 1
}
```

Source

CProxyGroup::delete() in *ui/include/classes/api/services/CProxyGroup.php*.

proxygroup.get

Description

`integer/array proxygroup.get(object parameters)`

The method allows you to retrieve proxy groups according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information. Only proxy groups allowed by the *Proxy access list* of the user's user groups are returned.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
proxy_groupids	ID/array	Return only proxy groups with the given IDs.
proxyids	ID/array	Return only proxy groups that contain the given proxies.
selectProxies	query	Return a `proxies` property with the proxies that belong to the proxy group. Supports `count`.
sortfield	string/array	Sort the result by the given properties. Possible values: `proxy_groupid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve all proxy groups

Retrieve all configured proxy groups with proxies.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxygroup.get",
    "params": {
        "output": "extend",
        "selectProxies": ["proxyid", "name"]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "proxy_groupid": "1",
            "name": "Proxy group 1",
```

```
            "failover_delay": "1m",
            "min_online": "3",
            "description": "",
            "state": "1",
            "proxies": [
                {
                    "proxyid": "1",
                    "name": "proxy 1"
                },
                {
                    "proxyid": "2",
                    "name": "proxy 2"
                }
            ]
        },
        {
            "proxy_groupid": "2",
            "name": "Proxy group 2",
            "failover_delay": "10m",
            "min_online": "3",
            "description": "",
            "state": "3",
            "proxies": [
                {
                    "proxyid": "3",
                    "name": "proxy 3"
                },
                {
                    "proxyid": "4",
                    "name": "proxy 4"
                },
                {
                    "proxyid": "5",
                    "name": "proxy 5"
                }
            ]
        }
    ],
    "id": 1
}
```


See also

- Proxy

Source

CProxyGroup::get() in *ui/include/classes/api/services/CProxyGroup.php*.

proxygroup.update

Description

`object proxygroup.update(object/array proxyGroups)`

This method allows you to update existing proxy groups.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Proxy group properties to be updated.

The `proxy_groupid` property must be defined for each proxy group, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

The method accepts proxy groups with the standard proxy group properties.

Return values

`(object)` Returns an object containing the IDs of the updated proxy groups under the `proxy_groupids` property.

Examples

Change minimum number of online proxies

Change minimum number of online proxies required for the group to be online.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "proxygroup.update",
    "params": {
        "proxy_groupid": "5",
        "min_online": "3"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "proxy_groupids": [
            "5"
        ]
    },
    "id": 1
}
```

Source

CProxyGroup::update() in *ui/include/classes/api/services/CProxyGroup.php*.

Regular expression

This class is designed to work with global regular expressions.

Object references:

- Regular expression
- Expressions

Available methods:

- regexp.create - create new regular expressions
- regexp.delete - delete regular expressions
- regexp.get - retrieve regular expressions
- regexp.update - update regular expressions

Regular expression object

The following objects are directly related to the `regexp` API.

Regular expression

The global regular expression object has the following properties.

Property	Type	Description
regexpid	ID	ID of the regular expression. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the regular expression. Property behavior: - *required* for create operations
description	text	Description of the regular expression
test_string	string	Test string

Expressions

The expressions object has the following properties.

Property	Type	Description
expression	string	Regular expression. Property behavior: - *required*
expression_type	integer	Type of Regular expression. Possible values: 0 - Contains string; 1 - Contains any substring from list; 2 - Does not contain string; 3 - Matches regular expression; 4 - Does not match regular expression. Property behavior: - *required*
exp_delimiter	string	Expression delimiter. Default value: `","`. Possible values: `","` or `"."`, or `"/"`. Property behavior: - *supported* if `expression_type` is set to *Contains any substring from list*.
case_sensitive	integer	Case sensitivity. Default value: 0. Possible values: 0 - Case insensitive; 1 - Case sensitive.

regexp.create

Description

`object regexp.create(object/array regularExpressions)`

This method allows you to create new global regular expressions.

Note:
This method is only available to *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Regular expressions to create.

Additionally to the standard properties, the method accepts the following parameters.

Parameter	Type	Description
expressions	array	Expressions options. Parameter behavior: *- required*

Return values

`(object)` Returns an object containing the IDs of the created regular expressions under the `regexpids` property.

Examples

Creating a new global regular expression.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "regexp.create",
    "params": {
        "name": "Storage devices for SNMP discovery",
        "description": "Global regular expression used to discover and identify memory resources on the mo
        "test_string": "/boot",
        "expressions": [
            {
                "expression": "^(Physical memory|Virtual memory|Memory buffers|Cached memory|Swap space)$"
                "expression_type": "4",
                "case_sensitive": "1"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "regexpids": [
            "9"
        ]
    },
    "id": 1
}
```

Source

CRegexp::create() in *ui/include/classes/api/services/CRegexp.php*.

regexp.delete

Description

`object regexp.delete(array regexpids)`

This method allows you to delete global regular expressions.

> **Note:**
> This method is only available to *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the regular expressions to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted regular expressions under the `regexpids` property.

Examples

Deleting multiple global regular expressions.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "regexp.delete",
    "params": [
        "16",
        "17"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "regexpids": [
            "16",
            "17"
        ]
    },
    "id": 1
}
```

Source

CRegexp::delete() in *ui/include/classes/api/services/CRegexp.php*.

regexp.get

Description

`integer/array regexp.get(object parameters)`

The method allows you to retrieve global regular expressions according to the given parameters.

> **Note:**
> This method is available only to *Super Admin*. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
regexpids	ID/array	Return only regular expressions with the given IDs.
selectExpressions	query	Return an `expressions` property.
sortfield	string/array	Sort the result by the given properties. Possible values: `regexpid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	

Parameter	Type	Description
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving global regular expressions.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "regexp.get",
    "params": {
        "output": ["regexpid","name","description"],
        "selectExpressions": ["expression","expression_type"],
        "regexpids": [1,2],
        "preservekeys": true
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "1": {
            "regexpid": "1",
            "name": "File systems for discovery",
            "description": "",
            "expressions": [
                {
                    "expression": "^(btrfs|ext2|ext3|ext4|reiser|xfs|ffs|ufs|jfs|jfs2|vxfs|hfs|apfs|refs|n
                    "expression_type": "3"
                }
            ]
        },
        "2": {
            "regexpid": "2",
            "name": "Network interfaces for discovery",
            "description": "",
            "expressions": [
                {
                    "expression": "^Software Loopback Interface",
                    "expression_type": "4"
                },
                {
                    "expression": "^(In)?[Ll]oop[Bb]ack[0-9._]*$",
                    "expression_type": "4"
                },
                {
```

```
                    "expression": "^NULL[0-9.]*$",
                    "expression_type": "4"
                },
                {
                    "expression": "^[Ll]o[0-9.]*$",
                    "expression_type": "4"
                },
                {
                    "expression": "^[Ss]ystem$",
                    "expression_type": "4"
                },
                {
                    "expression": "^Nu[0-9.]*$",
                    "expression_type": "4"
                }
            ]
        }
    },
    "id": 1
}
```


Source

CRegexp::get() in *ui/include/classes/api/services/CRegexp.php*.

regexp.update

Description

`object regexp.update(object/array regularExpressions)`

This method allows you to update existing global regular expressions.

> **Note:**
> This method is only available to *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Regular expression properties to be updated.

The `regexpid` property must be defined for each object, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard properties, the method accepts the following parameters.

Parameter	Type	Description
expressions	array	Expressions options.

Return values

`(object)` Returns an object containing the IDs of the updated regular expressions under the `regexpids` property.

Examples

Updating global regular expression for file systems discovery.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "regexp.update",
    "params": {
        "regexpid": "1",
        "name": "File systems for discovery",
        "description": "Global regular expression used to discover and identify file systems on the monito
```

```
        "test_string": "",
        "expressions": [
            {
                "expression": "^(btrfs|ext2|ext3|ext4|reiser|xfs|ffs|ufs|jfs|jfs2|vxfs|hfs|apfs|refs|zfs)$
                "expression_type": "3",
                "case_sensitive": "0"
            },
            {
                "expression": "^(ntfs|fat32|fat16)$",
                "expression_type": "3",
                "case_sensitive": "0"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "regexpids": [
            "1"
        ]
    },
    "id": 1
}
```

Source

CRegexp::update() in *ui/include/classes/api/services/CRegexp.php*.

Report

This class is designed to work with scheduled reports.

Object references:

- Report
- Users
- User groups

Available methods:

- report.create - create new scheduled reports
- report.delete - delete scheduled reports
- report.get - retrieve scheduled reports
- report.update - update scheduled reports

Report object

The following objects are directly related to the `report` API.

Report

The report object has the following properties:

Property	Type	Description
reportid	ID	ID of the report. Property behavior: - *read-only* - *required* for update operations

Property	Type	Description
userid	ID	ID of the user who created the report. Property behavior: - *required* for create operations
name	string	Unique name of the report. Property behavior: - *required* for create operations
dashboardid	ID	ID of the dashboard that the report is based on. Property behavior: - *required* for create operations
period	integer	Period for which the report will be prepared. Possible values: 0 - *(default)* previous day; 1 - previous week; 2 - previous month; 3 - previous year.
cycle	integer	Period repeating schedule. Possible values: 0 - *(default)* daily; 1 - weekly; 2 - monthly; 3 - yearly.
start_time	integer	Time of the day, in seconds, when the report will be prepared for sending. Default: 0.
weekdays	integer	Days of the week for sending the report. Possible bitmap values: 1 - Monday; 2 - Tuesday; 4 - Wednesday; 8 - Thursday; 16 - Friday; 32 - Saturday; 64 - Sunday. This is a bitmask field; any sum of possible bitmap values is acceptable (for example, 21 for Monday, Wednesday, and Friday). Default: 0. Property behavior: - *required* if `cycle` is set to ”weekly”.
active_since	string	On which date to start. Possible values: empty string - *(default)* not specified (stored as 0); specific date in YYYY-MM-DD format (stored as a timestamp of the beginning of a day (00:00:00)).
active_till	string	On which date to end. Possible values: empty string - *(default)* not specified (stored as 0); specific date in YYYY-MM-DD format (stored as a timestamp of the end of a day (23:59:59)).
subject	string	Report message subject.

Property	Type	Description
message	string	Report message text.
status	integer	Whether the report is enabled or disabled. Possible values: 0 - *(default)* Enabled; 1 - Disabled.
description	text	Description of the report.
state	integer	State of the report. Possible values: 0 - *(default)* report was not yet processed; 1 - report was generated and successfully sent to all recipients; 2 - report generating failed; ”info” contains error information; 3 - report was generated, but sending to some (or all) recipients failed; ”info” contains error information. Property behavior: - *read-only*
lastsent	timestamp	Unix timestamp of the last successfully sent report. Property behavior: - *read-only*
info	string	Error description or additional information. Property behavior: - *read-only*

Users

The users object has the following properties:

Property	Type	Description
userid	ID	ID of user to send the report to. Property behavior: - *required*
access_userid	ID	ID of user on whose behalf the report will be generated. 0 - *(default)* Generate report by recipient.
exclude	integer	Whether to exclude the user from mailing list. Possible values: 0 - *(default)* Include; 1 - Exclude.

User groups

The user groups object has the following properties:

Property	Type	Description
usrgrpid	ID	ID of user group to send the report to. Property behavior: - *required*
access_userid	ID	ID of user on whose behalf the report will be generated. 0 - *(default)* Generate report by recipient.

report.create

Description

`object report.create(object/array reports)`

This method allows you to create new scheduled reports.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Scheduled reports to create.

Additionally to the standard scheduled report properties, the method accepts the following parameters.

Parameter	Type	Description
users	object/array	Users to send the report to. Parameter behavior: - *required* if `user_groups` is not set
user_groups	object/array	User groups to send the report to. Parameter behavior: - *required* if `users` is not set

Return values

`(object)` Returns an object containing the IDs of the created scheduled reports under the `reportids` property. The order of the returned IDs matches the order of the passed scheduled reports.

Examples

Creating a scheduled report

Create a weekly report that will be prepared for the previous week every Monday-Friday at 12:00 from 2021-04-01 to 2021-08-31.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "report.create",
    "params": {
        "userid": "1",
        "name": "Weekly report",
        "dashboardid": "1",
        "period": "1",
        "cycle": "1",
        "start_time": "43200",
        "weekdays": "31",
        "active_since": "2021-04-01",
        "active_till": "2021-08-31",
        "subject": "Weekly report",
        "message": "Report accompanying text",
        "status": "1",
        "description": "Report description",
        "users": [
            {
                "userid": "1",
                "access_userid": "1",
                "exclude": "0"
            },
            {
                "userid": "2",
```

```
                "access_userid": "0",
                "exclude": "1"
            }
        ],
        "user_groups": [
            {
                "usrgrpid": "7",
                "access_userid": "0"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "reportids": [
            "1"
        ]
    },
    "id": 1
}
```

See also

- Users
- User groups

Source

CReport::create() in *ui/include/classes/api/services/CReport.php*.

report.delete

Description

`object report.delete(array reportids)`

This method allows you to delete scheduled reports.

> **Note:**
> This method is only available to *Admin* and *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the scheduled reports to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted scheduled reports under the `reportids` property.

Examples

Deleting multiple scheduled reports

Delete two scheduled reports.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "report.delete",
    "params": [
        "1",
        "2"
```

```
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "reportids": [
            "1",
            "2"
        ]
    },
    "id": 1
}
```

Source

CReport::delete() in *ui/include/classes/api/services/CReport.php*.

report.get

Description

`integer/array report.get(object parameters)`

The method allows you to retrieve scheduled reports according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
reportids	ID/array	Return only scheduled reports with the given report IDs.
expired	boolean	If set to `true` return only expired scheduled reports.
selectUsers	query	Return a `users` property with users the report is configured to be sent to.
selectUserGroups	query	Return a `user_groups` property with user groups the report is configured to be sent to.
sortfield	string/array	Sort the result by the given properties. Possible values: `reportid`, `name`, `status`.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving report data

Request:

```
{
    "jsonrpc": "2.0",
    "method": "report.get",
    "params": {
        "output": "extend",
        "selectUsers": "extend",
        "selectUserGroups": "extend",
        "reportids": ["1", "2"]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "reportid": "1",
            "userid": "1",
            "name": "Weekly report",
            "dashboardid": "1",
            "period": "1",
            "cycle": "1",
            "start_time": "43200",
            "weekdays": "31",
            "active_since": "2021-04-01",
            "active_till": "2021-08-31",
            "subject": "Weekly report",
            "message": "Report accompanying text",
            "status": "1",
            "description": "Report description",
            "state": "1",
            "lastsent": "1613563219",
            "info": "",
            "users": [
                {
                    "userid": "1",
                    "access_userid": "1",
                    "exclude": "0"
                },
                {
                    "userid": "2",
                    "access_userid": "0",
                    "exclude": "1"
                }
            ],
            "user_groups": [
                {
                    "usrgrpid": "7",
                    "access_userid": "0"
                }
            ]
        },
        {
            "reportid": "2",
            "userid": "1",
```

```
            "name": "Monthly report",
            "dashboardid": "2",
            "period": "2",
            "cycle": "2",
            "start_time": "0",
            "weekdays": "0",
            "active_since": "2021-05-01",
            "active_till": "",
            "subject": "Monthly report",
            "message": "Report accompanying text",
            "status": "1",
            "description": "",
            "state": "0",
            "lastsent": "0",
            "info": "",
            "users": [
                {
                    "userid": "1",
                    "access_userid": "1",
                    "exclude": "0"
                }
            ],
            "user_groups": []
        }
    ],
    "id": 1
}
```


See also

- Users
- User groups

Source

CReport::get() in *ui/include/classes/api/services/CReport.php*.

report.update

Description

`object report.update(object/array reports)`

This method allows you to update existing scheduled reports.

> **Note:**
> This method is only available to *Admin* and *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Scheduled report properties to be updated.

The `reportid` property must be defined for each scheduled report, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard scheduled report properties the method accepts the following parameters.

Parameter	Type	Description
users	object/array	Users to replace the current users assigned to the scheduled report. Parameter behavior: - *required* if `user_groups` is not set

Parameter	Type	Description
user_groups	object/array	User groups to replace the current user groups assigned to the scheduled report. Parameter behavior: - *required* if `users` is not set

Return values

`(object)` Returns an object containing the IDs of the updated scheduled reports under the `reportids` property.

Examples

Disabling scheduled report

Request:

```
{
    "jsonrpc": "2.0",
    "method": "report.update",
    "params": {
        "reportid": "1",
        "status": "0"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "reportids": [
            "1"
        ]
    },
    "id": 1
}
```

See also

- Users
- User groups

Source

CReport::update() in *ui/include/classes/api/services/CReport.php*.

Role

This class is designed to work with user roles.

Object references:

- Role
- Role rules
  - UI element
  - Service
  - Service tag
  - Module
  - Action

Available methods:

- role.create - create new user roles
- role.delete - delete user roles
- role.get - retrieve user roles
- role.update - update user roles

Role object

The following objects are directly related to the `role` API.

Role

The role object has the following properties:

Property	Type	Description
roleid	ID	ID of the role. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the role. Property behavior: - *required* for create operations
type	integer	User type. Possible values: 1 - *(default)* User; 2 - Admin; 3 - Super admin. Property behavior: - *required* for create operations
readonly	integer	Whether the role is readonly. Possible values: 0 - *(default)* No; 1 - Yes. Property behavior: - *read-only*

Role rules

The role rules object has the following properties:

Property	Type	Description
ui	array	Array of the UI element objects.
ui.default_access	integer	Whether access to new UI elements is enabled. Possible values: 0 - Disabled; 1 - *(default)* Enabled.
profile.redirect.enforce	integer	Whether role-level redirect URL after authentication is enforced and cannot be overridden by the user. Possible values: 0 - *(default)* Disabled; 1 - Enabled.
profile.redirect.url	string	URL used for user redirection after authentication.
services.read.mode	integer	Read-only access to services. Possible values: 0 - Read-only access to the services, specified by the `services.read.list` or matched by the `services.read.tag` properties; 1 - *(default)* Read-only access to all services.

Property	Type	Description
services.read.list	array	Array of Service objects. The specified services, including child services, will be granted a read-only access to the user role. Read-only access will not override read-write access to the services. Property behavior: - *supported* if `services.read.mode` is set to ”0”.
services.read.tag	object	Array of Service tag objects. The tag matched services, including child services, will be granted a read-only access to the user role. Read-only access will not override read-write access to the services. Property behavior: - *supported* if `services.read.mode` is set to ”0”.
services.write.mode	integer	Read-write access to services. Possible values: 0 - *(default)* Read-write access to the services, specified by the `services.write.list` or matched by the `services.write.tag` properties; 1 - Read-write access to all services.
services.write.list	array	Array of Service objects. The specified services, including child services, will be granted a read-write access to the user role. Read-write access will override read-only access to the services. Property behavior: - *supported* if `services.write.mode` is set to ”0”.
services.write.tag	object	Array of Service tag objects. The tag matched services, including child services, will be granted a read-write access to the user role. Read-write access will override read-only access to the services. Property behavior: - *supported* if `services.write.mode` is set to ”0”.
modules	array	Array of the module objects.
modules.default_access	integer	Whether access to new modules is enabled. Possible values: 0 - Disabled; 1 - *(default)* Enabled.
api.access	integer	Whether access to API is enabled. Possible values: 0 - *(default)* Disabled; 1 - Enabled.
api.mode	integer	Mode for treating API methods listed in the `api` property. Possible values: 0 - *(default)* Deny list; 1 - Allow list. Property behavior: - *supported* if `api.access` is set to ”1”.

Property	Type	Description
api	array	Array of API methods. Wildcards are supported, for example: `*` (any method), `dashboard.*`, `*.export`. If no methods are specified, access to all API methods is denied if `api.mode` is set to "1" (Allow list), and allowed if `api.mode` is set to "0" (Deny list). Property behavior: - *supported* if `api.access` is set to "1".
actions	array	Array of the action objects.
actions.default_access	integer	Whether access to new actions is enabled. Possible values: 0 - Disabled; 1 - *(default)* Enabled.
devices.access	integer	Whether access to user's own mobile devices for use with Zabbix mobile app is enabled. Possible values: 0 - (*default for User and Admin*) Disabled; 1 - (*default for Super admin*) Enabled. Property behavior: - *supported* if `EnableMobileDevices` is enabled
devices.actions	array	Array of Device action objects. Property behavior: - *supported* if `devices.access` is set to "1"
devices.actions.default_access	integer	Whether access to new device actions is enabled. Possible values: 0 - (*default for User and Admin*) Disabled; 1 - (*default for Super admin*) Enabled. Property behavior: - *supported* if `EnableMobileDevices` is enabled

UI element

The UI element object has the following properties:

Property	Type	Description
name	string	Name of the UI element. Possible values if `type` of the Role object is set to ”User”, ”Admin”, or ”Super admin”: `monitoring.dashboard` - *Dashboards*; `monitoring.problems` - *Monitoring > Problems*; `monitoring.hosts` - *Monitoring > Hosts*; `monitoring.latest_data` - *Monitoring > Latest data*; `monitoring.maps` - *Monitoring > Maps*; `services.services` - *Services > Services*; `services.sla_report` - *Services > SLA report*; `inventory.overview` - *Inventory > Overview*; `inventory.hosts` - *Inventory > Hosts*; `reports.availability_report` - *Reports > Availability report*; `reports.top_triggers` - *Reports > Triggers top 100*. Possible values if `type` of the Role object is set to ”Admin” or ”Super admin”: `monitoring.discovery` - *Monitoring > Discovery*; `services.sla` - *Services > SLA*; `reports.scheduled_reports` - *Reports > Scheduled reports*; `reports.notifications` - *Reports > Notifications*; `configuration.template_groups` - *Data collection > Template groups*; `configuration.host_groups` - *Data collection > Host groups*; `configuration.templates` - *Data collection > Templates*; `configuration.hosts` - *Data collection > Hosts*; `configuration.maintenance` - *Data collection > Maintenance*; `configuration.event_processing` - *Data collection > Event processing*; `configuration.discovery` - *Data collection > Discovery*; `configuration.trigger_actions` - *Alerts > Actions > Trigger actions*; `configuration.service_actions` - *Alerts > Actions > Service actions*; `configuration.discovery_actions` - *Alerts > Actions > Discovery actions*; `configuration.autoregistration_actions` - *Alerts > Actions > Autoregistration actions*; `configuration.internal_actions` - *Alerts > Actions > Internal actions*. Possible values if `type` of the Role object is set to ”Super admin”: `reports.system_info` - *Reports > System information*; `reports.audit` - *Reports > Audit log*; `reports.action_log` - *Reports > Action log*; `administration.media_types` - *Alerts > Media types*; `administration.scripts` - *Alerts > Scripts*; `administration.user_groups` - *Users > User groups*; `administration.user_roles` - *Users > User roles*; `administration.users` - *Users > Users*; `administration.api_tokens` - *Users > API tokens*; `administration.authentication` - *Users > Authentication*; `administration.general` - *Administration > General*; `administration.audit_log` - *Administration > Audit log*; `administration.housekeeping` - *Administration > Housekeeping*; `administration.proxy_groups` - *Administration > Proxy groups*; `administration.proxies` - *Administration > Proxies*; `administration.macros` - *Administration > Macros*; `administration.queue` - *Administration > Queue*. Property behavior: *- required*

Property	Type	Description
name	string	Name of the UI element. Possible values if `type` of the Role object is set to ”User”, ”Admin”, or ”Super admin”: `monitoring.dashboard` - *Dashboards*; `monitoring.problems` - *Monitoring > Problems*; `monitoring.hosts` - *Monitoring > Hosts*; `monitoring.latest_data` - *Monitoring > Latest data*; `monitoring.maps` - *Monitoring > Maps*; `services.services` - *Services > Services*; `services.sla_report` - *Services > SLA report*; `inventory.overview` - *Inventory > Overview*; `inventory.hosts` - *Inventory > Hosts*; `reports.availability_report` - *Reports > Availability report*; `reports.top_triggers` - *Reports > Triggers top 100*. Possible values if `type` of the Role object is set to ”Admin” or ”Super admin”: `monitoring.discovery` - *Monitoring > Discovery*; `services.sla` - *Services > SLA*; `reports.scheduled_reports` - *Reports > Scheduled reports*; `reports.notifications` - *Reports > Notifications*; `configuration.template_groups` - *Data collection > Template groups*; `configuration.host_groups` - *Data collection > Host groups*; `configuration.templates` - *Data collection > Templates*; `configuration.hosts` - *Data collection > Hosts*; `configuration.maintenance` - *Data collection > Maintenance*; `configuration.discovery` - *Data collection > Discovery*; `configuration.trigger_actions` - *Alerts > Actions > Trigger actions*; `configuration.service_actions` - *Alerts > Actions > Service actions*; `configuration.discovery_actions` - *Alerts > Actions > Discovery actions*; `configuration.autoregistration_actions` - *Alerts > Actions > Autoregistration actions*; `configuration.internal_actions` - *Alerts > Actions > Internal actions*. Possible values if `type` of the Role object is set to ”Super admin”: `reports.system_info` - *Reports > System information*; `reports.audit` - *Reports > Audit log*; `reports.action_log` - *Reports > Action log*; `configuration.event_correlation` - *Data collection > Event correlation*; `administration.media_types` - *Alerts > Media types*; `administration.scripts` - *Alerts > Scripts*; `administration.user_groups` - *Users > User groups*; `administration.user_roles` - *Users > User roles*; `administration.users` - *Users > Users*; `administration.api_tokens` - *Users > API tokens*; `administration.authentication` - *Users > Authentication*; `administration.linked_devices` - *Users > Devices*; `administration.general` - *Administration > General*; `administration.audit_log` - *Administration > Audit log*; `administration.housekeeping` - *Administration > Housekeeping*; `administration.proxy_groups` - *Administration > Proxy groups*; `administration.proxies` - *Administration > Proxies*; `administration.macros` - *Administration > Macros*; `administration.queue` - *Administration > Queue*. Property behavior: - *required*

Property	Type	Description
status	integer	Whether access to the UI element is enabled. Possible values: 0 - Disabled; 1 - *(default)* Enabled.

Service

Property	Type	Description
serviceid	ID	ID of the Service. Property behavior: - *required*.

Service tag

Property	Type	Description
tag	string	Tag name. If empty string is specified, the service tag will not be used for service matching. Property behavior: - *required*.
value	string	Tag value. If no value or empty string is specified, only the tag name will be used for service matching.

Module

The module object has the following properties:

Property	Type	Description
moduleid	ID	ID of the module. Property behavior: - *required*.
status	integer	Whether access to the module is enabled. Possible values: 0 - Disabled; 1 - *(default)* Enabled.

Action

The action object has the following properties:

Property	Type	Description
name	string	Name of the action. Possible values if `type` of the Role object is set to ”User”, ”Admin”, or ”Super admin: `edit_dashboards` - Create and edit dashboards; `edit_maps` - Create and edit maps; `add_problem_comments` - Add problem comments; `change_severity` - Change problem severity; `acknowledge_problems` - Acknowledge problems; `suppress_problems` - Suppress problems; `close_problems` - Close problems; `execute_scripts` - Execute scripts; `manage_api_tokens` - Manage API tokens; `change_problem_ranking` - Change the problem ranking from cause to symptom, and vice versa; `edit_own_media` - Allow users to create and edit their own media. Possible values if `type` of the Role object is set to ”Admin” or ”Super admin”: `edit_maintenance` - Create and edit maintenances; `manage_scheduled_reports` - Manage scheduled reports, `manage_sla` - Manage SLA, `select_server_for_monitoring` - Allow users to create hosts monitored by the server and network discovery rules executed by the server. Possible values if `type` of the Role object is set to ”User” or ”Admin”: `invoke_execute_now` - allows users with only read permissions on a host to execute item checks. Possible values if `type` of the Role object is set to ”Super admin”: `edit_user_media` - Allow users to create and edit media for other users. Property behavior: - *required*.
status	integer	Whether access to perform the action is enabled. Possible values: 0 - Disabled; 1 - *(default)* Enabled.

Device action

The device action object has the following properties:

Property	Type	Description
name	string	Name of the device action. Possible values if `type` of the Role object is set to ”User”, ”Admin”, or ”Super admin: `manage_own` - Allow users to add or remove their own mobile devices. Possible values if `type` of the Role object is set to ”Super admin”: `manage_user` - Allow users to add or remove mobile devices for all users. Property behavior: - *required*

Property	Type	Description
status	integer	Whether access to perform the action is enabled. Possible values: 0 - Disabled; 1 - (*default*) Enabled.

role.create

Description

`object role.create(object/array roles)`

This method allows you to create new roles.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Roles to create.

Additionally to the standard role properties, the method accepts the following parameters.

Parameter	Type	Description
rules	array	Role rules to be created for the role.

Return values

`(object)` Returns an object containing the IDs of the created roles under the `roleids` property. The order of the returned IDs matches the order of the passed roles.

Examples

Creating a role

Create a role with type "User" and denied access to two UI elements.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "role.create",
    "params": {
        "name": "Operator",
        "type": "1",
        "rules": {
            "ui": [
                {
                    "name": "monitoring.hosts",
                    "status": "0"
                },
                {
                    "name": "monitoring.maps",
                    "status": "0"
                }
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "roleids": [
            "5"
        ]
    },
    "id": 1
}
```

See also

- Role rules
- UI element
- Module
- Action

Source

CRole::create() in *ui/include/classes/api/services/CRole.php*.

role.delete

Description

`object role.delete(array roleids)`

This method allows you to delete roles.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the roles to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted roles under the `roleids` property.

Examples

Deleting multiple user roles

Delete two user roles.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "role.delete",
    "params": [
        "4",
        "5"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "roleids": [
            "4",
            "5"
        ]
    },
```

```
    "id": 1
}
```

Source

CRole::delete() in *ui/include/classes/api/services/CRole.php*.

role.get

Description

`integer/array role.get(object parameters)`

The method allows you to retrieve roles according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
roleids	ID/array	Return only roles with the given IDs.
current	boolean	If set to `true`, returns only the current user's role.
selectRules	query	Return a `rules` property with the role rules.
selectUsers	query	Return a `users` property with the users that the role is assigned to. See `user.get` for restrictions based on user type.
sortfield	string/array	Sort the result by the given properties. Possible values: `roleid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving role data

Retrieve "Super admin role" role data and its access rules.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "role.get",
    "params": {
```

```
        "output": "extend",
        "selectRules": "extend",
        "roleids": "3"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "roleid": "3",
            "name": "Super admin role",
            "type": "3",
            "readonly": "1",
            "rules": {
                "ui": [
                    {
                        "name": "monitoring.dashboard",
                        "status": "1"
                    },
                    {
                        "name": "monitoring.problems",
                        "status": "1"
                    },
                    {
                        "name": "monitoring.hosts",
                        "status": "1"
                    },
                    {
                        "name": "monitoring.latest_data",
                        "status": "1"
                    },
                    {
                        "name": "monitoring.maps",
                        "status": "1"
                    },
                    {
                        "name": "services.services",
                        "status": "1"
                    },
                    {
                        "name": "services.sla_report",
                        "status": "1"
                    },
                    {
                        "name": "inventory.overview",
                        "status": "1"
                    },
                    {
                        "name": "inventory.hosts",
                        "status": "1"
                    },
                    {
                        "name": "reports.availability_report",
                        "status": "1"
                    },
                    {
                        "name": "reports.top_triggers",
                        "status": "1"
                    },
```

```
                {
                    "name": "monitoring.discovery",
                    "status": "1"
                },
                {
                    "name": "services.sla",
                    "status": "1"
                },
                {
                    "name": "reports.scheduled_reports",
                    "status": "1"
                },
                {
                    "name": "reports.notifications",
                    "status": "1"
                },
                {
                    "name": "configuration.template_groups",
                    "status": "1"
                },
                {
                    "name": "configuration.host_groups",
                    "status": "1"
                },
                {
                    "name": "configuration.templates",
                    "status": "1"
                },
                {
                    "name": "configuration.hosts",
                    "status": "1"
                },
                {
                    "name": "configuration.maintenance",
                    "status": "1"
                },
                {
                    "name": "configuration.discovery",
                    "status": "1"
                },
                {
                    "name": "configuration.trigger_actions",
                    "status": "1"
                },
                {
                    "name": "configuration.service_actions",
                    "status": "1"
                },
                {
                    "name": "configuration.discovery_actions",
                    "status": "1"
                },
                {
                    "name": "configuration.autoregistration_actions",
                    "status": "1"
                },
                {
                    "name": "configuration.internal_actions",
                    "status": "1"
                },
                {
                    "name": "reports.system_info",
```

```
            "status": "1"
        },
        {
            "name": "reports.audit",
            "status": "1"
        },
        {
            "name": "reports.action_log",
            "status": "1"
        },
        {
            "name": "configuration.event_correlation",
            "status": "1"
        },
        {
            "name": "administration.media_types",
            "status": "1"
        },
        {
            "name": "administration.scripts",
            "status": "1"
        },
        {
            "name": "administration.user_groups",
            "status": "1"
        },
        {
            "name": "administration.user_roles",
            "status": "1"
        },
        {
            "name": "administration.users",
            "status": "1"
        },
        {
            "name": "administration.api_tokens",
            "status": "1"
        },
        {
            "name": "administration.authentication",
            "status": "1"
        },
        {
            "name": "administration.general",
            "status": "1"
        },
        {
            "name": "administration.audit_log",
            "status": "1"
        },
        {
            "name": "administration.housekeeping",
            "status": "1"
        },
        {
            "name": "administration.proxy_groups",
            "status": "1"
        },
        {
            "name": "administration.proxies",
            "status": "1"
        },
```

```
                {
                    "name": "administration.macros",
                    "status": "1"
                },
                {
                    "name": "administration.queue",
                    "status": "1"
                }
            ],
            "ui.default_access": "1",
            "profile.redirect.enforce": "0",
            "profile.redirect.url": "",
            "services.read.mode": "1",
            "services.read.list": [],
            "services.read.tag": {
                "tag": "",
                "value": ""
            },
            "services.write.mode": "1",
            "services.write.list": [],
            "services.write.tag": {
                "tag": "",
                "value": ""
            },
            "modules": [
                {
                    "moduleid": 1,
                    "status": "1"
                },
                {
                    "moduleid": 2,
                    "status": "1"
                },
                {
                    "moduleid": 3,
                    "status": "1"
                },
                {
                    "moduleid": 4,
                    "status": "1"
                },
                {
                    "moduleid": 5,
                    "status": "1"
                },
                {
                    "moduleid": 6,
                    "status": "1"
                },
                {
                    "moduleid": 7,
                    "status": "1"
                },
                {
                    "moduleid": 8,
                    "status": "1"
                },
                {
                    "moduleid": 9,
                    "status": "1"
                },
                {
```

```
            "moduleid": 10,
            "status": "1"
        },
        {
            "moduleid": 11,
            "status": "1"
        },
        {
            "moduleid": 12,
            "status": "1"
        },
        {
            "moduleid": 13,
            "status": "1"
        },
        {
            "moduleid": 14,
            "status": "1"
        },
        {
            "moduleid": 15,
            "status": "1"
        },
        {
            "moduleid": 16,
            "status": "1"
        },
        {
            "moduleid": 17,
            "status": "1"
        },
        {
            "moduleid": 18,
            "status": "1"
        },
        {
            "moduleid": 19,
            "status": "1"
        },
        {
            "moduleid": 20,
            "status": "1"
        },
        {
            "moduleid": 21,
            "status": "1"
        },
        {
            "moduleid": 22,
            "status": "1"
        },
        {
            "moduleid": 23,
            "status": "1"
        },
        {
            "moduleid": 24,
            "status": "1"
        },
        {
            "moduleid": 25,
            "status": "1"
```

```
            },
            {
                "moduleid": 26,
                "status": "1"
            },
            {
                "moduleid": 27,
                "status": "1"
            },
            {
                "moduleid": 28,
                "status": "1"
            },
            {
                "moduleid": 29,
                "status": "1"
            },
            {
                "moduleid": 30,
                "status": "1"
            },
            {
                "moduleid": 31,
                "status": "1"
            },
            {
                "moduleid": 32,
                "status": "1"
            },
            {
                "moduleid": 33,
                "status": "1"
            }
        ],
        "modules.default_access": "1",
        "api.access": "1",
        "api.mode": "0",
        "api": [],
        "actions": [
            {
                "name": "edit_dashboards",
                "status": "1"
            },
            {
                "name": "edit_maps",
                "status": "1"
            },
            {
                "name": "acknowledge_problems",
                "status": "1"
            },
            {
                "name": "suppress_problems",
                "status": "1"
            },
            {
                "name": "close_problems",
                "status": "1"
            },
            {
                "name": "change_severity",
                "status": "1"
```

```
                    },
                    {
                        "name": "add_problem_comments",
                        "status": "1"
                    },
                    {
                        "name": "execute_scripts",
                        "status": "1"
                    },
                    {
                        "name": "manage_api_tokens",
                        "status": "1"
                    },
                    {
                        "name": "edit_maintenance",
                        "status": "1"
                    },
                    {
                        "name": "manage_scheduled_reports",
                        "status": "1"
                    },
                    {
                        "name": "manage_sla",
                        "status": "1"
                    },
                    {
                        "name": "invoke_execute_now",
                        "status": "1"
                    },
                    {
                        "name": "change_problem_ranking",
                        "status": "1"
                    },
                    {
                        "name": "edit_own_media",
                        "status": "1"
                    },
                    {
                        "name": "edit_user_media",
                        "status": "1"
                    }
                ],
                "actions.default_access": "1"
            }
        }
    ],
    "id": 1
}
```


See also

- Role rules
- User

Source

CRole::get() in *ui/include/classes/api/services/CRole.php*.

role.update

Description

`object role.update(object/array roles)`

This method allows you to update existing roles.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Role properties to be updated.

The `roleid` property must be defined for each role, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard role properties the method accepts the following parameters.

Parameter	Type	Description
rules	array	Access rules to update for the role.

Return values

`(object)` Returns an object containing the IDs of the updated roles under the `roleids` property.

Examples

Disabling ability to execute scripts

Update role with ID "5", disable ability to execute scripts.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "role.update",
    "params": [
        {
            "roleid": "5",
            "rules": {
                "actions": [
                    {
                        "name": "execute_scripts",
                        "status": "0"
                    }
                ]
            }
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "roleids": [
            "5"
        ]
    },
    "id": 1
}
```

Limiting access to API

Update role with ID "5", deny to call any "create", "update" or "delete" methods.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "role.update",
```

```
    "params": [
        {
            "roleid": "5",
            "rules": {
                "api.access": "1",
                "api.mode": "0",
                "api": ["*.create", "*.update", "*.delete"]
            }
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "roleids": [
            "5"
        ]
    },
    "id": 1
}
```

Source

CRole::update() in *ui/include/classes/api/services/CRole.php*.

Script

This class is designed to work with scripts.

Object references:

- Script
  - Webhook parameters
- Debug
  - Log entry

Available methods:

- script.create - create new scripts
- script.delete - delete scripts
- script.execute - run scripts
- script.get - retrieve scripts
- script.getscriptsbyevents - retrieve scripts for events
- script.getscriptsbyhosts - retrieve scripts for hosts
- script.update - update scripts

Script object

The following objects are directly related to the `script` API.

Script

The script object has the following properties.

Property	Type	Description
scriptid	ID	ID of the script. Property behavior: - *read-only* - *required* for update operations

Property	Type	Description
name	string	Name of the script. Property behavior: - *required* for create operations
type	integer	Script type. Possible values if `scope` is set to ”action operation”: 0 - Script; 1 - IPMI; 2 - SSH; 3 - TELNET; 5 - Webhook. Possible values if `scope` is set to ”manual host action” or ”manual event action”: 6 - URL. Property behavior: - *required* for create operations
command	string	Command to run. Property behavior: - *required* if `type` is set to ”Script”, ”IPMI”, ”SSH”, ”TELNET”, or ”Webhook”
scope	integer	Script scope. Possible values: 1 - action operation; 2 - manual host action; 4 - manual event action. Property behavior: - *required* for create operations
execute_on	integer	Where to run the script. Possible values: 0 - run on Zabbix agent; 1 - run on Zabbix server. It is *supported* only if execution of global scripts is enabled on Zabbix server; 2 - *(default)* run on Zabbix server or proxy. Property behavior: - *supported* if `type` is set to ”Script”
menu_path	string	Folders separated by slash that form a menu like navigation in frontend when clicked on host or event. Property behavior: - *supported* if `scope` is set to ”manual host action” or ”manual event action”
authtype	integer	Authentication method used for SSH script type. Possible values: 0 - password; 1 - public key. Property behavior: - *supported* if `type` is set to ”SSH”
username	string	User name used for authentication. Property behavior: - *required* if `type` is set to ”SSH” or ”TELNET”

Property	Type	Description
password	string	Password used for SSH scripts with password authentication and TELNET scripts. Property behavior: - *supported* if `type` is set to ”SSH” and `authtype` is set to ”password”, or `type` is set to ”TELNET”
publickey	string	Name of the public key file used for SSH scripts with public key authentication. Property behavior: - *required* if `type` is set to ”SSH” and `authtype` is set to ”public key”
privatekey	string	Name of the private key file used for SSH scripts with public key authentication. Property behavior: - *required* if `type` is set to ”SSH” and `authtype` is set to ”public key”
port	string	Port number used for SSH and TELNET scripts. Property behavior: - *supported* if `type` is set to ”SSH” or ”TELNET”
groupid	ID	ID of the host group that the script can be run on. If set to ”0”, the script will be available on all host groups. Default: 0.
usrgrpid	ID	ID of the user group that will be allowed to run the script. If set to ”0”, the script will be available for all user groups. Default: 0. Property behavior: - *supported* if `scope` is set to ”manual host action” or ”manual event action”
host_access	integer	Host permissions needed to run the script. Possible values: 2 - *(default)* read; 3 - write. Property behavior: - *supported* if `scope` is set to ”manual host action” or ”manual event action”
confirmation	string	Confirmation pop up text. The pop up will appear when trying to run the script from the Zabbix frontend. Property behavior: - *supported* if `scope` is set to ”manual host action” or ”manual event action”
timeout	string	Webhook script execution timeout in seconds. Time suffixes are supported (e.g., `30s`, `1m`). Possible values: 1-60s. Default: 30s. Property behavior: - *required* if `type` is set to ”Webhook”

Property	Type	Description
parameters	array	Array of webhook input parameters. Property behavior: - *supported* if `type` is set to ”Webhook”
description	string	Description of the script.
url	string	User defined URL. Property behavior: - *required* if `type` is set to ”URL”
new_window	integer	Open URL in a new window. Possible values: 0 - No; 1 - *(default)* Yes. Property behavior: - *supported* if `type` is set to ”URL”
manualinput	integer	Indicates whether the script accepts user-provided input. Possible values: 0 - *(default)* Disabled; 1 - Enabled; Property behavior: - *supported* if `scope` is set to ”manual host action” or ”manual event action”
manualinput_prompt	string	Manual input prompt text. Property behavior: - *required* if `manualinput` is set to ”Enabled”
manualinput_validator	string	A character string field used to validate the user provided input. The string consists of either a regular expression or a set of values separated by commas. Property behavior: - *required* if `manualinput` is set to ”Enabled”
manualinput_validator_type	integer	Determines the type of user input expected. Possible values: 0 - *(default)* String. Indicates that manualinput_validator is to be treated as a regular expression; 1 - List. Indicates that manualinput_validator is to be treated as a comma-separated list of possible input values. Property behavior: - *supported* if `manualinput` is set to ”Enabled”
manualinput_default_value	string	Default value for auto-filling user input. Property behavior: - *supported* if `manualinput_validator_type` is set to ”String”

Webhook parameters

Parameters passed to webhook script when it is called have the following properties.

Property	Type	Description
name	string	Parameter name. Property behavior: - *required*
value	string	Parameter value. Supports macros.

Debug

Debug information of executed webhook script. The debug object has the following properties.

Property	Type	Description
logs	array	Array of log entries.
ms	string	Script execution duration in milliseconds.

Log entry

The log entry object has the following properties.

Property	Type	Description
level	integer	Log level.
ms	string	The time elapsed in milliseconds since the script was run before log entry was added.
message	string	Log message.

script.create

Description

`object script.create(object/array scripts)`

This method allows you to create new scripts.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Scripts to create.

The method accepts scripts with the standard script properties.

Return values

`(object)` Returns an object containing the IDs of the created scripts under the `scriptids` property. The order of the returned IDs matches the order of the passed scripts.

Examples

Create a webhook script

Create a webhook script that sends HTTP request to external service.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.create",
    "params": {
        "name": "Webhook script",
        "command": "try {\n var request = new HttpRequest(),\n response,\n data;\n\n request.addHeader('Co
        "scope": 1,
        "type": 5,
        "timeout": "40s",
        "parameters": [
            {
                "name": "token",
                "value": "{$WEBHOOK.TOKEN}"
            },
            {
                "name": "host",
                "value": "{HOST.HOST}"
```

```
            },
            {
                "name": "v",
                "value": "2.2"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "scriptids": [
            "3"
        ]
    },
    "id": 1
}
```

Create an SSH script

Create an SSH script with public key authentication that can be executed on a host and has a context menu.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.create",
    "params": {
        "name": "SSH script",
        "command": "my script command",
        "type": 2,
        "authtype": 1,
        "username": "John",
        "publickey": "pub.key",
        "privatekey": "priv.key",
        "password": "secret",
        "port": "12345",
        "scope": 2,
        "menu_path": "All scripts/SSH",
        "usrgrpid": "7",
        "groupid": "4"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "scriptids": [
            "5"
        ]
    },
    "id": 1
}
```

Create a custom script

Create a custom script that will reboot a server. The script will require write access to the host and will prompt the user for manual input. Upon successful input submission, script will display confirmation message in the frontend.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.create",
    "params": {
        "name": "Reboot server",
        "command": "reboot server {MANUALINPUT}",
        "type": 0,
        "scope": 2,
        "confirmation": "Are you sure you would like to reboot the server {MANUALINPUT}?",
        "manualinput": 1,
        "manualinput_prompt": "Which server you want to reboot?",
        "manualinput_validator": "[1-9]",
        "manualinput_validator_type": 0,
        "manualinput_default_value": "1"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "scriptids": [
            "4"
        ]
    },
    "id": 1
}
```

Create a URL type script

Create a URL type script for host scope that remains in the same window and has confirmation text.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.create",
    "params": {
        "name": "URL script",
        "type": 6,
        "scope": 2,
        "url": "http://zabbix/ui/zabbix.php?action=host.edit&hostid={HOST.ID}",
        "confirmation": "Edit host {HOST.NAME}?",
        "new_window": 0
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "scriptids": [
            "56"
        ]
    },
    "id": 1
}
```

Create a URL type script with manual input

Create a URL type script for event scope that opens in a new window and has manual input.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.create",
    "params": {
        "name": "URL script with manual input",
        "type": 6,
        "scope": 4,
        "url": "http://zabbix/ui/zabbix.php?action={MANUALINPUT}",
        "new_window": 1,
        "manualinput": 1,
        "manualinput_prompt": "Select a page to open:",
        "manualinput_validator": "dashboard.view,script.list,actionlog.list",
        "manualinput_validator_type": 1
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "scriptids": [
            "57"
        ]
    },
    "id": 1
}
```

Source

CScript::create() in *ui/include/classes/api/services/CScript.php*.

script.delete

Description

`object script.delete(array scriptIds)`

This method allows you to delete scripts.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the scripts to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted scripts under the `scriptids` property.

Examples

Delete multiple scripts

Delete two scripts.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.delete",
    "params": [
        "3",
        "4"
    ],
```

```
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "scriptids": [
            "3",
            "4"
        ]
    },
    "id": 1
}
```

Source

CScript::delete() in *ui/include/classes/api/services/CScript.php*.

script.execute

Description

`object script.execute(object parameters)`

This method allows you to run a script on a host or event. Except for URL type scripts; those are not executable.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters containing the ID of the script to run, either the ID of the host or the ID of the event and manualinput value.

Parameter	Type	Description
scriptid	ID	ID of the script to run. Parameter behavior: - *required*
hostid	ID	ID of the host to run the script on. Parameter behavior: - *required* if `eventid` is not set
eventid	ID	ID of the event to run the script on. Parameter behavior: - *required* if `hostid` is not set
manualinput	string	User-provided value to run the script with, substituting the {MANUALINPUT} macro.

Return values

`(object)` Returns the result of script execution.

Property	Type	Description
response	string	Whether the script was run successfully. Possible value - `success`.
value	string	Script output.
debug	object	Contains a debug object if a webhook script is executed. For other script types, it contains empty object.

Examples

Run a webhook script

Run a webhook script that sends HTTP request to external service.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.execute",
    "params": {
        "scriptid": "4",
        "hostid": "30079"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "response": "success",
        "value": "{\"status\":\"sent\",\"timestamp\":\"1611235391\"}",
        "debug": {
            "logs": [
                {
                    "level": 3,
                    "ms": 480,
                    "message": "[Webhook Script] HTTP status: 200."
                }
            ],
            "ms": 495
        }
    },
    "id": 1
}
```

Run a custom script

Run a ”ping” script on a host.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.execute",
    "params": {
        "scriptid": "1",
        "hostid": "30079"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "response": "success",
        "value": "PING 127.0.0.1 (127.0.0.1) 56(84) bytes of data.\n64 bytes from 127.0.0.1: icmp_req=1 tt
        "debug": []
    },
    "id": 1
}
```

Run a custom script with manual input

Run a ”ping” script with command ”ping -c {MANUALINPUT} {HOST.CONN}; case $? in [01]) true;; *) false;; esac” on a host.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.execute",
    "params": {
        "scriptid": "7",
        "hostid": "30079",
        "manualinput": "2"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "response": "success",
        "value": "PING 127.0.0.1 (127.0.0.1) 56(84) bytes of data.\n64 bytes from 127.0.0.1: icmp_seq=1 tt
        "debug": []
    },
    "id": 1
}
```

Source

CScript::execute() in *ui/include/classes/api/services/CScript.php*.

script.get

Description

`integer/array script.get(object parameters)`

The method allows you to retrieve scripts according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
groupids	ID/array	Return only scripts that can be run on the given host groups.
hostids	ID/array	Return only scripts that can be run on the given hosts.
scriptids	ID/array	Return only scripts with the given IDs.
usrgrpids	ID/array	Return only scripts that can be run by users in the given user groups.
selectHostGroups	query	Return a `hostgroups` property with host groups that the script can be run on.
selectHosts	query	Return a `hosts` property with hosts that the script can be run on.
selectActions	query	Return an `actions` property with actions that the script is associated with.
sortfield	string/array	Sort the result by the given properties. Possible values: `scriptid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	

Parameter	Type	Description
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve all scripts

Retrieve all configured scripts.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.get",
    "params": {
        "output": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "scriptid": "1",
            "name": "Ping",
            "command": "/bin/ping -c 3 {HOST.CONN} 2>&1",
            "host_access": "2",
            "usrgrpid": "0",
            "groupid": "0",
            "description": "",
            "confirmation": "",
            "type": "0",
            "execute_on": "1",
            "timeout": "30s",
            "scope": "2",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "",
            "new_window": "1",
            "manualinput": "0",
            "manualinput_prompt": "",
            "manualinput_validator": "",
            "manualinput_validator_type": "0",
            "manualinput_default_value": "",
```

```
            "parameters": []
        },
        {
            "scriptid": "2",
            "name": "Traceroute",
            "command": "/usr/bin/traceroute {HOST.CONN} 2>&1",
            "host_access": "2",
            "usrgrpid": "0",
            "groupid": "0",
            "description": "",
            "confirmation": "",
            "type": "0",
            "execute_on": "1",
            "timeout": "30s",
            "scope": "2",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "",
            "new_window": "1",
            "manualinput": "0",
            "manualinput_prompt": "",
            "manualinput_validator": "",
            "manualinput_validator_type": "0",
            "manualinput_default_value": "",
            "parameters": []
        },
        {
            "scriptid": "3",
            "name": "Detect operating system",
            "command": "sudo /usr/bin/nmap -O {HOST.CONN} 2>&1",
            "host_access": "2",
            "usrgrpid": "7",
            "groupid": "0",
            "description": "",
            "confirmation": "",
            "type": "0",
            "execute_on": "1",
            "timeout": "30s",
            "scope": "2",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "",
            "new_window": "1",
            "manualinput": "0",
            "manualinput_prompt": "",
            "manualinput_validator": "",
            "manualinput_validator_type": "0",
            "manualinput_default_value": "",
            "parameters": []
        },
        {
            "scriptid": "4",
```

```
        "name": "Webhook",
        "command": "try {\n var request = new HttpRequest(),\n response,\n data;\n\n request.addHeader
        "host_access": "2",
        "usrgrpid": "7",
        "groupid": "0",
        "description": "",
        "confirmation": "",
        "type": "5",
        "execute_on": "1",
        "timeout": "30s",
        "scope": "2",
        "port": "",
        "authtype": "0",
        "username": "",
        "password": "",
        "publickey": "",
        "privatekey": "",
        "menu_path": "",
        "url": "",
        "new_window": "1",
        "manualinput": "0",
        "manualinput_prompt": "",
        "manualinput_validator": "",
        "manualinput_validator_type": "0",
        "manualinput_default_value": "",
        "parameters": [
            {
                "name": "token",
                "value": "{$WEBHOOK.TOKEN}"
            },
            {
                "name": "host",
                "value": "{HOST.HOST}"
            },
            {
                "name": "v",
                "value": "2.2"
            }
        ]
    },
    {
        "scriptid": "5",
        "name": "URL",
        "command": "",
        "host_access": "2",
        "usrgrpid": "0",
        "groupid": "0",
        "description": "",
        "confirmation": "Go to {HOST.NAME}?",
        "type": "6",
        "execute_on": "1",
        "timeout": "30s",
        "scope": "4",
        "port": "",
        "authtype": "0",
        "username": "",
        "password": "",
        "publickey": "",
        "privatekey": "",
        "menu_path": "",
        "url": "http://zabbix/ui/zabbix.php?action=latest.view&hostids[]={HOST.ID}",
        "new_window": "0",
```

```
            "manualinput": "0",
            "manualinput_prompt": "",
            "manualinput_validator": "",
            "manualinput_validator_type": "0",
            "manualinput_default_value": "",
            "parameters": []
        },
        {
            "scriptid": "6",
            "name": "URL with user input",
            "command": "",
            "host_access": "2",
            "usrgrpid": "0",
            "groupid": "0",
            "description": "",
            "confirmation": "Open zabbix page {MANUALINPUT}?",
            "type": "6",
            "execute_on": "1",
            "timeout": "30s",
            "scope": "2",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "http://zabbix/ui/zabbix.php?action={MANUALINPUT}",
            "new_window": "0",
            "manualinput": "1",
            "manualinput_prompt": "Select a page to open:",
            "manualinput_validator": "dashboard.view,script.list,actionlog.list",
            "manualinput_validator_type": "1",
            "parameters": []
        }
    ],
    "id": 1
}
```


See also

- Host
- Host group

Source

CScript::get() in *ui/include/classes/api/services/CScript.php*.

script.getscriptsbyevents

Description

`object script.getscriptsbyevents(object parameters)`

This method allows you to retrieve all available scripts on the given event or specific script if script ID is provided. When manualinput is provided, it substitutes the `{MANUALINPUT}` macro with the specified value.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` The method accepts object or array of objects with the following parameters.

Parameter	Type	Description
eventid	ID	ID of event to return scripts for. Must be unique. Parameter behavior: - *required*
scriptid	ID	ID of script to return.
manualinput	string	Value of the user-provided {MANUALINPUT} macro value.

Return values

(object) Returns an object with event IDs as properties and arrays of available scripts as values. If script ID is provided, the associated value is an array containing the specific script.

Note:
The method will automatically expand macros in the `confirmation` text, `manualinput prompt` text and `url`.
If the manualinput parameter is provided, the {MANUALINPUT} macro will be resolved to the specified value.

Examples

Retrieve scripts by event IDs

Retrieve all scripts available to events ”632” and ”614”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.getscriptsbyevents",
    "params": [
      {
         "eventid":  "632"
      },
      {
         "eventid":  "614"
      }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "632": [
            {
                "scriptid": "3",
                "name": "Detect operating system",
                "command": "sudo /usr/bin/nmap -O {HOST.CONN} 2>&1",
                "host_access": "2",
                "usrgrpid": "7",
                "groupid": "0",
                "description": "",
                "confirmation": "",
                "type": "0",
                "execute_on": "1",
                "timeout": "30s",
                "scope": "4",
                "port": "",
                "authtype": "0",
                "username": "",
                "password": "",
                "publickey": "",
                "privatekey": "",
```

```
            "menu_path": "",
            "url": "",
            "new_window": "1",
            "manualinput": "0",
            "manualinput_prompt": "",
            "manualinput_validator_type": "0",
            "manualinput_validator": "",
            "manualinput_default_value": "",
            "parameters": []
        },
        {
            "scriptid": "1",
            "name": "Ping",
            "command": "/bin/ping -c 3 {HOST.CONN} 2>&1",
            "host_access": "2",
            "usrgrpid": "0",
            "groupid": "0",
            "description": "",
            "confirmation": "",
            "type": "0",
            "execute_on": "1",
            "timeout": "30s",
            "scope": "4",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "",
            "new_window": "1",
            "manualinput": "0",
            "manualinput_prompt": "",
            "manualinput_validator_type": "0",
            "manualinput_validator": "",
            "manualinput_default_value": "",
            "parameters": []
        },
        {
            "scriptid": "4",
            "name": "Open Zabbix page",
            "command": "",
            "host_access": "2",
            "usrgrpid": "0",
            "groupid": "0",
            "description": "",
            "confirmation": "Are you sure you want to open page *UNKNOWN*?",
            "type": "6",
            "execute_on": "2",
            "timeout": "30s",
            "scope": "4",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "http://localhost/ui/zabbix.php?action=*UNKNOWN*",
            "new_window": "1",
            "manualinput": "1",
```

```
            "manualinput_prompt": "Zabbix page to open:",
            "manualinput_validator_type": "1",
            "manualinput_validator": "dashboard.view,discovery.view",
            "manualinput_default_value": "",
            "parameters": []
        },
        {
            "scriptid": "2",
            "name": "Traceroute",
            "command": "/usr/bin/traceroute {HOST.CONN} 2>&1",
            "host_access": "2",
            "usrgrpid": "0",
            "groupid": "0",
            "description": "",
            "confirmation": "",
            "type": "0",
            "execute_on": "1",
            "timeout": "30s",
            "scope": "4",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "",
            "new_window": "1",
            "manualinput": "0",
            "manualinput_prompt": "",
            "manualinput_validator_type": "0",
            "manualinput_validator": "",
            "manualinput_default_value": "",
            "parameters": []
        }
    ],
    "614": [
        {
            "scriptid": "3",
            "name": "Detect operating system",
            "command": "sudo /usr/bin/nmap -O {HOST.CONN} 2>&1",
            "host_access": "2",
            "usrgrpid": "7",
            "groupid": "0",
            "description": "",
            "confirmation": "",
            "type": "0",
            "execute_on": "1",
            "timeout": "30s",
            "scope": "4",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "",
            "new_window": "1",
            "manualinput": "0",
            "manualinput_prompt": "",
            "manualinput_validator_type": "1",
```

```
            "manualinput_validator": "",
            "manualinput_default_value": "",
            "parameters": []
        },
        {
            "scriptid": "1",
            "name": "Ping",
            "command": "/bin/ping -c 3 {HOST.CONN} 2>&1",
            "host_access": "2",
            "usrgrpid": "0",
            "groupid": "0",
            "description": "",
            "confirmation": "",
            "type": "0",
            "execute_on": "1",
            "timeout": "30s",
            "scope": "4",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "",
            "new_window": "1",
            "manualinput": "0",
            "manualinput_prompt": "",
            "manualinput_validator_type": "0",
            "manualinput_validator": "",
            "manualinput_default_value": "",
            "parameters": []
        },
        {
            "scriptid": "4",
            "name": "Open Zabbix page",
            "command": "",
            "host_access": "2",
            "usrgrpid": "0",
            "groupid": "0",
            "description": "",
            "confirmation": "Are you sure you want to open page *UNKNOWN*?",
            "type": "6",
            "execute_on": "2",
            "timeout": "30s",
            "scope": "4",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "http://localhost/ui/zabbix.php?action=*UNKNOWN*",
            "new_window": "1",
            "manualinput": "1",
            "manualinput_prompt": "Zabbix page to open:",
            "manualinput_validator_type": "1",
            "manualinput_validator": "dashboard.view,discovery.view",
            "manualinput_default_value": "",
            "parameters": []
        },
```

```
            {
                "scriptid": "2",
                "name": "Traceroute",
                "command": "/usr/bin/traceroute {HOST.CONN} 2>&1",
                "host_access": "2",
                "usrgrpid": "0",
                "groupid": "0",
                "description": "",
                "confirmation": "",
                "type": "0",
                "execute_on": "1",
                "timeout": "30s",
                "scope": "4",
                "port": "",
                "authtype": "0",
                "username": "",
                "password": "",
                "publickey": "",
                "privatekey": "",
                "menu_path": "",
                "url": "",
                "new_window": "1",
                "manualinput": "0",
                "manualinput_prompt": "",
                "manualinput_validator_type": "0",
                "manualinput_validator": "",
                "manualinput_default_value": "",
                "parameters": []
            }
        ]
    },
    "id": 1
}
```

Retrieve specific script with manualinput value.

Retrieve script with ID "4" on event "632" with manualinput value "dashboard.view".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.getscriptsbyevents",
    "params": [
      {
        "eventid":  "632",
        "scriptid": "4",
        "manualinput": "dashboard.view"
      }
  ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
      "632": [
        {
          "scriptid": "4",
          "name": "Open Zabbix page",
          "command": "",
          "host_access": "2",
          "usrgrpid": "0",
```

```
            "groupid": "0",
            "description": "",
            "confirmation": "Are you sure you want to open page dashboard.view?",
            "type": "6",
            "execute_on": "2",
            "timeout": "30s",
            "scope": "4",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "http://localhost/ui/zabbix.php?action=dashboard.view",
            "new_window": "1",
            "manualinput": "1",
            "manualinput_prompt": "Zabbix page to open:",
            "manualinput_validator_type": "1",
            "manualinput_validator": "dashboard.view,discovery.view",
            "manualinput_default_value": "",
            "parameters": []
        }
      ]
    },
    "id": 1
}
```


Source

CScript::getScriptsByEvents() in *ui/include/classes/api/services/CScript.php*.

script.getscriptsbyhosts

Description

`object script.getscriptsbyhosts(object parameters)`

This method allows you to retrieve all available scripts on the given host or specific script if script ID is provided. When manualinput is provided, it substitutes the `{MANUALINPUT}` macro with the specified value.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` The method accepts object or array of objects with the following parameters.

Parameter	Type	Description
hostid	ID	ID of host to return scripts for. Must be unique. Parameter behavior: - *required*
scriptid	ID	ID of script to return.
manualinput	string	Value of the user-provided {MANUALINPUT} macro value.

Return values

`(object)` Returns an object with host IDs as properties and arrays of available scripts as values. If script ID is provided, the associated value is an array containing the specific script.

Note:
The method will automatically expand macros in the `confirmation` text, `manualinput prompt` text and `url`.
If the manualinput parameter is provided, the {MANUALINPUT} macro will be resolved to the specified value.

Examples

Retrieve scripts by host IDs

Retrieve all scripts available on hosts ”30079” and ”30073”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.getscriptsbyhosts",
    "params": [
       {
          "hostid":  "30079"
       },
       {
          "hostid":  "30073"
       }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "30079": [
            {
                "scriptid": "3",
                "name": "Detect operating system",
                "command": "sudo /usr/bin/nmap -O {HOST.CONN} 2>&1",
                "host_access": "2",
                "usrgrpid": "7",
                "groupid": "0",
                "description": "",
                "confirmation": "",
                "type": "0",
                "execute_on": "1",
                "timeout": "30s",
                "scope": "2",
                "port": "",
                "authtype": "0",
                "username": "",
                "password": "",
                "publickey": "",
                "privatekey": "",
                "menu_path": "",
                "url": "",
                "new_window": "1",
                "manualinput": "0",
                "manualinput_prompt": "",
                "manualinput_validator_type": "0",
                "manualinput_validator": "",
                "manualinput_default_value": "",
                "parameters": []
            },
            {
                "scriptid": "1",
                "name": "Ping",
                "command": "/bin/ping -c 3 {HOST.CONN} 2>&1",
```

```
        "host_access": "2",
        "usrgrpid": "0",
        "groupid": "0",
        "description": "",
        "confirmation": "",
        "type": "0",
        "execute_on": "1",
        "timeout": "30s",
        "scope": "2",
        "port": "",
        "authtype": "0",
        "username": "",
        "password": "",
        "publickey": "",
        "privatekey": "",
        "menu_path": "",
        "url": "",
        "new_window": "1",
        "manualinput": "0",
        "manualinput_prompt": "",
        "manualinput_validator_type": "0",
        "manualinput_validator": "",
        "manualinput_default_value": "",
        "parameters": []
    },
    {
        "scriptid": "4",
        "name": "Open Zabbix page",
        "command": "",
        "host_access": "2",
        "usrgrpid": "0",
        "groupid": "0",
        "description": "",
        "confirmation": "Are you sure you want to open page *UNKNOWN*?",
        "type": "6",
        "execute_on": "2",
        "timeout": "30s",
        "scope": "2",
        "port": "",
        "authtype": "0",
        "username": "",
        "password": "",
        "publickey": "",
        "privatekey": "",
        "menu_path": "",
        "url": "http://localhost/ui/zabbix.php?action=*UNKNOWN*",
        "new_window": "1",
        "manualinput": "0",
        "manualinput_prompt": "Zabbix page to open:",
        "manualinput_validator_type": "0",
        "manualinput_validator": "dashboard.view,discovery.view",
        "manualinput_default_value": "",
        "parameters": []
    },
    {
        "scriptid": "2",
        "name": "Traceroute",
        "command": "/usr/bin/traceroute {HOST.CONN} 2>&1",
        "host_access": "2",
        "usrgrpid": "0",
        "groupid": "0",
        "description": "",
```

```
            "confirmation": "",
            "type": "0",
            "execute_on": "1",
            "timeout": "30s",
            "scope": "2",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "",
            "new_window": "1",
            "manualinput": "0",
            "manualinput_prompt": "",
            "manualinput_validator_type": "0",
            "manualinput_validator": "",
            "manualinput_default_value": "",
            "parameters": []
        }
    ],
    "30073": [
        {
            "scriptid": "3",
            "name": "Detect operating system",
            "command": "sudo /usr/bin/nmap -O {HOST.CONN} 2>&1",
            "host_access": "2",
            "usrgrpid": "7",
            "groupid": "0",
            "description": "",
            "confirmation": "",
            "type": "0",
            "execute_on": "1",
            "timeout": "30s",
            "scope": "2",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "",
            "new_window": "1",
            "manualinput": "0",
            "manualinput_prompt": "",
            "manualinput_validator_type": "0",
            "manualinput_validator": "",
            "manualinput_default_value": "",
            "parameters": []
        },
        {
            "scriptid": "1",
            "name": "Ping",
            "command": "/bin/ping -c 3 {HOST.CONN} 2>&1",
            "host_access": "2",
            "usrgrpid": "0",
            "groupid": "0",
            "description": "",
            "confirmation": "",
            "type": "0",
```

```
            "execute_on": "1",
            "timeout": "30s",
            "scope": "2",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "",
            "new_window": "1",
            "manualinput": "0",
            "manualinput_prompt": "",
            "manualinput_validator_type": "0",
            "manualinput_validator": "",
            "manualinput_default_value": "",
            "parameters": []
        },
        {
            "scriptid": "4",
            "name": "Open Zabbix page",
            "command": "",
            "host_access": "2",
            "usrgrpid": "0",
            "groupid": "0",
            "description": "",
            "confirmation": "Are you sure you want to open page *UNKNOWN*?",
            "type": "6",
            "execute_on": "2",
            "timeout": "30s",
            "scope": "2",
            "port": "",
            "authtype": "0",
            "username": "",
            "password": "",
            "publickey": "",
            "privatekey": "",
            "menu_path": "",
            "url": "http://localhost/ui/zabbix.php?action=*UNKNOWN*",
            "new_window": "1",
            "manualinput": "1",
            "manualinput_prompt": "Zabbix page to open:",
            "manualinput_validator_type": "1",
            "manualinput_validator": "dashboard.view,discovery.view",
            "manualinput_default_value": "",
            "parameters": []
        },
        {
            "scriptid": "2",
            "name": "Traceroute",
            "command": "/usr/bin/traceroute {HOST.CONN} 2>&1",
            "host_access": "2",
            "usrgrpid": "0",
            "groupid": "0",
            "description": "",
            "confirmation": "",
            "type": "0",
            "execute_on": "1",
            "timeout": "30s",
            "scope": "2",
            "port": "",
```

```
                "authtype": "0",
                "username": "",
                "password": "",
                "publickey": "",
                "privatekey": "",
                "menu_path": "",
                "url": "",
                "new_window": "1",
                "manualinput": "0",
                "manualinput_prompt": "",
                "manualinput_validator_type": "0",
                "manualinput_validator": "",
                "manualinput_default_value": "",
                "parameters": []
            }
        ]
    },
    "id": 1
}
```


Retrieve specific script with manualinput value.

Retrieve script with ID ”4” on host ”30079” with manualinput value ”dashboard.view”.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "script.getscriptsbyhosts",
    "params": [
      {
        "hostid":  "30079",
        "scriptid": "4",
        "manualinput": "dashboard.view"
      }
  ],
    "id": 1
}
```


Response:


```
{
    "jsonrpc": "2.0",
    "result": {
      "30079": [
        {
          "scriptid": "4",
          "name": "Open Zabbix page",
          "command": "",
          "host_access": "2",
          "usrgrpid": "0",
          "groupid": "0",
          "description": "",
          "confirmation": "Are you sure you want to open page dashboard.view?",
          "type": "6",
          "execute_on": "2",
          "timeout": "30s",
          "scope": "2",
          "port": "",
          "authtype": "0",
          "username": "",
          "password": "",
          "publickey": "",
          "privatekey": "",
          "menu_path": "",
```

```
            "url": "http://localhost/ui/zabbix.php?action=dashboard.view",
            "new_window": "1",
            "manualinput": "1",
            "manualinput_prompt": "Zabbix page to open:",
            "manualinput_validator_type": "1",
            "manualinput_validator": "dashboard.view,discovery.view",
            "manualinput_default_value": "",
            "parameters": []
        }
      ]
    },
    "id": 1
}
```

Source

CScript::getScriptsByHosts() in *ui/include/classes/api/services/CScript.php*.

script.update

Description

`object script.update(object/array scripts)`

This method allows you to update existing scripts.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Script properties to be updated.

The `scriptid` property must be defined for each script, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged. An exception is `type` property change from 5 (Webhook) to other: the `parameters` property will be cleaned.

Return values

`(object)` Returns an object containing the IDs of the updated scripts under the `scriptids` property.

Examples

Change script command

Change the command of the script to "/bin/ping -c 10 {HOST.CONN} 2>&1".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.update",
    "params": {
        "scriptid": "1",
        "command": "/bin/ping -c 10 {HOST.CONN} 2>&1"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "scriptids": [
            "1"
        ]
    },
```

```
    "id": 1
}
```

Change script command and add manual input

Change the command of the script to "/bin/ping -c {MANUALINPUT} {HOST.CONN} 2>&1".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "script.update",
    "params": {
      "scriptid": "1",
      "command": "/bin/ping -c {MANUALINPUT} {HOST.CONN} 2>&1",
      "manualinput": "1",
      "manualinput_prompt": "Specify the number of ICMP packets to send with the ping command",
      "manualinput_validator": "^(?:[1-9]|10)$",
      "manualinput_validator_type": "0",
      "manualinput_default_value": "10"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "scriptids": [
            "1"
        ]
    },
    "id": 1
}
```

Source

CScript::update() in *ui/include/classes/api/services/CScript.php*.

Server info

This class is designed to retrieve information about Zabbix server.

Available methods:

- serverinfo.get - retrieve the time of the last update from Zabbix server

serverinfo.get

This method allows you to retrieve the time of the last update from the given Zabbix server.

The method is available to *User*, *Admin*, *Super admin*, and unauthenticated users.

Method signature: `object serverinfo.get(object parameters)`

Method source: `CServerInfo::get()` in `ui/include/classes/api/services/CServerInfo.php`

Parameters

Parameter	Type	Description
serverid	ID	ID of the Zabbix server. Parameter behavior: - Required
output	query	This parameter is described in the reference commentary.

Return values

Property	Type	Description
serverid	ID	Zabbix server ID.
lastaccess	timestamp	Time of the last update from Zabbix server. For a high availability (HA) cluster, this shows the `lastaccess` value of the active node.

Examples

Retrieve server info

Retrieve information about the given Zabbix server.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "serverinfo.get",
    "params": {
        "serverid": "019fb1f9-22bd-74ea-bc03-51efb8b76aba",
        "output": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "serverid": "019fb1f9-22bd-74ea-bc03-51efb8b76aba",
        "lastaccess": 1785337245
    },
    "id": 1
}
```

Service

This class is designed to work with IT infrastructure/business services.

Object references:

- Service
- Status rule
- Service tag
- Service alarm
- Problem tag

Available methods:

- service.create - create new services
- service.delete - delete services
- service.get - retrieve services
- service.update - update services

Service object

The following objects are directly related to the `service` API.

Service

The service object has the following properties.

Property	Type	Description
serviceid	ID	ID of the service. Property behavior: - *read-only* - *required* for update operations
algorithm	integer	Status calculation rule. Only applicable if child services exist. Possible values: 0 - set status to OK; 1 - most critical if all children have problems; 2 - most critical of child services. Property behavior: - *required* for create operations
name	string	Name of the service. Property behavior: - *required* for create operations
sortorder	integer	Position of the service used for sorting. Possible values: 0-999. Property behavior: - *required* for create operations
weight	integer	Service weight. Possible values: 0-1000000. Default: 0.
propagation_rule	integer	Status propagation rule. Possible values: 0 - *(default)* propagate service status as is - without any changes; 1 - increase the propagated status by a given `propagation_value` (by 1 to 5 severities); 2 - decrease the propagated status by a given `propagation_value` (by 1 to 5 severities); 3 - ignore this service - the status is not propagated to the parent service at all; 4 - set fixed service status using a given `propagation_value`. Property behavior: - *required* if `propagation_value` is set

Property	Type	Description
propagation_value	integer	Status propagation value. Possible values if `propagation_rule` is set to "0" or "3": 0 - Not classified. Possible values if `propagation_rule` is set to "1" or "2": 1 - Information; 2 - Warning; 3 - Average; 4 - High; 5 - Disaster. Possible values if `propagation_rule` is set to "4": -1 - OK; 0 - Not classified; 1 - Information; 2 - Warning; 3 - Average; 4 - High; 5 - Disaster. Property behavior: - *required* if `propagation_rule` is set
status	integer	Whether the service is in OK or problem state. If the service is in a problem state, `status` equals either the severity of the most critical problem or the highest status of a child service in a problem state. If the service is in OK state, `status` is equal to: -1. Property behavior: - *read-only*
description	string	Description of the service.
uuid	string	Universal unique identifier, used for linking imported services to already existing ones. Auto-generated, if not given.
created_at	integer	Unix timestamp when service was created.
readonly	integer	Access to the service. Possible values: 0 - Read-write; 1 - Read-only. Property behavior: - *read-only*

Status rule

The status rule object has the following properties.

Property	Type	Description
type	integer	Condition for setting (New status) status. Possible values: 0 - if at least (N) child services have (Status) status or above; 1 - if at least (N%) of child services have (Status) status or above; 2 - if less than (N) child services have (Status) status or below; 3 - if less than (N%) of child services have (Status) status or below; 4 - if weight of child services with (Status) status or above is at least (W); 5 - if weight of child services with (Status) status or above is at least (N%); 6 - if weight of child services with (Status) status or below is less than (W); 7 - if weight of child services with (Status) status or below is less than (N%). Where: - N (W) is `limit_value`; - (Status) is `limit_status`; - (New status) is `new_status`. Property behavior: - *required*
limit_value	integer	Limit value. Possible values: - for N and W: 1-100000; - for N%: 1-100. Property behavior: - *required*
limit_status	integer	Limit status. Possible values: -1 - OK; 0 - Not classified; 1 - Information; 2 - Warning; 3 - Average; 4 - High; 5 - Disaster. Property behavior: - *required*
new_status	integer	New status value. Possible values: 0 - Not classified; 1 - Information; 2 - Warning; 3 - Average; 4 - High; 5 - Disaster. Property behavior: - *required*

Service tag

The service tag object has the following properties.

Property	Type	Description
tag	string	Service tag name. Property behavior: - *required*
value	string	Service tag value.

Service alarm

Note:
Service alarms cannot be directly created, updated or deleted via the Zabbix API.

The service alarm objects represent a service's state change. It has the following properties.

Property	Type	Description
clock	timestamp	Time when the service state change has happened.
value	integer	Status of the service. Refer to the service `status` property for a list of possible values.

Problem tag

Problem tags allow linking services with problem events. The problem tag object has the following properties.

Property	Type	Description
tag	string	Problem tag name. Property behavior: - *required*
operator	integer	Mapping condition operator. Possible values: 0 - *(default)* Equals; 2 - Contains.
value	string	Problem tag value.

service.create

Description

`object service.create(object/array services)`

This method allows you to create new services.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` services to create.

Additionally to the standard service properties, the method accepts the following parameters.

Parameter	Type	Description
children	array	Child services to be linked to the service. The child services must have only the `serviceid` property defined.
parents	array	Parent services to be linked to the service. The parent services must have only the `serviceid` property defined.

Parameter	Type	Description
tags	array	Service tags to be created for the service.
problem_tags	array	Problem tags to be created for the service.
status_rules	array	Status rules to be created for the service.

Return values

(object) Returns an object containing the IDs of the created services under the serviceids property. The order of the returned IDs matches the order of the passed services.

Examples

Creating a service

Create a service that will be switched to problem state, if at least one child has a problem.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "service.create",
    "params": {
        "name": "Server 1",
        "algorithm": 1,
        "sortorder": 1
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "serviceids": [
            "5"
        ]
    },
    "id": 1
}
```

Source

CService::create() in *ui/include/classes/api/services/CService.php*.

service.delete

Description

`object service.delete(array serviceIds)`

This method allows you to delete services.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

(array) IDs of the services to delete.

Return values

(object) Returns an object containing the IDs of the deleted services under the serviceids property.

Examples

Deleting multiple services

Delete two services.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "service.delete",
    "params": [
        "4",
        "5"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "serviceids": [
            "4",
            "5"
        ]
    },
    "id": 1
}
```

Source

CService::delete() in *ui/include/classes/api/services/CService.php*.

service.get

Description

`integer/array service.get(object parameters)`

The method allows you to retrieve services according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
serviceids	ID/array	Return only services with the given IDs.
parentids	ID/array	Return only services that are linked to the given parent services.
deep_parentids	flag	Return all direct and indirect child services. Used together with `parentids`.
childids	ID/array	Return only services that are linked to the given child services.
evaltype	integer	Tag evaluation method. Possible values: 0 - *(default)* And/Or; 2 - Or.

Parameter	Type	Description
tags	object/array	Return only services with the given tags. Format: `[{"tag": "<tag>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all services. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
problem_tags	object/array	Return only services with the given problem tags. Format: `[{"tag": "<tag>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all services. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
without_problem_tags	flag	Return only services without problem tags.
slaids	ID/array	Return only services that are linked to the specific SLA(s).
selectChildren	query	Return a `children` property with the child services. Supports `count`.
selectParents	query	Return a `parents` property with the parent services. Supports `count`.
selectTags	query	Return a `tags` property with service tags. Supports `count`.
selectProblemEvents	query	Return a `problem_events` property with an array of problem event objects. The problem event object has the following properties: `eventid` - `(ID)` Event ID; `severity` - `(string)` Current event severity; `name` - `(string)` Resolved event name. Supports `count`.
selectProblemTags	query	Return a `problem_tags` property with problem tags. Supports `count`.
selectStatusRules	query	Return a `status_rules` property with status rules. Supports `count`.
selectStatusTimeline	object/array	Return a `status_timeline` property containing service state changes for given periods. Format `[{"period_from": "<period_from>", "period_to": "<period_to>"}, ...]` - `period_from` being a starting date (inclusive; integer timestamp) and `period_to` being an ending date (exclusive; integer timestamp) for the period you're interested in. Returns an array of entries containing a `start_value` property and an `alarms` array for the state changes within specified periods.

Parameter	Type	Description
sortfield	string/array	Sort the result by the given properties. Possible values: `serviceid`, `name`, `status`, `sortorder`, `created_at`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving all services

Retrieve all data about all services and their dependencies.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "service.get",
    "params": {
        "output": "extend",
        "selectChildren": "extend",
        "selectParents": ["serviceid", "name"]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "serviceid": "1",
            "name": "Zabbix cluster",
            "status": "-1",
            "algorithm": "2",
            "sortorder": "0",
            "weight": "0",
            "propagation_rule": "0",
            "propagation_value": "0",
            "description": "",
            "uuid": "8d77bd91b62347e4b79382912eb5df95",
            "created_at": "1761742392",
            "readonly": false,
            "parents": [],
            "children": [
                {
                    "serviceid": "2",
```

```
                "name": "Zabbix server node 1",
                "status": "-1",
                "algorithm": "2",
                "sortorder": "0",
                "weight": "0",
                "propagation_rule": "0",
                "propagation_value": "0",
                "description": "",
                "uuid": "195911d26d7f4e218d6217079bcd5929",
                "created_at": "1761742623",
                "readonly": false
            },
            {
                "serviceid": "3",
                "name": "Zabbix server node 2",
                "status": "-1",
                "algorithm": "2",
                "sortorder": "0",
                "weight": "0",
                "propagation_rule": "0",
                "propagation_value": "0",
                "description": "",
                "uuid": "9fc659a30fe244f690dff25fc2a9db5c",
                "created_at": "1761742654",
                "readonly": false
            }
        ]
    },
    {
        "serviceid": "2",
        "name": "Zabbix server node 1",
        "status": "-1",
        "algorithm": "2",
        "sortorder": "0",
        "weight": "0",
        "propagation_rule": "0",
        "propagation_value": "0",
        "description": "",
        "uuid": "195911d26d7f4e218d6217079bcd5929",
        "created_at": "1761742623",
        "readonly": false,
        "parents": [
            {
                "serviceid": "1",
                "name": "Zabbix cluster"
            }
        ],
        "children": []
    },
    {
        "serviceid": "3",
        "name": "Zabbix server node 2",
        "status": "-1",
        "algorithm": "2",
        "sortorder": "0",
        "weight": "0",
        "propagation_rule": "0",
        "propagation_value": "0",
        "description": "",
        "uuid": "9fc659a30fe244f690dff25fc2a9db5c",
        "created_at": "1761742654",
        "readonly": false,
```

```
            "parents": [
                {
                    "serviceid": "1",
                    "name": "Zabbix cluster"
                }
            ],
            "children": []
        }
    ],
    "id": 1
}
```

Source

CService::get() in *ui/include/classes/api/services/CService.php*.

service.update

Description

`object service.update(object/array services)`

This method allows you to update existing services.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` service properties to be updated.

The `serviceid` property must be defined for each service, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard service properties, the method accepts the following parameters.

Parameter	Type	Description
children	array	Child services to replace the current child services. The child services must have only the `serviceid` property defined.
parents	array	Parent services to replace the current parent services. The parent services must have only the `serviceid` property defined.
tags	array	Service tags to replace the current service tags.
problem_tags	array	Problem tags to replace the current problem tags.
status_rules	array	Status rules to replace the current status rules.

Return values

`(object)` Returns an object containing the IDs of the updated services under the `serviceids` property.

Examples

Setting the parent for a service

Make service with ID "3" to be the parent for service with ID "5".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "service.update",
    "params": {
        "serviceid": "5",
        "parents": [
            {
```

```
                "serviceid": "3"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "serviceids": [
            "5"
        ]
    },
    "id": 1
}
```

Adding a scheduled downtime

Add a downtime for service with ID ”4” scheduled weekly from Monday 22:00 till Tuesday 10:00.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "service.update",
    "params": {
        "serviceid": "4",
        "times": [
            {
                "type": "1",
                "ts_from": "165600",
                "ts_to": "201600"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "serviceids": [
            "4"
        ]
    },
    "id": 1
}
```

Source

CService::update() in *ui/include/classes/api/services/CService.php*.

Settings

This class is designed to work with common administration settings.

Object references:

- Settings

Available methods:

- settings.get - retrieve settings

- settings.update - update settings

Settings object

The following objects are directly related to the `settings` API.

Settings

The settings object has the following properties.

Property	Type	Description
default_lang	string	System language by default. Default: `en_US`.
default_timezone	string	System time zone by default. Default: `system` - system default. For the full list of supported time zones please refer to PHP documentation.
default_theme	string	Default theme. Possible values: `blue-theme` - *(default)* Blue; `blue-classic-theme` - Blue (classic); `dark-theme` - Dark; `dark-blue-theme` - Dark blue; `dark-classic-theme` - Dark (classic); `hc-light` - High-contrast light; `hc-dark` - High-contrast dark.
search_limit	integer	Limit for search and filter results. Default: 1000.
max_overview_table_size	integer	Max number of columns and rows in Data overview and Trigger overview dashboard widgets. Default: 50.
max_in_table	integer	Max count of elements to show inside table cell. Default: 50.
server_check_interval	integer	Show warning if Zabbix server is down. Possible values: 0 - Do not show warning; 10 - *(default)* Show warning.
work_period	string	Working time. Default: 1-5,09:00-18:00.
show_technical_errors	integer	Show technical errors (PHP/SQL) to non-Super admin users and to users that are not part of user groups with debug mode enabled. Possible values: 0 - *(default)* Do not technical errors; 1 - Show technical errors.
history_period	string	Max period to display history data in Latest data, Web, and Data overview dashboard widgets. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., `30s`, `1m`, `2h`, `1d`). Default: 24h.

Property	Type	Description
period_default	string	Time filter default period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d, 1M, 1y), including support for months and years. Default: 1h.
max_period	string	Max period for time filter. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h, 1d, 1M, 1y), including support for months and years. Default: 2y.
severity_color_0	string	Color for ”Not classified” severity as a hexadecimal color code. Default: 97AAB3.
severity_color_1	string	Color for ”Information” severity as a hexadecimal color code. Default: 7499FF.
severity_color_2	string	Color for ”Warning” severity as a hexadecimal color code. Default: FFC859.
severity_color_3	string	Color for ”Average” severity as a hexadecimal color code. Default: FFA059.
severity_color_4	string	Color for ”High” severity as a hexadecimal color code. Default: E97659.
severity_color_5	string	Color for ”Disaster” severity as a hexadecimal color code. Default: E45959.
severity_name_0	string	Name for ”Not classified” severity. Default: Not classified.
severity_name_1	string	Name for ”Information” severity. Default: Information.
severity_name_2	string	Name for ”Warning” severity. Default: Warning.
severity_name_3	string	Name for ”Average” severity. Default: Average.
severity_name_4	string	Name for ”High” severity. Default: High.
severity_name_5	string	Name for ”Disaster” severity. Default: Disaster.
custom_color	integer	Use custom event status colors. Possible values: 0 - *(default)* Do not use custom event status colors; 1 - Use custom event status colors.
ok_period	string	Display OK triggers period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h). Default: 5m.

Property	Type	Description
blink_period	string	Blink period on change of trigger state. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h). Default: 2m.
problem_unack_color	string	Color for unacknowledged PROBLEM events as a hexadecimal color code. Default: CC0000.
problem_ack_color	string	Color for acknowledged PROBLEM events as a hexadecimal color code. Default: CC0000.
ok_unack_color	string	Color for unacknowledged RESOLVED events as a hexadecimal color code. Default: 009900.
ok_ack_color	string	Color for acknowledged RESOLVED events as a hexadecimal color code. Default: 009900.
problem_unack_style	integer	Blinking for unacknowledged PROBLEM events. Possible values: 0 - Do not show blinking; 1 - *(default)* Show blinking.
problem_ack_style	integer	Blinking for acknowledged PROBLEM events. Possible values: 0 - Do not show blinking; 1 - *(default)* Show blinking.
ok_unack_style	integer	Blinking for unacknowledged RESOLVED events. Possible values: 0 - Do not show blinking; 1 - *(default)* Show blinking.
ok_ack_style	integer	Blinking for acknowledged RESOLVED events. Possible values: 0 - Do not show blinking; 1 - *(default)* Show blinking.
url	string	Frontend URL.
discovery_groupid	ID	ID of the host group to which will be automatically placed discovered hosts.
default_inventory_mode	integer	Default host inventory mode. Possible values: -1 - *(default)* Disabled; 0 - Manual; 1 - Automatic.
alert_usrgrpid	ID	ID of the user group to which the database down alarm message will be sent. If set to ”0”, the alarm message will not be sent.
snmptrap_logging	integer	Log unmatched SNMP traps. Possible values: 0 - Do not log unmatched SNMP traps; 1 - *(default)* Log unmatched SNMP traps.
login_attempts	integer	Number of failed login attempts after which login form will be blocked. Default: 5.

Property	Type	Description
login_block	string	Time interval during which login form will be blocked if number of failed login attempts exceeds defined in login_attempts field. Accepts a number of seconds (e.g., 1, 60), or a value with a time suffix (e.g., 30s, 1m). Default: 30s.
validate_uri_schemes	integer	Validate URI schemes. Possible values: 0 - Do not validate; 1 - *(default)* Validate.
uri_valid_schemes	string	Valid URI schemes. Default: http,https,ftp,file,mailto,tel,ssh.
x_frame_options	string	X-Frame-Options HTTP header. Default: SAMEORIGIN.
iframe_sandboxing_enabled	integer	Use iframe sandboxing. Possible values: 0 - Do not use; 1 - *(default)* Use.
iframe_sandboxing_exceptions	string	Iframe sandboxing exceptions.
connect_timeout	string	Connection timeout with Zabbix server. Possible values range: 1-30s. Default: 3s. Property behavior: - *required*
socket_timeout	string	Network default timeout. Possible values range: 1-300s. Default: 3s. Property behavior: - *required*
media_type_test_timeout	string	Network timeout for media type test. Possible values range: 1-300s. Default: 65s. Property behavior: - *required*
item_test_timeout	string	Network timeout for item tests. Possible value range: 1-600s. Default: 60s. Property behavior: - *required*

Property	Type	Description
script_timeout	string	Network timeout for script execution. Possible values range: 1-300s. Default: 60s. Property behavior: - *required*
report_test_timeout	string	Network timeout for scheduled report test. Possible values range: 1-300s. Default: 60s. Property behavior: - *required*
device_link_timeout	string	Timeout for adding a mobile device for use with Zabbix mobile app; applies to communication between the Zabbix API and bridge adapter. Possible values range: 1-300s. Default: 60s. Property behavior: - *required*
auditlog_enabled	integer	Whether to enable audit logging. Possible values: 0 - Disable; 1 - *(default)* Enable.
auditlog_mode	integer	Whether to enable audit logging of low-level discovery, network discovery and autoregistration activities performed by the server (System user). Possible values: 0 - Disable; 1 - *(default)* Enable.
ha_failover_delay	string	Failover delay in seconds. Default: 1m. Property behavior: - *read-only*
geomaps_tile_provider	string	Geomap tile provider. Possible values: `OpenStreetMap.Mapnik` - *(default)* OpenStreetMap Mapnik; `OpenTopoMap` - OpenTopoMap; `Stamen.TonerLite` - Stamen Toner Lite; `Stamen.Terrain` - Stamen Terrain; `USGS.USTopo` - USGS US Topo; `USGS.USImagery` - USGS US Imagery. Supports empty string to specify custom values of `geomaps_tile_url`, `geomaps_max_zoom` and `geomaps_attribution`.
geomaps_tile_url	string	Geomap tile URL. Property behavior: - *supported* if `geomaps_tile_provider` is set to empty string

Property	Type	Description
geomaps_max_zoom	integer	Geomap max zoom level. Possible values: 0-30. Property behavior: - *supported* if `geomaps_tile_provider` is set to empty string
geomaps_attribution	string	Geomap attribution text. Property behavior: - *supported* if `geomaps_tile_provider` is set to empty string
vault_provider	integer	Vault provider. Possible values: 0 - *(default)* HashiCorp Vault; 1 - CyberArk Vault.
proxy_secrets_provider	integer	Resolve secret macro values by: Possible values: 0 - *(default)* server only; 1 - server and proxies independently.
timeout_zabbix_agent timeout_simple_check timeout_snmp_agent timeout_external_check timeout_db_monitor timeout_http_agent timeout_ssh_agent timeout_telnet_agent timeout_script timeout_browser timeout_telemetry_query	string	Spend no more than `timeout_*` seconds on processing. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., `30s`, `1m`). Also accepts user macros. Possible values range: 1-600s. Default: 3s. Default for `timeout_browser`: 60s. Property behavior: - *required*
apm_global_db	object	Global data source configuration.

Global data source

The global data source object has the following properties:

Property	Type	Description
status	integer	Global data source status. Possible values: 0 - *(default)* not configured; 1 - configured.
url	string	URL used to connect to the data source. Property behavior: - *required* if user changes `status` to 1 (configured).

Property	Type	Description
authentication_type	integer	Type of authentication. Possible values: 0 - *(default)* using a username and password; 1 - no authentication used. Property behavior: - *required* if user changes `status` to 1 (configured).
username	string	Username used to log in to the data source. Property behavior: - *supported* if `status` is set to 1 (configured) and `authentication_type` is set to 0 (username and password). - *required* if user changes `authentication_type` to 0 (username and password).
password	string	Password used to log in to the data source. The value of this parameter can be empty. Property behavior: - *supported* if `status` is set to 1 (configured) and `authentication_type` is set to 0 (username and password).
db	string	Database name. The value of this parameter can be empty. Property behavior: - *supported* if `status` is set to 1 (configured).
ssl_verify_peer	integer	Verify server certificate. Possible values: 0 - *(default)* disabled; 1 - enabled. Property behavior: - *supported* if `status` is set to 1 (configured) and `url` uses the `https://` scheme.
ssl_verify_host	integer	Verify host. Possible values: 0 - disabled; 1 - *(default)* enabled. Property behavior: - *supported* if `status` is set to 1 (configured) and `url` uses the `https://` scheme.

settings.get

Description

`object settings.get(object parameters)`

The method allows you to retrieve settings object according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports only one parameter.

Parameter	Type	Description
output	query	This parameter is described in the reference commentary.

Return values

`(object)` Returns settings object.

Examples

Request:

```
{
    "jsonrpc": "2.0",
    "method": "settings.get",
    "params": {
        "output": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "alert_usrgrpid": "7",
        "apm_global_db": {
            "status": 1,
            "url": "https://localhost:8443",
            "authentication_type": 0,
            "username": "zabbix",
            "db": "",
            "ssl_verify_peer": 1,
            "ssl_verify_host": 1
        },
        "auditlog_enabled": "1",
        "auditlog_mode": "1",
        "blink_period": "2m",
        "connect_timeout": "3s",
        "custom_color": "0",
        "default_inventory_mode": "-1",
        "default_lang": "en_US",
        "default_theme": "blue-theme",
        "default_timezone": "system",
        "device_link_timeout": "60s",
        "discovery_groupid": "5",
        "geomaps_attribution": "",
        "geomaps_max_zoom": "0",
        "geomaps_tile_provider": "OpenStreetMap.Mapnik",
        "geomaps_tile_url": "",
        "ha_failover_delay": "1m",
        "history_period": "24h",
        "iframe_sandboxing_enabled": "1",
        "iframe_sandboxing_exceptions": "",
        "item_test_timeout": "60s",
        "login_attempts": "5",
        "login_block": "30s",
        "max_in_table": "50",
        "max_overview_table_size": "50",
        "max_period": "2y",
        "media_type_test_timeout": "65s",
        "ok_ack_color": "009900",
```

```
        "ok_ack_style": "1",
        "ok_period": "5m",
        "ok_unack_color": "009900",
        "ok_unack_style": "1",
        "period_default": "1h",
        "problem_ack_color": "CC0000",
        "problem_ack_style": "1",
        "problem_unack_color": "CC0000",
        "problem_unack_style": "1",
        "proxy_secrets_provider": "0",
        "report_test_timeout": "60s",
        "script_timeout": "60s",
        "search_limit": "1000",
        "server_check_interval": "10",
        "serverid": "019f4707-1a3e-70e6-aca2-286a2b696358",
        "severity_color_0": "97AAB3",
        "severity_color_1": "7499FF",
        "severity_color_2": "FFC859",
        "severity_color_3": "FFA059",
        "severity_color_4": "E97659",
        "severity_color_5": "E45959",
        "severity_name_0": "Not classified",
        "severity_name_1": "Information",
        "severity_name_2": "Warning",
        "severity_name_3": "Average",
        "severity_name_4": "High",
        "severity_name_5": "Disaster",
        "show_technical_errors": "0",
        "snmptrap_logging": "1",
        "socket_timeout": "3s",
        "timeout_browser": "60s",
        "timeout_db_monitor": "3s",
        "timeout_external_check": "3s",
        "timeout_http_agent": "3s",
        "timeout_script": "3s",
        "timeout_simple_check": "3s",
        "timeout_snmp_agent": "3s",
        "timeout_ssh_agent": "3s",
        "timeout_telemetry_query": "3s",
        "timeout_telnet_agent": "3s",
        "timeout_zabbix_agent": "3s",
        "uri_valid_schemes": "http,https,ftp,file,mailto,tel,ssh",
        "url": "",
        "validate_uri_schemes": "1",
        "vault_provider": "0",
        "work_period": "1-5,09:00-18:00",
        "x_frame_options": "SAMEORIGIN"
    },
    "id": 1
}
```


Source

CSettings::get() in *ui/include/classes/api/services/CSettings.php*.

settings.update

Description

`object settings.update(object settings)`

This method allows you to update existing common settings.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Settings properties to be updated.

Return values

`(array)` Returns an array with the names of updated parameters.

Examples

Request:

```
{
    "jsonrpc": "2.0",
    "method": "settings.update",
    "params": {
        "login_attempts": "1",
        "login_block": "1m"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        "login_attempts",
        "login_block"
    ],
    "id": 1
}
```

Source

CSettings::update() in *ui/include/classes/api/services/CSettings.php*.

SLA

This class is designed to work with SLA (Service Level Agreement) objects used to estimate the performance of IT infrastructure and business services.

Object references:

- SLA
- SLA schedule
- SLA excluded downtime
- SLA service tag

Available methods:

- sla.create - create new SLAs
- sla.delete - delete SLAs
- sla.get - retrieve SLAs
- sla.getsli - retrieve Service Level Indicator (SLI) data for SLAs
- sla.update - update SLAs

SLA object

The following objects are directly related to the `sla` (Service Level Agreement) API.

SLA

The SLA object has the following properties.

Property	Type	Description
slaid	ID	ID of the SLA. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the SLA. Property behavior: - *required* for create operations
period	integer	Reporting period of the SLA. Possible values: 0 - daily; 1 - weekly; 2 - monthly; 3 - quarterly; 4 - annually. Property behavior: - *required* for create operations
slo	float	Minimum acceptable Service Level Objective expressed as a percent. If the Service Level Indicator (SLI) drops lower, the SLA is considered to be in problem/unfulfilled state. Possible values: 0-100 (up to 4 fractional digits). Property behavior: - *required* for create operations
effective_date	integer	Effective date of the SLA. Possible values: date timestamp in UTC.
timezone	string	Reporting time zone, for example: `Europe/London`, `UTC`. For the full list of supported time zones please refer to PHP documentation. Property behavior: - *required* for create operations
status	integer	Status of the SLA. Possible values: 0 - *(default)* disabled SLA; 1 - enabled SLA.
description	string	Description of the SLA.

SLA schedule

The SLA schedule object defines periods where the connected service(s) are scheduled to be in working order. It has the following properties.

Property	Type	Description
period_from	integer	Starting time of the recurrent weekly period of time (inclusive). Possible values: number of seconds (counting from Sunday). Property behavior: - *required*

Property	Type	Description
period_to	integer	Ending time of the recurrent weekly period of time (exclusive). Possible values: number of seconds (counting from Sunday). Property behavior: - *required*

SLA excluded downtime

The excluded downtime object defines periods where the connected service(s) are scheduled to be out of working order, without affecting SLI, e.g., undergoing planned maintenance. It has the following properties.

Property	Type	Description
name	string	Name of the excluded downtime. Property behavior: - *required*
period_from	integer	Starting time of the excluded downtime (inclusive). Possible values: timestamp. Property behavior: - *required*
period_to	integer	Ending time of the excluded downtime (exclusive). Possible values: timestamp. Property behavior: - *required*

SLA service tag

The SLA service tag object links services to include in the calculations for the SLA. It has the following properties.

Property	Type	Description
tag	string	SLA service tag name. Property behavior: - *required*
operator	integer	SLA service tag operator. Possible values: 0 - *(default)* Equals; 2 - Contains.
value	string	SLA service tag value.

sla.create

Description

`object sla.create(object/array SLAs)`

This method allows you to create new SLA objects.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

(object/array) SLA objects to create.

Additionally to the standard SLA properties, the method accepts the following parameters.

Parameter	Type	Description
service_tags	array	SLA service tags to be created for the SLA. Parameter behavior: - *required*
schedule	array	SLA schedule to be created for the SLA. Specifying an empty parameter will be interpreted as a 24x7 schedule. Default: 24x7 schedule.
excluded_downtimes	array	SLA excluded downtimes to be created for the SLA.

Return values

(object) Returns an object containing the IDs of the created SLAs under the slaids property. The order of the returned IDs matches the order of the passed SLAs.

Examples

Creating an SLA

Instruct to create an SLA entry for: * tracking uptime for SQL-engine related services; * custom schedule of all weekdays excluding last hour on Saturday; * an effective date of the last day of the year 2022; * with 1 hour and 15 minutes long planned downtime starting at midnight on the 4th of July; * SLA weekly report calculation will be on; * the minimum acceptable SLO will be 99.9995%.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "sla.create",
    "params": [
        {
            "name": "Database Uptime",
            "slo": "99.9995",
            "period": "1",
            "timezone": "America/Toronto",
            "description": "Provide excellent uptime for main database engines.",
            "effective_date": 1672444800,
            "status": 1,
            "schedule": [
                {
                    "period_from": 0,
                    "period_to": 601200
                }
            ],
            "service_tags": [
                {
                    "tag": "database",
                    "operator": "0",
                    "value": "mysql"
                },
                {
                    "tag": "database",
                    "operator": "0",
                    "value": "postgresql"
                }
            ],
            "excluded_downtimes": [
                {
                    "name": "Software version upgrade rollout",
                    "period_from": "1648760400",
                    "period_to": "1648764900"
                }
```

```
            ]
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "slaids": [
            "5"
        ]
    },
    "id": 1
}
```

Source

CSla::create() in *ui/include/classes/api/services/CSla.php*.

sla.delete

Description

`object sla.delete(array slaids)`

This method allows you to delete SLA entries.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the SLAs to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted SLAs under the `slaids` property.

Examples

Deleting multiple SLAs

Delete two SLA entries.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "sla.delete",
    "params": [
        "4",
        "5"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "slaids": [
            "4",
            "5"
        ]
```

```
    },
    "id": 1
}
```

Source

CSla::delete() in *ui/include/classes/api/services/CSla.php*.

sla.get

Description

`integer/array sla.get(object parameters)`

The method allows you to retrieve SLA objects according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
slaids	ID/array	Return only SLAs with the given IDs.
serviceids	ID/array	Return only SLAs matching the specific services.
selectSchedule	query	Return a `schedule` property with SLA schedules. Supports `count`.
selectExcludedDowntimes	query	Return an `excluded_downtimes` property with SLA excluded downtimes. Supports `count`.
selectServiceTags	query	Return a `service_tags` property with SLA service tags. Supports `count`.
sortfield	string/array	Sort the result by the given properties. Possible values: `slaid`, `name`, `period`, `slo`, `effective_date`, `timezone`, `status`, `description`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving all SLAs

Retrieve all data about all SLAs and their properties.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "sla.get",
    "params": {
        "output": "extend",
        "selectSchedule": ["period_from", "period_to"],
        "selectExcludedDowntimes": ["name", "period_from", "period_to"],
        "selectServiceTags": ["tag", "operator", "value"],
        "preservekeys": true
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "1": {
            "slaid": "1",
            "name": "Database Uptime",
            "period": "1",
            "slo": "99.9995",
            "effective_date": "1672444800",
            "timezone": "America/Toronto",
            "status": "1",
            "description": "Provide excellent uptime for main SQL database engines.",
            "service_tags": [
                {
                    "tag": "database",
                    "operator": "0",
                    "value": "mysql"
                },
                {
                    "tag": "database",
                    "operator": "0",
                    "value": "postgresql"
                }
            ],
            "schedule": [
                {
                    "period_from": "0",
                    "period_to": "601200"
                }
            ],
            "excluded_downtimes": [
                {
                    "name": "Software version upgrade rollout",
                    "period_from": "1648760400",
                    "period_to": "1648764900"
                }
            ]
        }
    },
    "id": 1
}
```

Source

CSla:get() in *ui/include/classes/api/services/CSla.php*.

sla.getsli

Description

`object sla.getsli(object parameters)`

This method allows you to calculate the Service Level Indicator (SLI) data for a Service Level Agreement (SLA).

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters containing the SLA ID, reporting periods and, optionally, the IDs of the services - to calculate the SLI for.

Parameter	Type	Description
slaid	ID	ID of the SLA to return availability information for. Parameter behavior: - *required*
period_from	timestamp	Starting timestamp (inclusive) to report the SLI for. Possible values: Unix timestamp.
period_to	timestamp	Ending timestamp (inclusive) to report the SLI for. Possible values: Unix timestamp.
periods	integer	Number of periods to report. Possible values: 1-100
serviceids	ID/array	IDs of services to return the SLI for.

Partitioning of periods

The following table demonstrates the arrangement of returned period slices based on combinations of parameters.

Note:
Returned periods do not precede the first available period (based on the effective date of the SLA) and do not exceed the current period.

Parameters			Returned periods
period_from	**period_to**	**periods**	
-	-	-	Last 20 periods, including the current one.
-	-	N	Last N periods.
-	specified	-	Last 20 periods before `period_to`.
-	specified	N	Last N periods before `period_to`.
specified	-	-	First 20 periods starting with `period_from`.
specified	-	N	First N periods starting with `period_from`.
specified	specified	-	Up to 100 periods in the specified range.
specified	specified	N	N periods in the specified range.

Return values

`(object)` Returns the results of the calculation.

Property	Type	Description
periods	array	List of the reported periods. Each reported period is represented as an object consisting of: - `period_from` - `(timestamp)` Starting timestamp (inclusive) of the reported period. - `period_to` - `(timestamp)` Ending timestamp (exclusive) of the reported period. Periods are sorted by `period_from`, with earliest periods appearing first.
serviceids	array	List of service IDs in the reported periods. The sorting order of the list is not defined. Even if `serviceids` parameter was passed to the `sla.getsli` method.
sli	array	SLI data (as a **two-dimensional array**) for each reported period and service. The index of the `periods` property is used as the **first** dimension of the `sli` property. The index of the `serviceids` property is used as the **second** dimension of the `sli` property.

SLI data

The SLI data returned for each reported period and service consists of:

Property	Type	Description
uptime	integer	Amount of time service spent in an *OK* state during scheduled uptime, less the excluded downtimes.
downtime	integer	Amount of time service spent in a *not OK* state during scheduled uptime, less the excluded downtimes.
sli	float	SLI (per cent of total uptime), based on uptime and downtime.
error_budget	integer	Error budget (in seconds), based on the SLI and the SLO.
excluded_downtimes	array	Array of excluded downtimes in this reporting period. Each object will contain the following parameters: - `name` - `(string)` Name of the excluded downtime. - `period_from` - `(timestamp)` Starting timestamp (inclusive) of the excluded downtime. - `period_to` - `(timestamp)` Ending timestamp (exclusive) of the excluded downtime. Excluded downtimes are sorted by `period_from`, with earliest periods appearing first.

Examples

Calculating SLI for a daily SLA

Retrieve SLI data for services with IDs "1" and "4" that are linked to the SLA with ID "1". Retrieve data for a single period until "1761861599" (Oct 30 2025 23:59:59 GMT+0200). Since the reporting period of the SLA is daily, SLI data is retrieved from "1761775200" (Oct 30 2025 00:00:00 GMT+0200) to "1761861600" (Oct 31 2025 00:00:00 GMT+0200).

Request:

```
{
    "jsonrpc": "2.0",
    "method": "sla.getsli",
    "params": {
        "slaid": "1",
        "serviceids": [
```

```
            1,
            4
        ],
        "periods": 1,
        "period_to": 1761861599
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "periods": [
            {
                "period_from": 1761775200,
                "period_to": 1761861600
            }
        ],
        "serviceids": [
            1,
            4
        ],
        "sli": [
            [
                {
                    "uptime": 43843,
                    "downtime": 0,
                    "sli": 100,
                    "error_budget": 0,
                    "excluded_downtimes": [
                        {
                            "name": "Maintenance OCT",
                            "period_from": 1761825600,
                            "period_to": 1761829200
                        }
                    ]
                },
                {
                    "uptime": 32225,
                    "downtime": 0,
                    "sli": 100,
                    "error_budget": 0,
                    "excluded_downtimes": []
                }
            ]
        ]
    },
    "id": 1
}
```

Calculating SLI for a monthly SLA

Retrieve SLI data on services with IDs ”50”, ”60” and ”70” that are linked to the SLA with ID ”5”. Retrieve data for three periods starting from ”1635724800” (Nov 01 2021 00:00:00 UTC). Since the reporting period of the SLA is monthly, SLI data is retrieved for the following three months:

- From ”1635724800” (Nov 01 2021 00:00:00 UTC) to ”1638316800” (Dec 01 2021 00:00:00 UTC)
- From ”1638316800” (Dec 01 2021 00:00:00 UTC) to ”1640995200” (Jan 01 2022 00:00:00 UTC)
- From ”1640995200” (Jan 01 2022 00:00:00 UTC) to ”1643673600” (Feb 01 2022 00:00:00 UTC)

Request:

```
{
    "jsonrpc": "2.0",
```

```
    "method": "sla.getsli",
    "params": {
        "slaid": "5",
        "serviceids": [
            50,
            60,
            70
        ],
        "periods": 3,
        "period_from": 1635724800
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "periods": [
            {
                "period_from": 1635724800,
                "period_to": 1638316800
            },
            {
                "period_from": 1638316800,
                "period_to": 1640995200
            },
            {
                "period_from": 1640995200,
                "period_to": 1643673600
            }
        ],
        "serviceids": [
            50,
            60,
            70
        ],
        "sli": [
            [
                {
                    "uptime": 1186212,
                    "downtime": 0,
                    "sli": 100,
                    "error_budget": 0,
                    "excluded_downtimes": [
                        {
                            "name": "Maintenance Nov 25 - Dec 01",
                            "period_from": 1637836212,
                            "period_to": 1638316800
                        }
                    ]
                },
                {
                    "uptime": 1186212,
                    "downtime": 0,
                    "sli": 100,
                    "error_budget": 0,
                    "excluded_downtimes": [
                        {
                            "name": "Maintenance Nov 25 - Dec 01",
                            "period_from": 1637836212,
                            "period_to": 1638316800
```

```
                    }
                ]
            },
            {
                "uptime": 1186212,
                "downtime": 0,
                "sli": 100,
                "error_budget": 0,
                "excluded_downtimes": [
                    {
                        "name": "Maintenance Nov 25 - Dec 01",
                        "period_from": 1637836212,
                        "period_to": 1638316800
                    }
                ]
            }
        ],
        [
            {
                "uptime": 1147548,
                "downtime": 0,
                "sli": 100,
                "error_budget": 0,
                "excluded_downtimes": [
                    {
                        "name": "Maintenance Dec 02 - Dec 10",
                        "period_from": 1638439200,
                        "period_to": 1639109652
                    }
                ]
            },
            {
                "uptime": 1147548,
                "downtime": 0,
                "sli": 100,
                "error_budget": 0,
                "excluded_downtimes": [
                    {
                        "name": "Maintenance Dec 02 - Dec 10",
                        "period_from": 1638439200,
                        "period_to": 1639109652
                    }
                ]
            },
            {
                "uptime": 1147548,
                "downtime": 0,
                "sli": 100,
                "error_budget": 0,
                "excluded_downtimes": [
                    {
                        "name": "Maintenance Dec 02 - Dec 10",
                        "period_from": 1638439200,
                        "period_to": 1639109652
                    }
                ]
            }
        ],
        [
            {
                "uptime": 1674000,
                "downtime": 0,
```

```
                    "sli": 100,
                    "error_budget": 0,
                    "excluded_downtimes": []
                },
                {
                    "uptime": 1674000,
                    "downtime": 0,
                    "sli": 100,
                    "error_budget": 0,
                    "excluded_downtimes": []
                },
                {
                    "uptime": 1674000,
                    "downtime": 0,
                    "sli": 100,
                    "error_budget": 0,
                    "excluded_downtimes": []
                }
            ]
        ]
    },
    "id": 1
}
```


Source

CSla::getSli() in *ui/include/classes/api/services/CSla.php*

sla.update

Description

`object sla.update(object/array slaids)`

This method allows you to update existing SLA entries.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` SLA properties to be updated.

The `slaid` property must be defined for each SLA, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard SLA properties, the method accepts the following parameters.

Parameter	Type	Description
service_tags	array	SLA service tags to replace the current SLA service tags.
schedule	array	SLA schedule to replace the current one. Specifying parameter as empty will be interpreted as a 24x7 schedule.
excluded_downtimes	array	SLA excluded downtimes to replace the current ones.

Return values

`(object)` Returns an object containing the IDs of the updated SLAs under the `slaids` property.

Examples

Updating service tags

Make SLA with ID "5" to be calculated at monthly intervals for NoSQL related services, without changing its schedule or excluded downtimes; set SLO to 95%.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "sla.update",
    "params": [
        {
            "slaid": "5",
            "name": "NoSQL Database engines",
            "slo": "95",
            "period": 2,
            "service_tags": [
                {
                    "tag": "database",
                    "operator": "0",
                    "value": "redis"
                },
                {
                    "tag": "database",
                    "operator": "0",
                    "value": "mongodb"
                }
            ]
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "slaids": [
            "5"
        ]
    },
    "id": 1
}
```

Changing the schedule of an SLA

Switch the SLA with ID ”5” to a 24x7 schedule.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "sla.update",
    "params": {
        "slaid": "5",
        "schedule": []
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "slaids": [
            "5"
        ]
    },
    "id": 1
}
```

Changing the excluded downtimes for an SLA

Add a planned 4 hour long RAM upgrade downtime on the 6th of April, 2022, while keeping (needs to be defined anew) a previously existing software upgrade planned on the 4th of July for the SLA with ID ”5”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "sla.update",
    "params": {
        "slaid": "5",
        "excluded_downtimes": [
            {
                "name": "Software version upgrade rollout",
                "period_from": "1648760400",
                "period_to": "1648764900"
            },
            {
                "name": "RAM upgrade",
                "period_from": "1649192400",
                "period_to": "1649206800"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "slaids": [
            "5"
        ]
    },
    "id": 1
}
```

Source

CSla::update() in *ui/include/classes/api/services/CSla.php*.

Task

This class is designed to work with tasks (such as checking items or low-level discovery rules without config reload).

Object references:

- Task
  - ’Execute now’ request object
  - ’Refresh proxy configuration’ request object
  - ’Diagnostic information’ request object
    - Statistic request object
    - Statistic result object

Available methods:

- task.create - create new tasks
- task.get - retrieve tasks

Task object

The following objects are directly related to the `task` API.

Task

The task object has the following properties:

Property	Type	Description
taskid	ID	ID of the task. Property behavior: - *read-only*
type	integer	Type of the task. Possible values: 1 - Diagnostic information; 2 - Refresh proxy configuration; 6 - Execute now. *Admin* and *User* type users may create only 'Execute now' tasks. Note that `task.get` always returns "7" (Task execution summary). Property behavior: - *required*
status	integer	Status of the task. Possible values: 1 - new task; 2 - task in progress; 3 - task is completed; 4 - task is expired. Property behavior: - *read-only*
clock	timestamp	Time when the task was created. Property behavior: - *read-only*
ttl	integer	The time in seconds after which task expires. Property behavior: - *read-only*
proxyid	ID	ID of the proxy about which diagnostic information statistic is collected. Property behavior: - *supported* if `type` is set to "Diagnostic information" or "Refresh proxy configuration"
request	object	Task request object according to the task type: Object of 'Execute now' task is described in detail below; Object of 'Refresh proxy configuration' task is described in detail below; Object of 'Diagnostic information' task is described in detail below. Property behavior: - *required*
result	object	Result object of the diagnostic information task. May contain NULL if result is not yet ready. Result object is described in detail below. Property behavior: - *read-only*

'Execute now' request object

The 'Execute now' task request object has the following properties.

Property	Type	Description
itemid	ID	ID of item and low-level discovery rules. *Admin* and *User* type users may 'Execute now' items on hosts for which they have *read-write* permission, or *read* permission and `invoke_execute_now` action enabled for their role. The same applies to *Admin* type users for low-level discovery (LLD) rules.

'Refresh proxy configuration' request object

The 'Refresh proxy configuration' task request object has the following properties.

Property	Type	Description
proxyids	array	Proxy IDs.

'Diagnostic information' request object

The diagnostic information task request object has the following properties. Statistic request object for all types of properties is described in detail below.

Property	Type	Description
historycache	object	History cache statistic request. Available on server and proxy.
valuecache	object	Items cache statistic request. Available on server.
preprocessing	object	Preprocessing manager statistic request. Available on server and proxy.
alerting	object	Alert manager statistic request. Available on server.
lld	object	LLD manager statistic request. Available on server.

Statistic request object

Statistic request object is used to define what type of information should be collected about server/proxy internal processes. It has the following properties.

Property	Type	Description
stats	query	Statistic object properties to be returned. The list of available fields for each type of diagnostic information statistic are described in detail below. Default: `extend` will return all available statistic fields.
top	object	Object to sort and limit returned statistic values. The list of available fields for each type of diagnostic information statistic are described in detail below. Example: { "source.alerts": 10 }

List of statistic fields available for each type of diagnostic information request

Following statistic fields can be requested for each type of diagnostic information request property.

Diagnostic type	Available fields	Description
historycache	items	Number of cached items.
	values	Number of cached values.
	memory	Shared memory statistics (free space, number of used chunks, number of free chunks, max size of free chunk).
	memory.data	History data cache shared memory statistics.
	memory.index	History index cache shared memory statistics.
valuecache	items	Number of cached items.
	values	Number of cached values.
	memory	Shared memory statistics (free space, number of used chunks, number of free chunks, max size of free chunk).

Diagnostic type	Available fields	Description
	mode	Value cache mode.
preprocessing	values	Number of queued values.
	preproc.values	Number of queued values with preprocessing steps.
alerting	alerts	Number of queued alerts.
lld	rules	Number of queued rules.
	values	Number of queued values.

List of sorting fields available for each type of diagnostic information request

Following statistic fields can be used to sort and limit requested information.

Diagnostic type	Available fields	Type
historycache	values	integer
valuecache	values	integer
	request.values	integer
preprocessing	values	integer
alerting	media.alerts	integer
	source.alerts	integer
lld	values	integer

Statistic result object

Statistic result object is retrieved in `result` field of task object.

Property	Type	Description
status	integer	Status of the task result. Possible values: -1 - error occurred during performing task; 0 - task result is created. Property behavior: - *read-only*
data	string/object	Results according the statistic request object of particular diagnostic information task. Contains error message string if error occurred during performing task.

task.create

Description

`object task.create(object/array tasks)`

This method allows you to create tasks.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` A task to create.

The method accepts tasks with the standard task properties.

Note that 'Execute now' tasks can be created only for the following types of items/discovery rules:

- Zabbix agent (passive)
- Simple check
- SNMP agent (v1/v2/v3)
- Zabbix internal

- External check
- Database monitor
- HTTP agent
- IPMI agent
- SSH agent
- TELNET agent
- JMX agent
- Calculated
- Dependent item
- Script
- Browser

If the item/discovery rule is of type ”Dependent item”, then its master item must also be one of the above types.

Return values

`(object)` Returns an object containing the IDs of the created tasks under the `taskids` property. One task is created for each item and low-level discovery rule. The order of the returned IDs matches the order of the passed `itemids`.

Examples

Creating a task

Create an ’Execute now’ task for an item and a low-level discovery rule.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "task.create",
    "params": [
        {
            "type": 6,
            "request": {
                "itemid": "10092"
            }
        },
        {
            "type": 6,
            "request": {
                "itemid": "10093"
            }
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "taskids": [
            "1",
            "2"
        ]
    },
    "id": 1
}
```

Create a ’Refresh proxy configuration’ task for two proxies.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "task.create",
    "params": [
        {
```

```
            "type": 2,
            "request": {
                "proxyids": ["10459", "10460"]
            }
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "taskids": [
            "1"
        ]
    },
    "id": 1
}
```

Create a 'Diagnostic information' task.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "task.create",
    "params": [
        {
            "type": 1,
            "request": {
                "alerting": {
                    "stats": [
                        "alerts"
                    ],
                    "top": {
                        "media.alerts": 10
                    }
                },
                "lld": {
                    "stats": "extend",
                    "top": {
                        "values": 5
                    }
                }
            },
            "proxyid": 0
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "taskids": [
            "3"
        ]
    },
    "id": 1
}
```

See also

- Task
- 'Execute now' request object
- 'Refresh proxy configuration' request object
- 'Diagnostic information' request object

Source

CTask::create() in *ui/include/classes/api/services/CTask.php*.

task.get

Description

`integer/array task.get(object parameters)`

The method allows you to retrieve tasks according to the given parameters.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
taskids	ID/array	Return only tasks with the given IDs.
output	query	These parameters are described in the reference commentary.
preservekeys	boolean	

Return values

`(integer/array)` Returns an array of objects.

Examples

Retrieve task by ID

Retrieve all the data about the task with the ID "1".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "task.get",
    "params": {
        "output": "extend",
        "taskids": "1"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "taskid": "1",
            "type": "7",
            "status": "3",
            "clock": "1601039076",
            "ttl": "3600",
            "proxyid": null,
            "request": {
                "alerting": {
```

```
                    "stats": [
                        "alerts"
                    ],
                    "top": {
                        "media.alerts": 10
                    }
                },
                "lld": {
                    "stats": "extend",
                    "top": {
                        "values": 5
                    }
                }
            },
            "result": {
                "data": {
                    "alerting": {
                        "alerts": 0,
                        "top": {
                            "media.alerts": []
                        },
                        "time": 0.000663
                    },
                    "lld": {
                        "rules": 0,
                        "values": 0,
                        "top": {
                            "values": []
                        },
                        "time": 0.000442
                    }
                },
                "status": "0"
            }
        }
    ],
    "id": 1
}
```


See also

- Task
- 'Execute now' request object
- 'Refresh proxy configuration' request object
- 'Diagnostic information' request object

Source

CTask::get() in *ui/include/classes/api/services/CTask.php*.

Template

This class is designed to work with templates.

Object references:

- Template
- Template tag

Available methods:

- template.create - create new templates
- template.delete - delete templates
- template.get - retrieve templates
- template.massadd - add related objects to templates
- template.massremove - remove related objects from templates

- template.update - update templates

Template object

The following objects are directly related to the `template` API.

Template

The template object has the following properties.

Property	Type	Description
templateid	ID	ID of the template. Property behavior: - *read-only* - *required* for update operations
description	text	Description of the template.
host	string	Technical name of the template. Property behavior: - *required* for create operations
name	string	Visible name of the template. Default: `host` property value.
readme	text	Template-specific configuration instructions to display in the Host Wizard. Supports Markdown formatting.
uuid	string	Universal unique identifier, used for linking imported templates to already existing ones. Auto-generated, if not given.
vendor_name	string	Template vendor name. For create operations, both `vendor_name` and `vendor_version` should be either set or left empty. For update operations, `vendor_version` can be left empty if it has a value in the database.
vendor_version	string	Template vendor version. For create operations, both `vendor_name` and `vendor_version` should be either set or left empty. For update operations, `vendor_name` can be left empty if it has a value in the database.
wizard_ready	integer	Whether the template is available for selection in the Host Wizard. Possible values: 0 - *(default)* Not available. 1 - Available.

Template tag

The template tag object has the following properties.

Property	Type	Description
tag	string	Template tag name. Property behavior: - *required*
value	string	Template tag value. Property behavior: - *read-only*.

Property	Type	Description
object	integer	Type of object the tag was inherited from. Possible values: 0 - template. Property behavior: - *read-only*.
objectid	ID	ID of the object the tag was inherited from. Property behavior: - *read-only*.

template.create

Description

`object template.create(object/array templates)`

This method allows you to create new templates.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Templates to create.

Additionally to the standard template properties, the method accepts the following parameters.

Parameter	Type	Description
groups	object/array	Template groups to add the template to. The template groups must have only the `groupid` property defined. Parameter behavior: - *required*
tags	object/array	Template tags.
templates	object/array	Templates to be linked to the template. The templates must have only the `templateid` property defined.
macros	object/array	User macros to be created for the template.

Return values

`(object)` Returns an object containing the IDs of the created templates under the `templateids` property. The order of the returned IDs matches the order of the passed templates.

Examples

Creating a template

Create a template with tags and link two templates to this template.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.create",
    "params": {
        "host": "Linux template",
        "groups": {
            "groupid": 1
```

```
        },
        "templates": [
            {
                "templateid": "11115"
            },
            {
                "templateid": "11116"
            }
        ],
        "tags": [
            {
                "tag": "host-name",
                "value": "{HOST.NAME}"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templateids": [
            "11117"
        ]
    },
    "id": 1
}
```

Source

CTemplate::create() in *ui/include/classes/api/services/CTemplate.php*.

template.delete

Description

`object template.delete(array templateIds)`

This method allows you to delete templates.

Deleting a template also removes all of its entities (items, triggers, graphs, etc.) from any hosts or templates it is linked to. To keep these entities, first unlink the template from the relevant hosts or templates using the `host.update`, `host.massremove`, `template.update`, or `template.massremove` methods.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the templates to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted templates under the `templateids` property.

Examples

Deleting multiple templates

Delete two templates.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.delete",
    "params": [
        "13",
        "32"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templateids": [
            "13",
            "32"
        ]
    },
    "id": 1
}
```

Source

CTemplate::delete() in *ui/include/classes/api/services/CTemplate.php*.

template.get

Description

`integer/array template.get(object parameters)`

The method allows you to retrieve templates according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
templateids	ID/array	Return only templates with the given template IDs.
groupids	ID/array	Return only templates that belong to the given template groups.
parentTemplateids	ID/array	Return only templates that the given template is linked to.
hostids	ID/array	Return only templates that are linked to the given hosts/templates.
graphids	ID/array	Return only templates that contain the given graphs.
itemids	ID/array	Return only templates that contain the given items.
triggerids	ID/array	Return only templates that contain the given triggers.
with_items	flag	Return only templates that have items.
with_triggers	flag	Return only templates that have triggers.
with_graphs	flag	Return only templates that have graphs.
with_httptests	flag	Return only templates that have web scenarios.
evaltype	integer	Tag evaluation method. Possible values: 0 - *(default)* And/Or; 2 - Or.

Parameter	Type	Description
tags	object/array	Return only templates with the given tags. Format: `[{"tag": "<tag>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all templates. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
inheritedTags	boolean	Return templates that have given `tags` also in linked templates. Possible values: `true` - linked templates must also have given tags; `false` - *(default)* tags from linked templates are ignored.
selectInheritedTags	query	Return an `inheritedTags` property with tags that are on linked templates.
selectTags	query	Return template tags in the `tags` property.
selectDiscoveryRules	query	Return a `discoveryRules` property with template LLD rules. Supports `count`.
selectHosts	query	Return the hosts that are linked to the template in the `hosts` property. Supports `count`.
selectTemplateGroups	query	Return the template groups that the template belongs to in the `templategroups` property.
selectTemplates	query	Return templates to which the given template is linked in the `templates` property. Supports `count`.
selectParentTemplates	query	Return templates that are linked to the given template in the `parentTemplates` property. Supports `count`.
selectHttpTests	query	Return the web scenarios from the template in the `httpTests` property. Supports `count`.
selectItems	query	Return items from the template in the `items` property. Supports `count`.
selectTriggers	query	Return triggers from the template in the `triggers` property. Supports `count`.
selectGraphs	query	Return graphs from the template in the `graphs` property. Supports `count`.
selectMacros	query	Return the macros from the template in the `macros` property.
selectDashboards	query	Return dashboards from the template in the `dashboards` property. Supports `count`.
selectValueMaps	query	Return a `valuemaps` property with template value maps.

Parameter	Type	Description
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectTemplates` - results will be sorted by `name`; `selectHosts` - sorted by `host`; `selectParentTemplates` - sorted by `host`; `selectItems` - sorted by `name`; `selectDiscoveryRules` - sorted by `name`; `selectTriggers` - sorted by `description`; `selectGraphs` - sorted by `name`; `selectDashboards` - sorted by `name`.
sortfield	string/array	Sort the result by the given properties. Possible values: `hostid`, `host`, `name`, `status`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	
selectDiscoveries	query	Return low-level discoveries from the template in the `discoveries` property. Supports `count`. This query is **deprecated**, please use `selectDiscoveryRules` instead.

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving templates by name

Retrieve all data about two templates named ”Linux by Zabbix agent” and ”Windows by Zabbix agent”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.get",
    "params": {
        "output": "extend",
        "filter": {
            "host": [
                "Linux by Zabbix agent",
                "Windows by Zabbix agent"
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "proxyid": "0",
            "host": "Linux by Zabbix agent",
            "status": "3",
            "ipmi_authtype": "-1",
            "ipmi_privilege": "2",
            "ipmi_username": "",
            "ipmi_password": "",
            "maintenanceid": "0",
            "maintenance_status": "0",
            "maintenance_type": "0",
            "maintenance_from": "0",
            "name": "Linux by Zabbix agent",
            "flags": "0",
            "templateid": "10001",
            "description": "This is an official Linux template. It requires Zabbix agent 8.0 or newer.\r\n
            "tls_connect": "1",
            "tls_accept": "1",
            "tls_issuer": "",
            "tls_subject": "",
            "tls_psk_identity": "",
            "tls_psk": "",
            "custom_interfaces": "0",
            "uuid": "f8f7908280354f2abeed07dc788c3747",
            "vendor_name": "Zabbix",
            "vendor_version": "8.2-2",
            "proxy_groupid": "0",
            "monitored_by": "0",
            "wizard_ready": "1",
            "readme": "## Overview\r\n\r\nThis is an official Linux template. It requires Zabbix agent 8.0
        },
        {
            "proxyid": "0",
            "host": "Windows by Zabbix agent",
            "status": "3",
            "ipmi_authtype": "-1",
            "ipmi_privilege": "2",
            "ipmi_username": "",
            "ipmi_password": "",
            "maintenanceid": "0",
            "maintenance_status": "0",
            "maintenance_type": "0",
            "maintenance_from": "0",
            "name": "Windows by Zabbix agent",
            "flags": "0",
            "templateid": "10081",
            "description": "This is an official Windows template. It requires Zabbix agent 8.0 or newer.\r
            "tls_connect": "1",
            "tls_accept": "1",
            "tls_issuer": "",
            "tls_subject": "",
            "tls_psk_identity": "",
            "tls_psk": "",
            "custom_interfaces": "0",
            "uuid": "13b06904a6bf41cbb795e3193d896340",
            "vendor_name": "Zabbix",
            "vendor_version": "8.2-2",
            "proxy_groupid": "0",
            "monitored_by": "0",
```

```
            "wizard_ready": "1",
            "readme": "## Overview\r\n\r\nThis is an official Windows template. It requires Zabbix agent &
        }
    ],
    "id": 1
}
```

Retrieving template groups

Retrieve template groups that the template "Linux by Zabbix agent" is a member of.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.get",
    "params": {
        "output": ["hostid"],
        "selectTemplateGroups": "extend",
        "filter": {
            "host": [
                "Linux by Zabbix agent"
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "templateid": "10001",
            "templategroups": [
                {
                    "groupid": "10",
                    "name": "Templates/Operating systems",
                    "uuid": "846977d1dfed4968bc5f8bdb363285bc"
                }
            ]
        }
    ],
    "id": 1
}
```

Retrieving hosts by template

Retrieve hosts that have the "10001" (*Linux by Zabbix agent*) template linked to them.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.get",
    "params": {
        "output": "templateid",
        "templateids": "10001",
        "selectHosts": ["hostid", "name"]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
```

```
    "result": [
        {
            "templateid": "10001",
            "hosts": [
                {
                    "hostid": "10084",
                    "name": "Zabbix server"
                },
                {
                    "hostid": "10603",
                    "name": "Host 1"
                },
                {
                    "hostid": "10604",
                    "name": "Host 2"
                }
            ]
        }
    ],
    "id": 1
}
```

Searching by template tags

Retrieve templates that have tag "host-name" equal to "{HOST.NAME}".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.get",
    "params": {
        "output": ["hostid"],
        "selectTags": "extend",
        "evaltype": 0,
        "tags": [
            {
                "tag": "host-name",
                "value": "{HOST.NAME}",
                "operator": 1
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "templateid": "10402",
            "tags": [
                {
                    "tag": "host-name",
                    "value": "{HOST.NAME}"
                }
            ]
        }
    ],
    "id": 1
}
```

See also

- Template group
- Template
- User macro
- Host interface

Source

CTemplate::get() in *ui/include/classes/api/services/CTemplate.php*.

template.massadd

Description

`object template.massadd(object parameters)`

This method allows you to simultaneously add multiple related objects to the given templates.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters containing the IDs of the templates to update and the objects to add to the templates.

The method accepts the following parameters.

Parameter	Type	Description
templates	object/array	Templates to be updated. The templates must have only the `templateid` property defined. Parameter behavior: - *required*
groups	object/array	Template groups to add the given templates to. The template groups must have only the `groupid` property defined.
macros	object/array	User macros to be created for the given templates.
templates_link	object/array	Templates to link to the given templates. The templates must have only the `templateid` property defined.

Return values

`(object)` Returns an object containing the IDs of the updated templates under the `templateids` property.

Examples

Link a group to templates

Add template group "2" to two templates.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.massadd",
    "params": {
        "templates": [
            {
                "templateid": "10085"
            },
            {
                "templateid": "10086"
            }
        ],
        "groups": [
```

```
            {
                "groupid": "2"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templateids": [
            "10085",
            "10086"
        ]
    },
    "id": 1
}
```

Link two templates to a template

Link templates ”10106” and ”10104” to template ”10073”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.massadd",
    "params": {
        "templates": [
            {
                "templateid": "10073"
            }
        ],
        "templates_link": [
            {
                "templateid": "10106"
            },
            {
                "templateid": "10104"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templateids": [
            "10073"
        ]
    },
    "id": 1
}
```

See also

- template.update
- Host
- Template group
- User macro

Source

CTemplate::massAdd() in *ui/include/classes/api/services/CTemplate.php*.

template.massremove

Description

`object template.massremove(object parameters)`

This method allows you to remove related objects from multiple templates.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters containing the IDs of the templates to update and the objects that should be removed.

Parameter	Type	Description
templateids	ID/array	IDs of the templates to be updated. Parameter behavior: *- required*
groupids	ID/array	IDs of the template groups from which to remove the given templates.
macros	string/array	IDs of the user macros to delete from the given templates.
templateids_clear	ID/array	IDs of the templates to unlink and clear from the given templates (upstream).
templateids_link	ID/array	IDs of the templates to unlink from the given templates (upstream).

Return values

`(object)` Returns an object containing the IDs of the updated templates under the `templateids` property.

Examples

Removing templates from a group

Remove two templates from group "2".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.massremove",
    "params": {
        "templateids": [
            "10085",
            "10086"
        ],
        "groupids": "2"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templateids": [
            "10085",
            "10086"
        ]
    },
```

```
    "id": 1
}
```

Unlinking templates from a host

Unlink templates "10106" and "10104" from template "10085".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.massremove",
    "params": {
        "templateids": "10085",
        "templateids_link": [
            "10106",
            "10104"
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templateids": [
            "10085"
        ]
    },
    "id": 1
}
```

See also

- template.update
- User macro

Source

CTemplate::massRemove() in *ui/include/classes/api/services/CTemplate.php*.

template.update

Description

`object template.update(object/array templates)`

This method allows you to update existing templates.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Template properties to be updated.

The `templateid` property must be defined for each template, all other properties are optional. Only the given properties will be updated, all others will remain unchanged.

Additionally to the standard template properties, the method accepts the following parameters.

Parameter	Type	Description
groups	object/array	Template groups to replace the current template groups the templates belong to. The template groups must have only the `groupid` property defined.
tags	object/array	Template tags to replace the current template tags.
macros	object/array	User macros to replace the current user macros on the given templates.
templates	object/array	Templates to replace the currently linked templates. Templates that are not passed are only unlinked. The templates must have only the `templateid` property defined.
templates_clear	object/array	Templates to unlink and clear from the given templates. The templates must have only the `templateid` property defined.

Return values

`(object)` Returns an object containing the IDs of the updated templates under the `templateids` property.

Examples

Changing the standard template properties

Change the technical name of the template to ”Linux by Zabbix agent Custom”, the visible name to ”My template”, and update the template description.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.update",
    "params": {
        "templateid": "10086",
        "host": "Linux by Zabbix agent Custom",
        "name": "My template",
        "description": "This is a custom Linux template."
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templateids": [
            "10086"
        ]
    },
    "id": 1
}
```

Updating template groups

Replace all template groups for the given template with a different one.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.update",
    "params": {
        "templateid": "10086",
        "groups": [
            {
                "groupid": "24"
```

```
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templateids": [
            "10086"
        ]
    },
    "id": 1
}
```

Updating multiple template groups

Replace all template groups for multiple templates with different ones.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.update",
    "params": [
        {
            "templateid": "10086",
            "groups": [
                {
                    "groupid": "24"
                }
            ]
        },
        {
            "templateid": "10087",
            "groups": [
                {
                    "groupid": "1"
                },
                {
                    "groupid": "12"
                }
            ]
        },
        {
            "templateid": "10088",
            "groups": [
                {
                    "groupid": "1"
                },
                {
                    "groupid": "12"
                }
            ]
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
```

```
    "result": {
        "templateids": [
            "10086",
            "10087",
            "10088"
        ]
    },
    "id": 1
}
```

Updating template tags

Replace all template tags with a different one.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.update",
    "params": {
        "templateid": "10086",
        "tags": [
            {
                "tag": "host-name",
                "value": "{HOST.NAME}"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templateids": [
            "10086"
        ]
    },
    "id": 1
}
```

Updating template macros

Replace all template macros with a different one.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.update",
    "params": {
        "templateid": "10086",
        "macros": [
            {
                "macro": "{$MY_MACRO}",
                "value": "new_value"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
```

```
    "result": {
        "templateids": [
            "10086"
        ]
    },
    "id": 1
}
```

Updating multiple template macros

Replace all user macros with the given user macro on multiple templates.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.update",
    "params": [
            {
                "templateid": "10087",
                "macros": [
                    {
                        "macro": "{$AGENT.TIMEOUT}",
                        "value": "5m",
                        "description": "Timeout after which agent is considered unavailable. Works only fo
                    }
                ]
            },
            {
                "templateid": "10088",
                "macros": [
                    {
                        "macro": "{$AGENT.TIMEOUT}",
                        "value": "5m",
                        "description": "Timeout after which agent is considered unavailable. Works only fo
                    }
                ]
            }
        ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templateids": [
            "10087",
            "10088"
        ]
    },
    "id": 1
}
```

Updating template linked templates

Unlink (without clearing) all templates from the given template and link a different one to it.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.update",
    "params": {
        "templateid": "10086",
        "templates": [
```

```
            {
                "templateid": "10001"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templateids": [
            "10086"
        ]
    },
    "id": 1
}
```

Clearing multiple template linked templates

Unlink and clear a specific linked template from the given templates.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "template.update",
    "params": [
        {
            "templateid": "10087",
            "templates_clear": [
                {
                    "templateid": "10001"
                }
            ]
        },
        {
            "templateid": "10088",
            "templates_clear": [
                {
                    "templateid": "10001"
                }
            ]
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "templateids": [
            "10087",
            "10088"
        ]
    },
    "id": 1
}
```

Source

CTemplate::update() in *ui/include/classes/api/services/CTemplate.php*.

Template dashboard

This class is designed to work with template dashboards.

Object references:

- Template dashboard
- Template dashboard page
  - Template dashboard widget
    - Template dashboard widget field

Available methods:

- templatedashboard.create - create new template dashboards
- templatedashboard.delete - delete template dashboards
- templatedashboard.get - retrieve template dashboards
- templatedashboard.update - update template dashboards

Template dashboard object

The following objects are directly related to the `templatedashboard` API.

Template dashboard

The template dashboard object has the following properties.

Property	Type	Description
dashboardid	ID	ID of the template dashboard. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the template dashboard. Property behavior: - *required* for create operations
templateid	ID	ID of the template the dashboard belongs to. Property behavior: - *constant* - *required* for create operations
display_period	integer	Default slideshow interval (in seconds). Possible values: 10, 30, 60, 120, 600, 1800, 3600. Default: 30.
auto_start	integer	Auto start slideshow. Possible values: 0 - do not auto start slideshow; 1 - *(default)* auto start slideshow.
uuid	string	Universal unique identifier, used for linking imported template dashboards to already existing ones. Auto-generated, if not given.

Template dashboard page

The template dashboard page object has the following properties.

Property	Type	Description
dashboard_pageid	ID	ID of the dashboard page. Property behavior: - *read-only*

Property	Type	Description
name	string	Dashboard page name. Default: empty string.
display_period	integer	Dashboard slideshow interval (in seconds). Possible values: 0, 10, 30, 60, 120, 600, 1800, 3600. Default: 0 (will use the default slideshow interval).
widgets	array	Array of the template dashboard widget objects.

Template dashboard widget

The template dashboard widget object has the following properties.

Property	Type	Description
widgetid	ID	ID of the dashboard widget. Property behavior: - *read-only*
type	string	Type of the dashboard widget. Possible values: actionlog - Action log; clock - Clock; discovery - Discovery status; favgraphs - Favorite graphs; favmaps - Favorite maps; gauge - Gauge; graph - Graph (classic); graphprototype - Graph prototype; honeycomb - Honeycomb; hostavail - Host availability; hostcard - Host card; hostnavigator - Host navigator; itemcard - Item card; itemnavigator - Item navigator; item - Item value; map - Map; navtree - Map Navigation Tree; piechart - Pie chart; plaintext - Plain text; problemhosts - Problem hosts; problems - Problems; problemsbysv - Problems by severity; scatterplot - Scatter plot; slareport - SLA report; svggraph - Graph; systeminfo - System information; tophosts - Top hosts; topitems - Top items; toptriggers - Top triggers; trigover - Trigger overview; url - URL; web - Web monitoring. Property behavior: - *required*
name	string	Custom widget name.

Property	Type	Description
x	integer	A horizontal position from the left side of the dashboard. Possible values range from 0 to 71.
y	integer	A vertical position from the top of the dashboard. Possible values range from 0 to 63.
width	integer	The widget width. Possible values range from 1 to 72.
height	integer	The widget height. Possible values range from 1 to 64.
view_mode	integer	The widget view mode. Possible values: 0 - *(default)* default widget view; 1 - with hidden header;
fields	array	Array of the template dashboard widget field objects.

Template dashboard widget field

The template dashboard widget field object has the following properties.

Property	Type	Description
type	integer	Type of the widget field. Possible values: 0 - Integer; 1 - String; 4 - Item; 5 - Item prototype; 6 - Graph; 7 - Graph prototype; 8 - Map; 9 - Service; 10 - SLA; 11 - User; 12 - Action; 13 - Media type. Property behavior: - *required*
name	string	Widget field name. Possible values: see Dashboard widget fields. Note that some host-related parameters (e.g., *Host groups*, *Exclude host groups* and *Hosts* in the *Problems* widget, *Host groups* in the *Host availability* widget, etc.) are not available when configuring the widget on a template dashboard. This is because template dashboards display data only from the host that the template is linked to. Property behavior: - *required*

Property	Type	Description
value	mixed	Widget field value depending on the type. Possible values: see Dashboard widget fields. Note that some host-related parameters (e.g., *Host groups*, *Exclude host groups* and *Hosts* in the *Problems* widget, *Host groups* in the *Host availability* widget, etc.) are not available when configuring the widget on a template dashboard. This is because template dashboards display data only from the host that the template is linked to. Property behavior: - *required*

templatedashboard.create

Description

`object templatedashboard.create(object/array templateDashboards)`

This method allows you to create new template dashboards.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Template dashboards to create.

Additionally to the standard template dashboard properties, the method accepts the following parameters.

Parameter	Type	Description
pages	array	Template dashboard pages to be created for the dashboard. Dashboard pages will be ordered in the same order as specified. Parameter behavior: - *required*

Return values

`(object)` Returns an object containing the IDs of the created template dashboards under the `dashboardids` property. The order of the returned IDs matches the order of the passed template dashboards.

Examples

Creating a template dashboard

Create a template dashboard named “Graphs” with one Graph widget on a single dashboard page.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "templatedashboard.create",
    "params": {
        "templateid": "10318",
        "name": "Graphs",
        "pages": [
            {
                "widgets": [
                    {
                        "type": "graph",
                        "x": 0,
                        "y": 0,
```

```
                        "width": 12,
                        "height": 5,
                        "view_mode": 0,
                        "fields": [
                            {
                                "type": 6,
                                "name": "graphid",
                                "value": "1123"
                            }
                        ]
                    }
                ]

            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "32"
        ]
    },
    "id": 1
}
```

See also

- Template dashboard page
- Template dashboard widget
- Template dashboard widget field

Source

CTemplateDashboard::create() in *ui/include/classes/api/services/CTemplateDashboard.php*.

templatedashboard.delete

Description

`object templatedashboard.delete(array templateDashboardIds)`

This method allows you to delete template dashboards.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the template dashboards to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted template dashboards under the `dashboardids` property.

Examples

Deleting multiple template dashboards

Delete two template dashboards.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "templatedashboard.delete",
    "params": [
        "45",
        "46"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "45",
            "46"
        ]
    },
    "id": 1
}
```

Source

CTemplateDashboard::delete() in *ui/include/classes/api/services/CTemplateDashboard.php*.

templatedashboard.get

Description

`integer/array templatedashboard.get(object parameters)`

The method allows you to retrieve template dashboards according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
dashboardids	ID/array	Return only template dashboards with the given IDs.
templateids	ID/array	Return only template dashboards that belong to the given templates.
selectPages	query	Return a `pages` property with template dashboard pages, correctly ordered.
sortfield	string/array	Sort the result by the given properties. Possible values: `dashboardid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	

Parameter	Type	Description
startSearch	boolean	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving template dashboards

Retrieve all template dashboards with widgets for a specified template.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "templatedashboard.get",
    "params": {
        "output": "extend",
        "selectPages": "extend",
        "templateids": "10001"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "dashboardid": "23",
            "name": "Docker overview",
            "templateid": "10001",
            "display_period": "30",
            "auto_start": "1",
            "uuid": "6dfcbe0bc5ad400ea9c1c2dd7649282f",
            "pages": [
                {
                    "dashboard_pageid": "1",
                    "name": "",
                    "display_period": "0",
                    "widgets": [
                        {
                            "widgetid": "220",
                            "type": "graph",
                            "name": "",
                            "x": "0",
                            "y": "0",
                            "width": "36",
                            "height": "5",
                            "view_mode": "0",
                            "fields": [
                                {
                                    "type": "6",
                                    "name": "graphid",
                                    "value": "1125"
                                }
                            ]
                        },
                        {
                            "widgetid": "221",
```

```
                "type": "graph",
                "name": "",
                "x": "12",
                "y": "0",
                "width": "36",
                "height": "5",
                "view_mode": "0",
                "fields": [
                    {
                        "type": "6",
                        "name": "graphid",
                        "value": "1129"
                    }
                ]
            },
            {
                "widgetid": "222",
                "type": "graph",
                "name": "",
                "x": "0",
                "y": "5",
                "width": "36",
                "height": "5",
                "view_mode": "0",
                "fields": [
                    {
                        "type": "6",
                        "name": "graphid",
                        "value": "1128"
                    }
                ]
            },
            {
                "widgetid": "223",
                "type": "graph",
                "name": "",
                "x": "12",
                "y": "5",
                "width": "36",
                "height": "5",
                "view_mode": "0",
                "fields": [
                    {
                        "type": "6",
                        "name": "graphid",
                        "value": "1126"
                    }
                ]
            },
            {
                "widgetid": "224",
                "type": "graph",
                "name": "",
                "x": "0",
                "y": "10",
                "width": "36",
                "height": "5",
                "view_mode": "0",
                "fields": [
                    {
                        "type": "6",
                        "name": "graphid",
```

```
                                "value": "1127"
                            }
                        ]
                    }
                ]
            }
        ]
    }
],
"id": 1
}
```


See also

- Template dashboard page
- Template dashboard widget
- Template dashboard widget field

Source

CTemplateDashboard::get() in *ui/include/classes/api/services/CTemplateDashboard.php*.

templatedashboard.update

Description

`object templatedashboard.update(object/array templateDashboards)`

This method allows you to update existing template dashboards.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Template dashboard properties to be updated.

The `dashboardid` property must be specified for each dashboard, all other properties are optional. Only the specified properties will be updated.

Additionally to the standard template dashboard properties, the method accepts the following parameters.

Parameter	Type	Description
pages	array	Template dashboard pages to replace the existing dashboard pages. Dashboard pages are updated by the `dashboard_pageid` property. New dashboard pages will be created for objects without `dashboard_pageid` property and the existing dashboard pages will be deleted if not reused. Dashboard pages will be ordered in the same order as specified. Only the specified properties of the dashboard pages will be updated. At least one dashboard page object is required for `pages` property.

Return values

`(object)` Returns an object containing the IDs of the updated template dashboards under the `dashboardids` property.

Examples

Renaming a template dashboard

Rename a template dashboard to ”Performance graphs”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "templatedashboard.update",
    "params": {
        "dashboardid": "23",
        "name": "Performance graphs"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "23"
        ]
    },
    "id": 1
}
```

Updating template dashboard pages

Rename the first dashboard page, replace widgets on the second dashboard page and add a new page as the third one. Delete all other dashboard pages.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "templatedashboard.update",
    "params": {
        "dashboardid": "2",
        "pages": [
            {
                "dashboard_pageid": 1,
                "name": "Renamed Page"
            },
            {
                "dashboard_pageid": 2,
                "widgets": [
                    {
                        "type": "clock",
                        "x": 0,
                        "y": 0,
                        "width": 12,
                        "height": 3
                    }
                ]
            },
            {
                "display_period": 60
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "dashboardids": [
            "2"
```

```
        ]
    },
    "id": 1
}
```

See also

- Template dashboard widget
- Template dashboard widget field

Source

CTemplateDashboard::update() in *ui/include/classes/api/services/CTemplateDashboard.php*.

Template group

This class is designed to work with template groups.

Object references:

- Template group

Available methods:

- templategroup.create - create new template groups
- templategroup.delete - delete template groups
- templategroup.get - retrieve template groups
- templategroup.massadd - add related objects to template groups
- templategroup.massremove - remove related objects from template groups
- templategroup.propagate - propagate permissions to template groups' subgroups
- templategroup.update - update template groups

Template group object

The following objects are directly related to the `templategroup` API.

Template group

The template group object has the following properties.

Property	Type	Description
groupid	ID	ID of the template group. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the template group. Property behavior: - *required* for create operations
uuid	string	Universal unique identifier, used for linking imported template groups to already existing ones. Auto-generated, if not given.

templategroup.create

Description

`object templategroup.create(object/array templateGroups)`

This method allows you to create new template groups.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

(object/array) Template groups to create. The method accepts template groups with the standard template group properties.

Return values

(object) Returns an object containing the IDs of the created template groups under the groupids property. The order of the returned IDs matches the order of the passed template groups.

Examples

Creating a template group

Create a template group called "Templates/Databases".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "templategroup.create",
    "params": {
        "name": "Templates/Databases"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
            "107820"
        ]
    },
    "id": 1
}
```

Source

CTemplateGroup::create() in *ui/include/classes/api/services/CTemplateGroup.php*.

templategroup.delete

Description

`object templategroup.delete(array templateGroupIds)`

This method allows you to delete template groups.

A template group cannot be deleted if it contains templates that belong to this group only.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

(array) IDs of the template groups to delete.

Return values

(object) Returns an object containing the IDs of the deleted template groups under the groupids property.

Examples

Deleting multiple template groups

Delete two template groups.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "templategroup.delete",
    "params": [
        "107814",
        "107815"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
            "107814",
            "107815"
        ]
    },
    "id": 1
}
```

Source

CTemplateGroup::delete() in *ui/include/classes/api/services/CTemplateGroup.php*.

templategroup.get

Description

`integer/array templategroup.get(object parameters)`

The method allows you to retrieve template groups according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
graphids	ID/array	Return only template groups that contain templates with the given graphs.
groupids	ID/array	Return only template groups with the given template group IDs.
templateids	ID/array	Return only template groups that contain the given templates.
triggerids	ID/array	Return only template groups that contain templates with the given triggers.
with_graphs	boolean	Return only template groups that contain templates with graphs.
with_graph_prototypes	boolean	Return only template groups that contain templates with graph prototypes.
with_httptests	boolean	Return only template groups that contain templates with web checks.
with_items	boolean	Return only template groups that contain templates with items. Overrides the `with_simple_graph_items` parameters.
with_item_prototypes	boolean	Return only template groups that contain templates with item prototypes. Overrides the `with_simple_graph_item_prototypes` parameter.

Parameter	Type	Description
with_simple_graph_item_prototypes	boolean	Return only template groups that contain templates with item prototypes, which are enabled for creation and have numeric type of information.
with_simple_graph_items	boolean	Return only template groups that contain templates with numeric items.
with_templates	boolean	Return only template groups that contain templates.
with_triggers	boolean	Return only template groups that contain templates with triggers.
selectTemplates	query	Return a `templates` property with the templates that belong to the template group. Supports `count`.
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectTemplates` - results will be sorted by `template`.
sortfield	string/array	Sort the result by the given properties. Possible values: `groupid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving data by name

Retrieve all data about two template groups named ”Templates/Databases” and ”Templates/Modules”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "templategroup.get",
    "params": {
        "output": "extend",
        "filter": {
            "name": [
                "Templates/Databases",
                "Templates/Modules"
            ]
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
```

```
    "result": [
        {
            "groupid": "13",
            "name": "Templates/Databases",
            "uuid": "748ad4d098d447d492bb935c907f652f"
        },
        {
            "groupid": "8",
            "name": "Templates/Modules",
            "uuid": "57b7ae836ca64446ba2c296389c009b7"
        }
    ],
    "id": 1
}
```


See also

- Template

Source

CTemplateGroup::get() in *ui/include/classes/api/services/CTemplateGroup.php*.

templategroup.massadd

Description

`object templategroup.massadd(object parameters)`

This method allows you to simultaneously add multiple related objects to all the given template groups.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters containing the IDs of the template groups to update and the objects to add to all the template groups.

The method accepts the following parameters.

Parameter	Type	Description
groups	object/array	Template groups to be updated. The template groups must have only the `groupid` property defined. Parameter behavior: - *required*
templates	object/array	Templates to add to all template groups. The templates must have only the `templateid` property defined. Parameter behavior: - *required*

Return values

`(object)` Returns an object containing the IDs of the updated template groups under the `groupids` property.

Examples

Adding templates to template groups

Add two templates to template groups with IDs 12 and 13.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "templategroup.massadd",
    "params": {
        "groups": [
            {
                "groupid": "12"
            },
            {
                "groupid": "13"
            }
        ],
        "templates": [
            {
                "templateid": "10486"
            },
            {
                "templateid": "10487"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
            "12",
            "13"
        ]
    },
    "id": 1
}
```

See also

- Template

Source

CTemplateGroup::massAdd() in *ui/include/classes/api/services/CTemplateGroup.php*.

templategroup.massremove

Description

`object templategroup.massremove(object parameters)`

This method allows you to remove related objects from multiple template groups.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters containing the IDs of the template groups to update and the objects that should be removed.

Parameter	Type	Description
groupids	ID/array	IDs of the template groups to be updated. Parameter behavior: - *required*
templateids	ID/array	IDs of the templates to remove from all template groups. Parameter behavior: - *required*

Return values

`(object)` Returns an object containing the IDs of the updated template groups under the `groupids` property.

Examples

Removing templates from template groups

Remove two templates from the given template groups.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "templategroup.massremove",
    "params": {
        "groupids": [
            "5",
            "6"
        ],
        "templateids": [
            "30050",
            "30001"
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
            "5",
            "6"
        ]
    },
    "id": 1
}
```

Source

CTemplateGroup::massRemove() in *ui/include/classes/api/services/CTemplateGroup.php*.

templategroup.propagate

Description

`object templategroup.propagate(object parameters)`

This method allows you to apply permissions to all template groups' subgroups.

Note:
This method is only available to *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
groups	object/array	Template groups to propagate. The template groups must have only the `groupid` property defined. Parameter behavior: *- required*
permissions	boolean	Set `true` if need to propagate permissions. Parameter behavior: *- required*

Return values

`(object)` Returns an object containing the IDs of the propagated template groups under the `groupids` property.

Examples

Propagating template group permissions to its subgroups.

Propagate template group permissions to its subgroups.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "templategroup.propagate",
    "params": {
        "groups": [
            {
                "groupid": "15"
            }
        ],
        "permissions": true
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
            "15",
        ]
    },
    "id": 1
}
```

See also

- templategroup.update
- templategroup.massadd
- Template

Source

CTemplateGroup::propagate() in *ui/include/classes/api/services/CTemplateGroup.php*.

templategroup.update

Description

`object templategroup.update(object/array templateGroups)`

This method allows you to update existing template groups.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Template group properties to be updated.

The `groupid` property must be defined for each template group, all other properties are optional. Only the given properties will be updated, all others will remain unchanged.

Return values

`(object)` Returns an object containing the IDs of the updated template groups under the `groupids` property.

Examples

Renaming a template group

Rename a template group to "Templates/Databases"

Request:

```
{
    "jsonrpc": "2.0",
    "method": "templategroup.update",
    "params": {
        "groupid": "7",
        "name": "Templates/Databases"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "groupids": [
            "7"
        ]
    },
    "id": 1
}
```

Source

CTemplateGroup::update() in *ui/include/classes/api/services/CTemplateGroup.php*.

Token

This class is designed to work with tokens.

Object references:

- Token

Available methods:

- token.create - create new tokens
- token.delete - delete tokens
- token.get - retrieve tokens
- token.update - update tokens

- token.generate - generate tokens

Token object

The following objects are directly related to the `token` API.

Token

The token object has the following properties.

Property	Type	Description
tokenid	ID	ID of the token. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the token. Property behavior: - *required* for create operations
description	text	Description of the token.
userid	ID	ID of the user that the token has been assigned to. Default: *current user*. Property behavior: - *constant*
lastaccess	timestamp	Most recent date and time the token was authenticated. "0" if the token has never been authenticated. Property behavior: - *read-only*
status	integer	Token status. Possible values: 0 - *(default)* enabled token; 1 - disabled token.
expires_at	timestamp	Token expiration date and time. "0" for never-expiring tokens.
created_at	timestamp	Token creation date and time. Property behavior: - *read-only*
creator_userid	ID	ID of the user that created the token. Property behavior: - *read-only*

token.create

Description

`object token.create(object/array tokens)`

This method allows you to create new tokens.

> **Note:**
> The *Manage API tokens* permission is required for the user role to manage tokens for other users.

Attention:
A token created by this method also has to be generated before it is usable.

Parameters

`(object/array)` Tokens to create.

The method accepts tokens with the standard token properties.

Return values

`(object)` Returns an object containing the IDs of the created tokens under the `tokenids` property. The order of the returned IDs matches the order of the passed tokens.

Examples

Create a token

Create an enabled token that never expires and authenticates user of ID 2.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "token.create",
    "params": {
        "name": "Your token",
        "userid": "2"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "tokenids": [
            "188"
        ]
    },
    "id": 1
}
```

Create a disabled token that expires at January 21st, 2021. This token will authenticate current user.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "token.create",
    "params": {
        "name": "Your token",
        "status": "1",
        "expires_at": "1611238072"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "tokenids": [
            "189"
        ]
    },
    "id": 1
}
```

Source

CToken::create() in *ui/include/classes/api/services/CToken.php*.

token.delete

Description

`object token.delete(array tokenids)`

This method allows you to delete tokens.

> **Note:**
> The *Manage API tokens* permission is required for the user role to manage tokens for other users.

Parameters

`(array)` IDs of the tokens to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted tokens under the `tokenids` property.

Examples

Delete multiple tokens

Delete two tokens.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "token.delete",
    "params": [
        "188",
        "192"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "tokenids": [
            "188",
            "192"
        ]
    },
    "id": 1
}
```

Source

CToken::delete() in *ui/include/classes/api/services/CToken.php*.

token.generate

Description

`object token.generate(array tokenids)`

This method allows you to generate tokens.

> **Note:**
> The *Manage API tokens* permission is required for the user role to manage tokens for other users.

Attention:
A token can be generated by this method only if it has been created.

Parameters

`(array)` IDs of the tokens to generate.

Return values

`(array)` Returns an array of objects containing the ID of the generated token under the `tokenid` property and generated authorization string under `token` property.

Property	Type	Description
tokenid	ID	ID of the token.
token	string	The generated authorization string for this token.

Examples

Generate multiple tokens

Generate two tokens.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "token.generate",
    "params": [
        "1",
        "2"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "tokenid": "1",
            "token": "bbcfce79a2d95037502f7e9a534906d3466c9a1484beb6ea0f4e7be28e8b8ce2"
        },
        {
            "tokenid": "2",
            "token": "fa1258a83d518eabd87698a96bd7f07e5a6ae8aeb8463cae33d50b91dd21bd6d"
        }
    ],
    "id": 1
}
```

Source

CToken::generate() in *ui/include/classes/api/services/CToken.php*.

token.get

Description

`integer/array token.get(object parameters)`

The method allows you to retrieve tokens according to the given parameters.

Note:
Only *Super admin* user type is allowed to view tokens for other users.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
tokenids	ID/array	Return only tokens with the given IDs.
userids	ID/array	Return only tokens created for the given users.
token	string	Return only tokens created for the given *Auth token*.
valid_at	timestamp	Return only tokens, which are valid (not expired) at the given date and time.
expired_at	timestamp	Return only tokens, which are expired (not valid) at the given date and time.
sortfield	string/array	Sort the result by the given properties. Possible values: `tokenid`, `name`, `lastaccess`, `status`, `expires_at`, `created_at`.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve a token

Retrieve all data for the token with ID "2".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "token.get",
    "params": {
        "output": "extend",
        "tokenids": "2"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "tokenid": "1",
            "name": "The Token",
            "description": "",
            "userid": "1",
            "lastaccess": "0",
            "status": "0",
```

```
            "expires_at": "1609406220",
            "created_at": "1611239454",
            "creator_userid": "1"
        }
    ],
    "id": 1
}
```

Source

CToken::get() in *ui/include/classes/api/services/CToken.php*.

token.update

Description

`object token.update(object/array tokens)`

This method allows you to update existing tokens.

> **Note:**
> The *Manage API tokens* permission is required for the user role to manage tokens for other users.

Parameters

`(object/array)` Token properties to be updated.

The `tokenid` property must be defined for each token, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

The method accepts tokens with the standard token properties.

Return values

`(object)` Returns an object containing the IDs of the updated tokens under the `tokenids` property.

Examples

Remove token expiry

Remove expiry date from token.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "token.update",
    "params": {
        "tokenid": "2",
        "expires_at": "0"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "tokenids": [
            "2"
        ]
    },
    "id": 1
}
```

Source

CToken::update() in *ui/include/classes/api/services/CToken.php*.

Trend

This class is designed to work with trend data.

Object references:

- Float trend
- Integer trend

Available methods:

- trend.get - retrieve trends

Trend object

The following objects are directly related to the `trend` API.

Note:
Trend objects differ depending on the item's type of information. They are created by the Zabbix server and cannot be modified via the API.

Float trend

The float trend object has the following properties.

Property	Type	Description
clock	timestamp	Timestamp of an hour for which the value was calculated. For example, timestamp of "04:00:00" means values calculated for period "04:00:00-04:59:59".
itemid	ID	ID of the related item.
num	integer	Number of values that were available for the hour.
value_min	float	Hourly minimum value.
value_avg	float	Hourly average value.
value_max	float	Hourly maximum value.

Integer trend

The integer trend object has the following properties.

Property	Type	Description
clock	timestamp	Timestamp of an hour for which the value was calculated. For example, timestamp of "04:00:00" means values calculated for period "04:00:00-04:59:59".
itemid	ID	ID of the related item.
num	integer	Number of values that were available for the hour.
value_min	integer	Hourly minimum value.
value_avg	integer	Hourly average value.
value_max	integer	Hourly maximum value.

trend.get

Description

`integer/array trend.get(object parameters)`

The method allows you to retrieve trend data according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
itemids	ID/array	Return only trends with the given item IDs.
time_from	timestamp	Return only values that have been collected after or at the given time.
time_till	timestamp	Return only values that have been collected before or at the given time.
countOutput	boolean	Count the number of retrieved objects.
limit	integer	Limit the amount of retrieved objects.
output	query	Set Trend object properties to be returned.

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving item trend data

Request:

```
{
    "jsonrpc": "2.0",
    "method": "trend.get",
    "params": {
        "output": [
            "itemid",
            "clock",
            "num",
            "value_min",
            "value_avg",
            "value_max"
        ],
        "itemids": [
            "23715"
        ],
        "limit": "1"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "itemid": "23715",
            "clock": "1446199200",
            "num": "60",
            "value_min": "0.165",
            "value_avg": "0.2168",
            "value_max": "0.35"
        }
    ],
    "id": 1
}
```

Source

CTrend::get() in *ui/include/classes/api/services/CTrend.php*.

Trigger

This class is designed to work with triggers.

Object references:

- Trigger
- Trigger tag

Available methods:

- trigger.create - create new triggers
- trigger.delete - delete triggers
- trigger.get - retrieve triggers
- trigger.update - update triggers

Trigger object

The following objects are directly related to the `trigger` API.

Trigger

The trigger object has the following properties.

Property	Type	Description
triggerid	ID	ID of the trigger. Property behavior: - *read-only* - *required* for update operations
description	string	Name of the trigger. Property behavior: - *required* for create operations
expression	string	Trigger expression. Items with `value_type` set to 5 (binary) or 6 (JSON) cannot be used in trigger expressions. Property behavior: - *required* for create operations
event_name	string	Event name generated by the trigger.
opdata	string	Operational data.
comments	string	Additional description of the trigger.
error	string	Error text if there have been any problems when updating the state of the trigger. Property behavior: - *read-only*
flags	integer	Origin of the trigger. Possible values: 0 - *(default)* a plain trigger; 4 - a trigger converted from prototype. Property behavior: - *read-only*
lastchange	timestamp	Time when the trigger last changed its state. Property behavior: - *read-only*

Property	Type	Description
priority	integer	Severity of the trigger. Possible values: 0 - *(default)* not classified; 1 - information; 2 - warning; 3 - average; 4 - high; 5 - disaster.
state	integer	State of the trigger expression. Possible values: 0 - *(default)* trigger state is up to date; 1 - current trigger state is unknown. Property behavior: - *read-only*
status	integer	Whether the trigger is enabled or disabled. Possible values: 0 - *(default)* enabled; 1 - disabled.
templateid	ID	ID of the parent template trigger. Property behavior: - *read-only*
type	integer	Whether the trigger can generate multiple problem events. Possible values: 0 - *(default)* do not generate multiple events; 1 - generate multiple events.
url	string	URL associated with the trigger.
url_name	string	Label for the URL associated with the trigger.
value	integer	Whether the trigger is in OK or problem state. Possible values: 0 - *(default)* OK; 1 - problem. Property behavior: - *read-only*
recovery_mode	integer	OK event generation mode. Possible values: 0 - *(default)* Expression; 1 - Recovery expression; 2 - None.
recovery_expression	string	Trigger recovery expression. Items with `value_type` set to 5 (binary) or 6 (JSON) cannot be used in trigger recovery expressions.
correlation_mode	integer	OK event closes. Possible values: 0 - *(default)* All problems; 1 - All problems if tag values match.
correlation_tag	string	Tag for matching.
manual_close	integer	Allow manual close. Possible values: 0 - *(default)* No; 1 - Yes.

Property	Type	Description
uuid	string	Universal unique identifier, used for linking imported triggers to already existing ones. Auto-generated, if not given. Property behavior: - *supported* if the trigger belongs to a template

Trigger tag

The trigger tag object has the following properties.

Property	Type	Description
tag	string	Trigger tag name. Property behavior: - *required*
value	string	Trigger tag value. Property behavior: - *read-only*.
object	integer	Type of object the tag was inherited from. Possible values: 0 - template; 1 - host; 3 - item. Property behavior: - *read-only*.
objectid	ID	ID of the object the tag was inherited from. Property behavior: - *read-only*.
automatic	integer	Type of a trigger tag: Possible values: 0 - *(default)* tag created by user; 1 - tag created by LLD;Property behavior: - *read-only*.

trigger.create

Description

`object trigger.create(object/array triggers)`

This method allows you to create new triggers.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Triggers to create.

Additionally to the standard trigger properties the method accepts the following parameters.

Parameter	Type	Description
dependencies	array	Triggers that the trigger is dependent on. The triggers must have only the `triggerid` property defined.
tags	array	Trigger tags.

Return values

`(object)` Returns an object containing the IDs of the created triggers under the `triggerids` property. The order of the returned IDs matches the order of the passed triggers.

Examples

Creating a trigger

Create two triggers, each depending on one other trigger.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "trigger.create",
    "params": [
        {
            "description": "Processor load is too high on {HOST.NAME}",
            "expression": "last(/Linux server/system.cpu.load[percpu,avg1])>5",
            "dependencies": [
                {
                    "triggerid": "17367"
                }
            ]
        },
        {
            "description": "Service has stopped",
            "expression": "length(last(/Linux server/log[/var/log/system,Service .* has stopped]))<>0",
            "dependencies": [
                {
                    "triggerid": "17368"
                }
            ],
            "tags": [
                {
                    "tag": "service",
                    "value": "{{ITEM.VALUE}.regsub(\"Service (.*) has stopped\", \"\\1\")}"
                },
                {
                    "tag": "error",
                    "value": ""
                }
            ]
        }
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "triggerids": [
            "17369",
            "17370"
        ]
    },
```

```
    "id": 1
}
```

Source

CTrigger::create() in *ui/include/classes/api/services/CTrigger.php*.

trigger.delete

Description

`object trigger.delete(array triggerIds)`

This method allows you to delete triggers.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the triggers to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted triggers under the `triggerids` property.

Examples

Delete multiple triggers

Delete two triggers.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "trigger.delete",
    "params": [
        "12002",
        "12003"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "triggerids": [
            "12002",
            "12003"
        ]
    },
    "id": 1
}
```

Source

CTrigger::delete() in *ui/include/classes/api/services/CTrigger.php*.

trigger.get

Description

`integer/array trigger.get(object parameters)`

The method allows you to retrieve triggers according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
triggerids	ID/array	Return only triggers with the given IDs.
groupids	ID/array	Return only triggers that belong to hosts or templates from the given host groups or template groups.
templateids	ID/array	Return only triggers that belong to the given templates.
hostids	ID/array	Return only triggers that belong to the given hosts.
itemids	ID/array	Return only triggers that contain the given items.
maintenanceids	ID/array	Return only triggers that are affected by the given maintenances.
functions	string/array	Return only triggers that use the given functions. Refer to the supported function page for a list of supported functions.
group	string	Return only triggers that belong to hosts or templates from the host group or template group with the given name.
host	string	Return only triggers that belong to host with the given technical name.
inherited	boolean	If set to `true` return only triggers inherited from a template.
templated	boolean	If set to `true` return only triggers that belong to templates.
dependent	boolean	If set to `true` return only triggers that have dependencies. If set to `false` return only triggers that do not have dependencies.
monitored	flag	Return only enabled triggers that belong to monitored hosts and contain only enabled items.
active	flag	Return only enabled triggers that belong to monitored hosts.
maintenance	boolean	If set to `true` return only enabled triggers that belong to hosts in maintenance.
withUnacknowledgedEvents	flag	Return only triggers that have unacknowledged events.
withAcknowledgedEvents	flag	Return only triggers with all events acknowledged.
withLastEventUnacknowledged	flag	Return only triggers with the last event unacknowledged.
skipDependent	flag	Skip triggers in a problem state that are dependent on other triggers. Note that the other triggers are ignored if disabled, have disabled items or disabled item hosts.
lastChangeSince	timestamp	Return only triggers that have changed their state after the given time.
lastChangeTill	timestamp	Return only triggers that have changed their state before the given time.
only_true	flag	Return only triggers that have been in a problem state recently (within the `ok_period`).
min_severity	integer	Return only triggers with severity greater or equal than the given severity.
evaltype	integer	Tag evaluation method. Possible values: 0 - *(default)* And/Or; 2 - Or.
tags	array	Return only triggers with the given tags. Format: `[{"tag": "<tag>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all triggers. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.

Parameter	Type	Description
expandComment	flag	Expand macros in the trigger description.
expandDescription	flag	Expand macros in the name of the trigger.
expandExpression	flag	Expand functions and macros in the trigger expression.
selectHostGroups	query	Return the host groups that the trigger belongs to in the `hostgroups` property.
selectHosts	query	Return the hosts that the trigger belongs to in the `hosts` property.
selectItems	query	Return items contained by the trigger in the `items` property.
selectFunctions	query	Return functions used in the trigger in the `functions` property. The function objects represent the functions used in the trigger expression and has the following properties: `functionid` - `(ID)` ID of the function; `itemid` - `(ID)` ID of the item used in the function; `function` - `(string)` name of the function; `parameter` - `(string)` parameter passed to the function. Query parameter is replaced by $ symbol in returned string.
selectDependencies	query	Return triggers that the trigger depends on in the `dependencies` property.
selectDiscoveryData	query	Return a `discoveryData` property with the trigger discovery object data. The trigger discovery object links a discovered trigger to a trigger prototype from which it was discovered. It has the following properties: `parent_triggerid` - `(ID)` ID of the trigger prototype from which the trigger has been created; `status` - `(int)` trigger discovery status: 0 - *(default)* trigger is discovered, 1 - trigger is not discovered anymore; `ts_delete` - `(timestamp)` time when a trigger that is no longer discovered will be deleted; `ts_disable` - `(timestamp)` time when a trigger that is no longer discovered will be disabled; `disable_source` - `(int)` indicator of whether trigger was disabled by an LLD rule or manually: 0 - *(default)* disabled automatically, 1 - disabled by an LLD rule.
selectDiscoveryRule	query	Return the low-level discovery rule that created the trigger in the `discoveryRule` property.
selectLastEvent	query	Return the last significant trigger event in the `lastEvent` property.
selectTags	query	Return the trigger tags in `tags` property.
inheritedTags	boolean	Return triggers that have given `tags` also in template/host/linked templates. Possible values: `true` - template/host/linked templates must also have given tags; `false` - *(default)* tags from template/host/linked templates are ignored.
selectInheritedTags	query	Return an `inheritedTags` property with tags inherited from templates, linked templates and hosts, and from items referenced in trigger or recovery expressions.
selectTemplateGroups	query	Return the template groups that the trigger belongs to in the `templategroups` property.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type. Supports additional properties: `host` - technical name of the host that the trigger belongs to; `hostid` - ID of the host that the trigger belongs to.

Parameter	Type	Description
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectHosts` - results will be sorted by `host`.
sortfield	string/array	Sort the result by the given properties. Possible values: `triggerid`, `description`, `status`, `priority`, `lastchange`, `hostname`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	
selectTriggerDiscovery	query	Return the trigger discovery object in the `triggerDiscovery` property. The trigger discovery objects link the trigger to a trigger prototype from which it was created. This query is **deprecated**, please use `selectDiscoveryData` instead.

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving data by trigger ID

Retrieve all data and the functions used in trigger "14062".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "trigger.get",
    "params": {
        "triggerids": "14062",
        "output": "extend",
        "selectFunctions": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "triggerid": "14062",
            "expression": "{13513}<10m",
            "description": "{HOST.NAME} has been restarted (uptime < 10m)",
            "url": "",
            "status": "0",
            "value": "0",
            "priority": "2",
```

```
            "lastchange": "0",
            "comments": "The host uptime is less than 10 minutes",
            "error": "",
            "templateid": "10016",
            "type": "0",
            "state": "0",
            "flags": "0",
            "recovery_mode": "0",
            "recovery_expression": "",
            "correlation_mode": "0",
            "correlation_tag": "",
            "manual_close": "0",
            "opdata": "",
            "event_name": "",
            "uuid": "",
            "url_name": "",
            "functions": [
                {
                    "functionid": "13513",
                    "itemid": "24350",
                    "triggerid": "14062",
                    "parameter": "$",
                    "function": "last"
                }
            ]
        }
    ],
    "id": 1
}
```

Retrieving triggers in problem state

Retrieve the ID, name and severity of all triggers in problem state and sort them by severity in descending order.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "trigger.get",
    "params": {
        "output": [
            "triggerid",
            "description",
            "priority"
        ],
        "filter": {
            "value": 1
        },
        "sortfield": "priority",
        "sortorder": "DESC"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "triggerid": "13907",
            "description": "Zabbix self-monitoring processes < 100% busy",
            "priority": "4"
        },
        {
```

```
            "triggerid": "13824",
            "description": "Zabbix discoverer processes more than 75% busy",
            "priority": "3"
        }
    ],
    "id": 1
}
```

Retrieving a specific trigger with tags

Retrieve a specific trigger with tags.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "trigger.get",
    "params": {
        "output": [
            "triggerid",
            "description"
        ],
        "selectTags": "extend",
        "triggerids": [
            "17578"
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "triggerid": "17370",
            "description": "Service status",
            "tags": [
                {
                    "tag": "service",
                    "value": "{{ITEM.VALUE}.regsub(\"Service (.*) has stopped\", \"\\1\")}",
                    "automatic": "1"
                },
                {
                    "tag": "error",
                    "value": "",
                    "automatic": "1"
                }
            ]
        }
    ],
    "id": 1
}
```

See also

- Discovery rule
- Item
- Host
- Host group
- Template group

Source

CTrigger::get() in *ui/include/classes/api/services/CTrigger.php*.

trigger.update

Description

`object trigger.update(object/array triggers)`

This method allows you to update existing triggers.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Trigger properties to be updated.

The `triggerid` property must be defined for each trigger, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard trigger properties the method accepts the following parameters.

Parameter	Type	Description
dependencies	array	Triggers that the trigger is dependent on. The triggers must have only the `triggerid` property defined.
tags	array	Trigger tags.

Return values

`(object)` Returns an object containing the IDs of the updated triggers under the `triggerids` property.

Examples

Enabling a trigger

Enable a trigger, that is, set its status to "0".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "trigger.update",
    "params": {
        "triggerid": "13938",
        "status": 0
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "triggerids": [
            "13938"
        ]
    },
    "id": 1
}
```

Replacing triggers tags

Replace tags for trigger.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "trigger.update",
    "params": {
        "triggerid": "13938",
        "tags": [
            {
                "tag": "service",
                "value": "{{ITEM.VALUE}.regsub(\"Service (.*) has stopped\", \"\\1\")}"
            },
            {
                "tag": "error",
                "value": ""
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "triggerids": [
            "13938"
        ]
    },
    "id": 1
}
```

Replacing dependencies

Replace dependencies for trigger.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "trigger.update",
    "params": {
        "triggerid": "22713",
        "dependencies": [
            {
                "triggerid": "22712"
            },
            {
                "triggerid": "22772"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "triggerids": [
            "22713"
        ]
    },
    "id": 1
}
```

Source

CTrigger::update() in *ui/include/classes/api/services/CTrigger.php*.

Trigger prototype

This class is designed to work with trigger prototypes.

Object references:

- Trigger prototype
- Trigger prototype tag

Available methods:

- triggerprototype.create - create new trigger prototypes
- triggerprototype.delete - delete trigger prototypes
- triggerprototype.get - retrieve trigger prototypes
- triggerprototype.update - update trigger prototypes

Trigger prototype object

The following objects are directly related to the `triggerprototype` API.

Trigger prototype

The trigger prototype object has the following properties.

Property	Type	Description
triggerid	ID	ID of the trigger prototype. Property behavior: - *read-only* - *required* for update operations
description	string	Name of the trigger prototype. Property behavior: - *required* for create operations
expression	string	Trigger expression. Must contain at least one item prototype. Item prototypes with `value_type` set to 5 (binary) or 6 (JSON) cannot be used in trigger expressions. Property behavior: - *required* for create operations
event_name	string	Event name generated by the trigger.
opdata	string	Operational data.
comments	string	Additional comments to the trigger prototype.
priority	integer	Severity of the trigger prototype. Possible values: 0 - *(default)* not classified; 1 - information; 2 - warning; 3 - average; 4 - high; 5 - disaster.
status	integer	Whether the trigger prototype is enabled or disabled. Possible values: 0 - *(default)* enabled; 1 - disabled.

Property	Type	Description
flags	integer	Origin of the trigger prototype. Possible values: 2 - a trigger prototype; 6 - a discovered trigger prototype. Property behavior: - *read-only*
templateid	ID	ID of the parent template trigger prototype. Property behavior: - *read-only*
type	integer	Whether the trigger prototype can generate multiple problem events. Possible values: 0 - *(default)* do not generate multiple events; 1 - generate multiple events.
url	string	URL associated with the trigger prototype.
url_name	string	Label for the URL associated with the trigger prototype.
recovery_mode	integer	OK event generation mode. Possible values: 0 - *(default)* Expression; 1 - Recovery expression; 2 - None.
recovery_expression	string	Trigger recovery expression. Must contain at least one item prototype. Item prototypes with `value_type` set to 5 (binary) or 6 (JSON) cannot be used in trigger recovery expressions.
correlation_mode	integer	OK event closes. Possible values: 0 - *(default)* All problems; 1 - All problems if tag values match.
correlation_tag	string	Tag for matching.
manual_close	integer	Allow manual close. Possible values: 0 - *(default)* No; 1 - Yes.
discover	integer	Trigger prototype discovery status. Possible values: 0 - *(default)* new triggers will be discovered; 1 - new triggers will not be discovered and existing triggers will be marked as lost.
uuid	string	Universal unique identifier, used for linking imported trigger prototypes to already existing ones. Auto-generated, if not given. Property behavior: - *supported* if the trigger prototype belongs to a template

Trigger prototype tag

The trigger prototype tag object has the following properties.

Property	Type	Description
tag	string	Trigger prototype tag name. Property behavior: - *required*
value	string	Trigger prototype tag value. Property behavior: - *read-only*.
object	integer	Type of object the tag was inherited from. Possible values: 0 - template; 1 - host; 3 - item; 4 - item prototype. Property behavior: - *read-only*.
objectid	ID	ID of the object the tag was inherited from. Property behavior: - *read-only*.

triggerprototype.create

Description

`object triggerprototype.create(object/array triggerPrototypes)`

This method allows you to create new trigger prototypes.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Trigger prototypes to create.

Additionally to the standard trigger prototype properties the method accepts the following parameters.

Parameter	Type	Description
dependencies	array	Triggers and trigger prototypes that the trigger prototype is dependent on. The triggers must have only the `triggerid` property defined.
tags	array	Trigger prototype tags.

Return values

`(object)` Returns an object containing the IDs of the created trigger prototypes under the `triggerids` property. The order of the returned IDs matches the order of the passed trigger prototypes.

Examples

Creating a trigger prototype

Create a trigger prototype to detect when a file system has less than 20% free disk space.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "triggerprototype.create",
    "params": {
        "description": "Free disk space is less than 20% on volume {#FSNAME}",
        "expression": "last(/Zabbix server/vfs.fs.size[{#FSNAME},pfree])<20",
        "tags": [
            {
                "tag": "volume",
                "value": "{#FSNAME}"
            },
            {
                "tag": "type",
                "value": "{#FSTYPE}"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "triggerids": [
            "17372"
        ]
    },
    "id": 1
}
```

Source

CTriggerPrototype::create() in *ui/include/classes/api/services/CTriggerPrototype.php*.

triggerprototype.delete

Description

`object triggerprototype.delete(array triggerPrototypeIds)`

This method allows you to delete trigger prototypes.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the trigger prototypes to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted trigger prototypes under the `triggerids` property.

Examples

Deleting multiple trigger prototypes

Delete two trigger prototypes.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "triggerprototype.delete",
    "params": [
```

```
        "12002",
        "12003"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "triggerids": [
            "12002",
            "12003"
        ]
    },
    "id": 1
}
```

Source

CTriggerPrototype::delete() in *ui/include/classes/api/services/CTriggerPrototype.php*.

triggerprototype.get

Description

`integer/array triggerprototype.get(object parameters)`

The method allows you to retrieve trigger prototypes according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
active	flag	Return only enabled trigger prototypes that belong to monitored hosts.
discoveryids	ID/array	Return only trigger prototypes that belong to the given LLD rules.
functions	string/array	Return only triggers that use the given functions. Refer to the Supported functions page for a list of supported functions.
group	string	Return only trigger prototypes that belong to hosts or templates from the host groups or template groups with the given name.
groupids	ID/array	Return only trigger prototypes that belong to hosts or templates from the given host groups or template groups.
host	string	Return only trigger prototypes that belong to hosts with the given name.
hostids	ID/array	Return only trigger prototypes that belong to the given hosts.
inherited	boolean	If set to `true` return only trigger prototypes inherited from a template.
maintenance	boolean	If set to `true` return only enabled trigger prototypes that belong to hosts in maintenance.
min_severity	integer	Return only trigger prototypes with severity greater or equal than the given severity.
monitored	flag	Return only enabled trigger prototypes that belong to monitored hosts and contain only enabled items.
templated	boolean	If set to `true` return only trigger prototypes that belong to templates.
templateids	ID/array	Return only trigger prototypes that belong to the given templates.
triggerids	ID/array	Return only trigger prototypes with the given IDs.
expandExpression	flag	Expand functions and macros in the trigger expression.

Parameter	Type	Description
selectDependencies	query	Return trigger prototypes and triggers that the trigger prototype depends on in the `dependencies` property.
selectDiscoveryData	query	Return a `discoveryData` property with the trigger prototype discovery object data. The trigger prototype discovery object links a discovered trigger prototype to a trigger prototype from which it was discovered. It has the following properties: `parent_triggerid` - `(ID)` ID of the trigger prototype from which the trigger prototype has been created; `status` - `(int)` trigger prototype discovery status: 0 - *(default)* trigger prototype is discovered, 1 - trigger prototype is not discovered anymore; `ts_delete` - `(timestamp)` time when a trigger prototype that is no longer discovered will be deleted; `ts_disable` - `(timestamp)` time when a trigger prototype that is no longer discovered will be disabled; `disable_source` - `(int)` indicator of whether trigger prototype was disabled by an LLD rule or manually: 0 - *(default)* disabled automatically, 1 - disabled by an LLD rule.
selectDiscoveryRule	query	Return the LLD rule that the trigger prototype belongs to in the `discoveryRule` property.
selectDiscoveryRulePrototype	query	Return a `discoveryRulePrototype` property with the parent LLD rule prototype that the trigger prototype belongs to.
selectFunctions	query	Return functions used in the trigger prototype in the `functions` property. The function objects represent the functions used in the trigger expression and has the following properties: `functionid` - `(ID)` ID of the function; `itemid` - `(ID)` ID of the item used in the function; `function` - `(string)` name of the function; `parameter` - `(string)` parameter passed to the function. Query parameter is replaced by $ symbol in returned string.
selectHostGroups	query	Return the host groups that the trigger prototype belongs to in the `hostgroups` property.
selectHosts	query	Return the hosts that the trigger prototype belongs to in the `hosts` property.
selectInheritedTags	query	Return an `inheritedTags` property with tags inherited from templates, linked templates and hosts, and from item prototypes referenced in trigger or recovery expressions.
selectItems	query	Return items and item prototypes used the trigger prototype in the `items` property.
selectTags	query	Return the trigger prototype tags in `tags` property.
selectTemplateGroups	query	Return the template groups that the trigger prototype belongs to in the `templategroups` property.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type. Supports additional properties: `host` - technical name of the host that the trigger prototype belongs to; `hostid` - ID of the host that the trigger prototype belongs to.
limitSelects	integer	Limits the number of records returned by subselects. Applies to the following subselects: `selectHosts` - results will be sorted by `host`.

Parameter	Type	Description
sortfield	string/array	Sort the result by the given properties. Possible values: `triggerid`, `description`, `status`, `priority`, `discovered`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieve trigger prototypes from an LLD rule

Retrieve all trigger prototypes and their functions from an LLD rule.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "triggerprototype.get",
    "params": {
        "output": "extend",
        "selectFunctions": "extend",
        "discoveryids": "22450"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "triggerid": "13272",
            "expression": "{12598}<20",
            "description": "Free inodes is less than 20% on volume {#FSNAME}",
            "url": "",
            "status": "0",
            "value": "0",
            "priority": "2",
            "lastchange": "0",
            "comments": "",
            "error": "",
            "templateid": "0",
            "type": "0",
            "state": "0",
            "flags": "2",
            "recovery_mode": "0",
            "recovery_expression": "",
            "correlation_mode": "0",
```

```
            "correlation_tag": "",
            "manual_close": "0",
            "opdata": "",
            "discover": "0",
            "event_name": "",
            "uuid": "6ce467d05e8745409a177799bed34bb3",
            "url_name": "",
            "functions": [
                {
                    "functionid": "12598",
                    "itemid": "22454",
                    "triggerid": "13272",
                    "parameter": "$",
                    "function": "last"
                }
            ]
        },
        {
            "triggerid": "13266",
            "expression": "{13500}<20",
            "description": "Free disk space is less than 20% on volume {#FSNAME}",
            "url": "",
            "status": "0",
            "value": "0",
            "priority": "2",
            "lastchange": "0",
            "comments": "",
            "error": "",
            "templateid": "0",
            "type": "0",
            "state": "0",
            "flags": "2",
            "recovery_mode": "0",
            "recovery_expression": "",
            "correlation_mode": "0",
            "correlation_tag": "",
            "manual_close": "0",
            "opdata": "",
            "discover": "0",
            "event_name": "",
            "uuid": "74a1fc62bfe24b7eabe4e244c70dc384",
            "url_name": "",
            "functions": [
                {
                    "functionid": "13500",
                    "itemid": "22686",
                    "triggerid": "13266",
                    "parameter": "$",
                    "function": "last"
                }
            ]
        }
    ],
    "id": 1
}
```


Retrieving a specific trigger prototype with tags

Request:

```
{
    "jsonrpc": "2.0",
    "method": "triggerprototype.get",
    "params": {
```

```
        "output": [
            "triggerid",
            "description"
        ],
        "selectTags": "extend",
        "triggerids": [
            "17373"
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "triggerid": "17373",
            "description": "Free disk space is less than 20% on volume {#FSNAME}",
            "tags": [
                {
                    "tag": "volume",
                    "value": "{#FSNAME}"
                },
                {
                    "tag": "type",
                    "value": "{#FSTYPE}"
                }
            ]
        }
    ],
    "id": 1
}
```

See also

- Discovery rule
- Item
- Host
- Host group
- Template group

Source

CTriggerPrototype::get() in *ui/include/classes/api/services/CTriggerPrototype.php*.

triggerprototype.update

Description

`object triggerprototype.update(object/array triggerPrototypes)`

This method allows you to update existing trigger prototypes.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Trigger prototype properties to be updated.

The `triggerid` property must be defined for each trigger prototype, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard trigger prototype properties the method accepts the following parameters.

Parameter	Type	Description
dependencies	array	Triggers and trigger prototypes that the trigger prototype is dependent on. The triggers must have only the `triggerid` property defined.
tags	array	Trigger prototype tags.

Return values

`(object)` Returns an object containing the IDs of the updated trigger prototypes under the `triggerids` property.

Examples

Enabling a trigger prototype

Enable a trigger prototype, that is, set its status to "0".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "triggerprototype.update",
    "params": {
        "triggerid": "13938",
        "status": 0
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "triggerids": [
            "13938"
        ]
    },
    "id": 1
}
```

Replacing trigger prototype tags

Replace tags for one trigger prototype.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "triggerprototype.update",
    "params": {
        "triggerid": "17373",
        "tags": [
            {
                "tag": "volume",
                "value": "{#FSNAME}"
            },
            {
                "tag": "type",
                "value": "{#FSTYPE}"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "triggerids": [
            "17373"
        ]
    },
    "id": 1
}
```

Source

CTriggerPrototype::update() in *ui/include/classes/api/services/CTriggerPrototype.php*.

User

This class is designed to work with users.

Object references:

- User
- Media

Available methods:

- user.checkauthentication - check and prolong user sessions
- user.create - create new users
- user.delete - delete users
- user.get - retrieve users
- user.login - log in to the API
- user.logout - log out of the API
- user.provision - provision LDAP users
- user.resettotp - reset user TOTP secrets
- user.unblock - unblock users
- user.update - update users

User object

The following objects are directly related to the `user` API.

User

The user object has the following properties.

Property	Type	Description
userid	ID	ID of the user. Property behavior: - *read-only* - *required* for update operations
username	string	User's name. Property behavior: - *required* for create operations - *read-only* for provisioned users if the user is linked to a user directory (`userdirectoryid` is set to a valid value that is not "0"), and user directory provisioning status is enabled (`provision_status` of User directory object is set to "1"), and authentication status of all LDAP or SAML provisioning is enabled (`ldap_jit_status` of Authentication object is set to "Enabled for configured LDAP IdPs" or `saml_jit_status` of Authentication object is set to "Enabled for configured SAML IdPs")

Property	Type	Description
passwd	string	User's password. The value of this parameter can be an empty string if the user is linked to a user directory. Property behavior: - *write-only*
roleid	ID	ID of the role of the user. Note that users without a role can log into Zabbix only using LDAP or SAML authentication, provided their LDAP/SAML information matches the user group mappings configured in Zabbix.
attempt_clock	timestamp	Time of the last unsuccessful login attempt. Property behavior: - *read-only*
attempt_failed	integer	Recent failed login attempt count. Property behavior: - *read-only*
attempt_ip	string	IP address from where the last unsuccessful login attempt came from. Property behavior: - *read-only*
autologin	integer	Whether to enable auto-login. Possible values: 0 - *(default)* auto-login disabled; 1 - auto-login enabled.
autologout	string	User session life time. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 2h). If set to 0s, the session will never expire. Default: 15m.
default_maintenance_period	string	Default maintenance period for a maintenance created via shortcut. Accepts seconds or time unit with suffix (300s, 5m, 1h, 1d, 1w). The given value will be rounded down to whole minutes. Default: 1h.
lang	string	Language code of the user's language, for example, en_US. Default: default - system default.
name	string	Name of the user.
provisioned	integer	Whether the user has been provisioned. Possible values: 0 - not provisioned; 1 - provisioned. Property behavior: - *read-only*
refresh	string	Automatic refresh period. Accepts a number of seconds (e.g., 1, 60, 300), or a value with a time suffix (e.g., 30s, 1m, 1h). Default: 30s.
rows_per_page	integer	Amount of object rows to show per page. Default: 50.
surname	string	Surname of the user.

Property	Type	Description
theme	string	User's theme. Possible values: `default` - *(default)* system default; `blue-theme` - Blue; `blue-classic-theme` - Blue (classic); `dark-theme` - Dark; `dark-blue-theme` - Dark blue; `dark-classic-theme` - Dark (classic); `hc-light` - High-contrast light; `hc-dark` - High-contrast dark.
ts_provisioned	timestamp	Time when the latest provisioning operation was made. Property behavior: - *read-only* - *supported* for *Super admin* type users
url	string	URL of the page to redirect the user to after logging in.
userdirectoryid	ID	ID of the user directory that the user in linked to. Used for provisioning (creating or updating), as well as to login a user that is linked to a user directory. For login operations the value of this property will have priority over the `userdirectoryid` property of user groups that the user belongs to. Default: 0. Property behavior: - *read-only* - *supported* for *Super admin* type users
timezone	string	User's time zone, for example, `Europe/London`, `UTC`. Default: `default` - system default. For the full list of supported time zones please refer to PHP documentation.

Media

The media object has the following properties.

Property	Type	Description
mediaid	ID	ID of the user's media. Property behavior: - *read-only*
mediatypeid	ID	ID of the media type used by the user's media. Property behavior: - *required*

Property	Type	Description
sendto	string/array	Address, user name or other identifier of the recipient. If `type` of Media type is set to ”Email” or ”Push”, values are represented as array. For ”Push”, the array can contain either a wildcard (*), or one or more mobile device IDs. If `type` of Media type is set to ”Script”, ”SMS”, or ”Webhook”, value is represented as a string. Property behavior: - *required*
active	integer	Whether the media is enabled. Possible values: 0 - *(default)* enabled; 1 - disabled.
severity	integer	Trigger severities to send notifications about. Possible bitmap values: 1 - Not classified; 2 - Information; 4 - Warning; 8 - Average; 16 - High; 32 - Disaster. This is a bitmask field; any sum of possible bitmap values is acceptable (for example, 48 for Average, High, and Disaster). Default: 63.
period	string	Time when the notifications can be sent as a time period or user macros separated by a semicolon. Default: 1-7,00:00-24:00.
provisioned	integer	Whether the user has been provisioned. Possible values: 0 - not provisioned; 1 - provisioned.
userdirectory_mediaid	ID	User directory media mapping ID for provisioned media. Property behavior: - *read-only* - *supported* for *Super admin* type users

user.checkAuthentication

Description

`object user.checkAuthentication`

This method checks and prolongs the user session.

> **Attention:**
> Calling the `user.checkAuthentication` method using the parameter `sessionid` prolongs the user session by default.

Parameters

The method accepts the following parameters.

Parameter	Type	Description
extend	boolean	Whether to prolong the user session. Default value: "true". Setting the value to "false" allows you to check the user session without prolonging it.
sessionid	string	Parameter behavior: - *supported* if `sessionid` is set User authentication token. Parameter behavior: - *required* if `token` is not set
token	string	User API token. Parameter behavior: - *required* if `sessionid` is not set

Return values

`(object)` Returns an object containing information about the user.

Additionally to the standard user properties, the following information is returned.

Property	Type	Description
auth_type	integer	Default authentication for the user. Refer to the `authentication_type` property of the Authentication object for a list of possible values.
debug_mode	integer	Whether debug mode is enabled or disabled for the user. Refer to the `debug_mode` property of the User group object for a list of possible values.
deprovisioned	boolean	Whether the user belongs to a deprovisioned users group.
gui_access	string	User's authentication method to the frontend. Refer to the `gui_access` property of the User group object for a list of possible values.
secret	string	Random 32 characters string. Is generated on user login. Property `secret` is not returned if the user session is checked using an API token.
sessionid	string	Authentication token, which must be used in the following API requests. Property `sessionid` is not returned if the user session is checked using an API token.
type	integer	User type. Refer to the `type` property of the Role object for a list of possible values.
userip	string	IP address of the user.

Examples

Check authentication using authentication token

Check and prolong a user session using the user authentication token, and return additional information about the user.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.checkAuthentication",
```

```
    "params": {
        "sessionid": "673b8ba11562a35da902c66cf5c23fa2"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userid": "1",
        "username": "Admin",
        "name": "Zabbix",
        "surname": "Administrator",
        "url": "",
        "autologin": "1",
        "autologout": "0",
        "lang": "ru_RU",
        "refresh": "0",
        "theme": "default",
        "attempt_failed": "0",
        "attempt_ip": "127.0.0.1",
        "attempt_clock": "1355919038",
        "rows_per_page": "50",
        "timezone": "Europe/Riga",
        "roleid": "3",
        "userdirectoryid": "0",
        "ts_provisioned": "0",
        "type": 3,
        "userip": "127.0.0.1",
        "debug_mode": 0,
        "gui_access": "0",
        "deprovisioned": false,
        "auth_type": 0,
        "sessionid": "673b8ba11562a35da902c66cf5c23fa2",
        "secret": "0e329b933e46984e49a5c1051ecd0751"
    },
    "id": 1
}
```

Check authentication using API token

Check a user session using the user API token, and return additional information about the user.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.checkAuthentication",
    "params": {
        "token": "00aff470e07c12d707e50d98cfe39edef9e6ec349c14728dbdfbc8ddc5ea3eae"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userid": "1",
        "username": "Admin",
        "name": "Zabbix",
        "surname": "Administrator",
        "url": "",
```

```
        "autologin": "1",
        "autologout": "0",
        "lang": "ru_RU",
        "refresh": "0",
        "theme": "default",
        "attempt_failed": "0",
        "attempt_ip": "127.0.0.1",
        "attempt_clock": "1355919338",
        "rows_per_page": "50",
        "timezone": "Europe/Riga",
        "roleid": "3",
        "userdirectoryid": "0",
        "ts_provisioned": "0",
        "debug_mode": 0,
        "deprovisioned": false,
        "gui_access": "1",
        "mfaid": 0,
        "auth_type": 0,
        "type": 3,
        "userip": "127.0.0.1"
    },
    "id": 1
}
```


Source

CUser::checkAuthentication() in *ui/include/classes/api/services/CUser.php*.

user.create

Description

`object user.create(object/array users)`

This method allows you to create new users.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Note:
The strength of user password is validated according to the password policy rules defined by Authentication API. See authentication object for more information.

Parameters

`(object/array)` Users to create.

Additionally to the standard user properties, the method accepts the following parameters.

Parameter	Type	Description
usrgrps	array	User groups to add the user to. The user groups must have only the `usrgrpid` property defined.
medias	array	User media to be created.

Return values

`(object)` Returns an object containing the IDs of the created users under the `userids` property. The order of the returned IDs matches the order of the passed users.

Examples

Creating a user

Create a new user, add him to a user group and create a new media for him.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.create",
    "params": {
        "username": "John",
        "passwd": "Doe12345",
        "roleid": "5",
        "usrgrps": [
            {
                "usrgrpid": "7"
            }
        ],
        "medias": [
            {
                "mediatypeid": "1",
                "sendto": [
                    "support@company.com"
                ],
                "active": 0,
                "severity": 63,
                "period": "1-7,00:00-24:00"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userids": [
            "12"
        ]
    },
    "id": 1
}
```

See also

- Authentication
- Media
- User group
- Role

Source

CUser::create() in *ui/include/classes/api/services/CUser.php*.

user.delete

Description

`object user.delete(array users)`

This method allows you to delete users.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

(array) IDs of users to delete.

Return values

(object) Returns an object containing the IDs of the deleted users under the userids property.

Examples

Deleting multiple users

Delete two users.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.delete",
    "params": [
        "1",
        "5"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userids": [
            "1",
            "5"
        ]
    },
    "id": 1
}
```

Source

CUser::delete() in *ui/include/classes/api/services/CUser.php*.

user.get

Description

integer/array user.get(object parameters)

The method allows you to retrieve users according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

(object) Parameters defining the desired output.

The method supports the following parameters.

> **Note:**
> When requesting user media, permissions, or role information, *Admin* and *User* type users may retrieve data only about their own user. For an example, see Retrieving users as *Admin*.

Parameter	Type	Description
mediaids	ID/array	Return only users that use the given media.
mediatypeids	ID/array	Return only users that use the given media types.
userids	ID/array	Return only users with the given IDs.
usrgrpids	ID/array	Return only users that belong to the given user groups.

Parameter	Type	Description
getAccess	flag	Adds additional information about user permissions. Adds the following properties for each user: `gui_access` - `(integer)` user's frontend authentication method. Refer to the `gui_access` property of the user group object for a list of possible values. `debug_mode` - `(integer)` indicates whether debug is enabled for the user. Possible values: 0 - debug disabled, 1 - debug enabled. `users_status` - `(integer)` indicates whether the user is disabled. Possible values: 0 - user enabled, 1 - user disabled.
selectMedias	query	Return media used by the user in the `medias` property.
selectMediatypes	query	Return media types used by the user in the `mediatypes` property. See `mediatype.get` for restrictions based on user type.
selectUsrgrps	query	Return user groups that the user belongs to in the `usrgrps` property. See `usergroup.get` for restrictions based on user type.
selectRole	query	Return user role in the `role` property.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values to match against. Does not support properties of `text` data type. Possible User object properties for *Admin* and *User* type users when requesting data on users in their user group: `userid`, `name`, `surname`, `username`.
output	query	User object properties to be returned. *Admin* and *User* type users may retrieve only the following properties: - For their own user: `userid`, `attempt_clock`, `attempt_failed`, `attempt_ip`, `autologin`, `autologout`, `lang`, `name`, `provisioned`, `refresh`, `roleid`, `rows_per_page`, `surname`, `theme`, `timezone`, `url`, `username`. - For users in their user group: `userid`, `name`, `surname`, `username`. Default: `extend`.
search	object	Return results that match the given pattern (case-insensitive). Accepts an object, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a `LIKE "%…%"` search. Supports only properties of `string` and `text` data type. Possible User object properties for *Admin* and *User* type users when requesting data on users in their user group: `name`, `surname`, `username`.
sortfield	string/array	Sort the result by the given properties. Possible values: `userid`, `username`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
limit	integer	
preservekeys	boolean	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

(integer/array) Returns either:

- an array of objects;
- the count of retrieved objects, if the countOutput parameter has been used.

Examples

Retrieving users

Retrieve all of the configured users.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.get",
    "params": {
        "output": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "userid": "1",
            "username": "Admin",
            "name": "Zabbix",
            "surname": "Administrator",
            "url": "",
            "autologin": "1",
            "autologout": "0",
            "lang": "en_US",
            "default_maintenance_period": "1h",
            "refresh": "30s",
            "theme": "default",
            "attempt_failed": "0",
            "attempt_ip": "",
            "attempt_clock": "0",
            "rows_per_page": "50",
            "timezone": "default",
            "roleid": "3",
            "userdirectoryid": "0",
            "ts_provisioned": "0",
            "provisioned": "0"
        },
        {
            "userid": "2",
            "username": "guest",
            "name": "",
            "surname": "",
            "url": "",
            "autologin": "0",
            "autologout": "15m",
            "lang": "default",
            "default_maintenance_period": "1h",
            "refresh": "30s",
            "theme": "default",
            "attempt_failed": "0",
            "attempt_ip": "",
            "attempt_clock": "0",
            "rows_per_page": "50",
```

```
            "timezone": "default",
            "roleid": "4",
            "userdirectoryid": "0",
            "ts_provisioned": "0",
            "provisioned": "0"
        }
    ],
    "id": 1
}
```

Retrieving users as *Admin*

As an *Admin* type user, retrieve detailed data about your own user and limited data for users in your user group.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.get",
    "params": {
        "output": "extend",
        "getAccess": true,
        "selectMedias": "extend",
        "selectMediatypes": "extend",
        "selectUsrgrps": "extend",
        "selectRole": "extend"
    },
    "id": 1
}
```

Response:

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "userid": "1",
      "username": "Admin",
      "name": "Zabbix",
      "surname": "Administrator",
      "url": "",
      "autologin": "1",
      "autologout": "0",
      "lang": "default",
      "default_maintenance_period": "1h",
      "refresh": "30s",
      "theme": "default",
      "attempt_failed": "0",
      "attempt_ip": "",
      "attempt_clock": "0",
      "rows_per_page": "50",
      "timezone": "default",
      "roleid": "3",
      "userdirectoryid": "0",
      "ts_provisioned": "0",
      "provisioned": "0",
      "gui_access": "1",
      "debug_mode": "0",
      "users_status": "0",
      "usrgrps": [
        {
          "usrgrpid": "7",
          "name": "Zabbix administrators",
          "gui_access": "0",
          "users_status": "0",
```

```
        "debug_mode": "0",
        "userdirectoryid": "0",
        "mfa_status": "0",
        "mfaid": "0",
        "proxy_mode": "0",
        "proxy_group_mode": "0"
      },
      {
        "usrgrpid": "13",
        "name": "Internal",
        "gui_access": "1",
        "users_status": "0",
        "debug_mode": "0",
        "userdirectoryid": "0",
        "mfa_status": "0",
        "mfaid": "0",
        "proxy_mode": "0",
        "proxy_group_mode": "0"
      }
    ],
    "medias": [],
    "mediatypes": [],
    "role": {
      "roleid": "3",
      "name": "Super admin role",
      "type": "3",
      "readonly": "1"
    }
  },
  {
    "userid": "2",
    "username": "guest",
    "name": "",
    "surname": "",
    "url": "",
    "autologin": "0",
    "autologout": "15m",
    "lang": "default",
    "default_maintenance_period": "1h",
    "refresh": "30s",
    "theme": "default",
    "attempt_failed": "0",
    "attempt_ip": "",
    "attempt_clock": "0",
    "rows_per_page": "50",
    "timezone": "default",
    "roleid": "4",
    "userdirectoryid": "0",
    "ts_provisioned": "0",
    "provisioned": "0",
    "gui_access": "1",
    "debug_mode": "0",
    "users_status": "1",
    "usrgrps": [
      {
        "usrgrpid": "8",
        "name": "Guests",
        "gui_access": "0",
        "users_status": "0",
        "debug_mode": "0",
        "userdirectoryid": "0",
        "mfa_status": "0",
```

```
          "mfaid": "0",
          "proxy_mode": "0",
          "proxy_group_mode": "0"
        },
        {
          "usrgrpid": "9",
          "name": "Disabled",
          "gui_access": "0",
          "users_status": "1",
          "debug_mode": "0",
          "userdirectoryid": "0",
          "mfa_status": "0",
          "mfaid": "0",
          "proxy_mode": "0",
          "proxy_group_mode": "0"
        },
        {
          "usrgrpid": "13",
          "name": "Internal",
          "gui_access": "1",
          "users_status": "0",
          "debug_mode": "0",
          "userdirectoryid": "0",
          "mfa_status": "0",
          "mfaid": "0",
          "proxy_mode": "0",
          "proxy_group_mode": "0"
        }
      ],
      "medias": [],
      "mediatypes": [],
      "role": {
        "roleid": "4",
        "name": "Guest role",
        "type": "1",
        "readonly": "0"
      }
    }
  ],
  "id": 1
}
```


Retrieving user data

Retrieve data of a user with ID ”2”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.get",
    "params": {
        "output": ["userid", "username"],
        "selectRole": "extend",
        "userids": "2"
    },
    "id": 1
}
```

Response:


```
{
  "jsonrpc": "2.0",
  "result": [
    {
```

```
      "userid": "2",
      "username": "guest",
      "role": {
        "roleid": "4",
        "name": "Guest role",
        "type": "1",
        "readonly": "0"
      }
    }
  ],
  "id": 1
}
```


See also

- Media
- Media type
- User group
- Role

Source

CUser::get() in *ui/include/classes/api/services/CUser.php*.

user.login

Description

`string/object user.login(object parameters)`

This method allows you to log in to the API and generate an authentication token.

Warning:
When using this method, you also need to do user.logout to prevent the generation of a large number of open session records.

Attention:
This method is only available to unauthenticated users who do not belong to any user group with enabled multi-factor authentication.

Parameters

`(object)` Parameters containing the user name and password.

The method accepts the following parameters.

Parameter	Type	Description
password	string	User password. Parameter behavior: - *required*
username	string	User name. Parameter behavior: - *required*
userData	flag	Return information about the authenticated user.

Return values

`(string/object)` If the `userData` parameter is used, returns an object containing information about the authenticated user.

Additionally to the standard user properties, the following information is returned:

Property	Type	Description
auth_type	integer	Default authentication for the user. Refer to the `authentication_type` property of the Authentication object for a list of possible values.
debug_mode	integer	Whether debug mode is enabled or disabled for the user. Refer to the `debug_mode` property of the User group object for a list of possible values.
deprovisioned	boolean	Whether the user belongs to a deprovisioned users group.
gui_access	string	User's authentication method to the frontend. Refer to the `gui_access` property of the User group object for a list of possible values.
mfaid	integer	ID of the MFA method to use for the user during login. Returns "0" if MFA is disabled globally or for all user groups that the user belongs to.
secret	string	Random 32 characters string. Is generated on user login.
sessionid	string	Authentication token, which must be used in the following API requests.
type	integer	User type. Refer to the `type` property of the Role object for a list of possible values.
userip	string	IP address of the user.

Note:
If a user has been successfully authenticated after one or more failed attempts, the method will return the current values for the `attempt_clock`, `attempt_failed` and `attempt_ip` properties and then reset them.

If the `userData` parameter is not used, the method returns an authentication token that is required for authentication.

Examples

Authenticating a user

Authenticate a user.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.login",
    "params": {
        "username": "Admin",
        "password": "zabbix"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": "0424bd59b807674191e7d77572075f33",
    "id": 1
}
```

Requesting authenticated user's information

Authenticate and return additional information about the user.

Request:

```
{
    "jsonrpc": "2.0",
```

```
    "method": "user.login",
    "params": {
        "username": "Admin",
        "password": "zabbix",
        "userData": true
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userid": "1",
        "username": "Admin",
        "name": "Zabbix",
        "surname": "Administrator",
        "url": "",
        "autologin": "1",
        "autologout": "0",
        "lang": "ru_RU",
        "refresh": "0",
        "theme": "default",
        "attempt_failed": "0",
        "attempt_ip": "127.0.0.1",
        "attempt_clock": "1355919038",
        "rows_per_page": "50",
        "timezone": "Europe/Riga",
        "roleid": "3",
        "userdirectoryid": "0",
        "type": 3,
        "userip": "127.0.0.1",
        "debug_mode": 0,
        "gui_access": "0",
        "mfaid": "1",
        "deprovisioned": false,
        "auth_type": 0,
        "sessionid": "5b56eee8be445e98f0bd42b435736e42",
        "secret": "cd0ba923319741c6586f3d866423a8f4"
    },
    "id": 1
}
```

See also

- user.logout

Source

CUser::login() in *ui/include/classes/api/services/CUser.php*.

user.logout

Description

`string/object user.logout(array)`

This method allows you to log out of the API and invalidates the current authentication token.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` The method accepts an empty array.

Return values

`(boolean)` Returns `true` if the user has been logged out successfully.

Examples

Logging out

Log out from the API.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.logout",
    "params": [],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": true,
    "id": 1
}
```

See also

- user.login

Source

CUser::login() in *ui/include/classes/api/services/CUser.php*.

user.provision

Description

`object user.provision(object/array users)`

This method allows you to provision LDAP users.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of users to provision.

Return values

`(object)` Returns an object containing the IDs of the provisioned users under the `userids` property.

Examples

Provisioning multiple users

Provision two users.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.provision",
    "params": [
        "1",
        "5"
    ],
```

```
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userids": [
            "1",
            "5"
        ]
    },
    "id": 1
}
```

Source

CUser::provision() in *ui/include/classes/api/services/CUser.php*.

user.resettotp

Description

`object user.resettotp(object/array users)`

This method allows you to reset user TOTP secrets.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

(`array`) IDs of users for which to reset TOTP secrets.

> **Note:**
> User sessions for the specified users will also be deleted (except for the user sending the request).

Return values

(`object`) Returns an object containing the IDs of the users for which TOTP secrets have been reset, under the `userids` property.

Examples

Resetting TOTP secrets for multiple users

Reset TOTP secrets for two users.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.resettotp",
    "params": [
        "1",
        "5"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userids": [
            "1",
```

```
            "5"
        ]
    },
    "id": 1
}
```

See also

- MFA object

Source

CUser::resettotp() in *ui/include/classes/api/services/CUser.php*.

user.unblock

Description

`object user.unblock(array userids)`

This method allows you to unblock users.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of users to unblock.

Return values

`(object)` Returns an object containing the IDs of the unblocked users under the `userids` property.

Examples

Unblocking multiple users

Unblock two users.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.unblock",
    "params": [
        "1",
        "5"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userids": [
            "1",
            "5"
        ]
    },
    "id": 1
}
```

Source

CUser::unblock() in *ui/include/classes/api/services/CUser.php*.

user.update

Description

`object user.update(object/array users)`

This method allows you to update existing users.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

> **Note:**
> The strength of user password is validated according the password policy rules defined by Authentication API. See Authentication API for more information.

Parameters

`(object/array)` User properties to be updated.

The `userid` property must be defined for each user, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard user properties, the method accepts the following parameters.

Parameter	Type	Description
current_passwd	string	User's current password. The value of this parameter can be an empty string if the user is linked to a user directory. Parameter behavior: - *write-only* - *required* if `passwd` of User object is set and user changes own user password
usrgrps	array	User groups to replace existing user groups. The user groups must have only the `usrgrpid` property defined.
medias	array	User media to replace existing, non-provisioned media. Provisioned media can be omitted when updating media.

Return values

`(object)` Returns an object containing the IDs of the updated users under the `userids` property.

Examples

Renaming a user

Rename a user to John Doe.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.update",
    "params": {
        "userid": "1",
        "name": "John",
        "surname": "Doe"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userids": [
            "1"
        ]
    },
    "id": 1
}
```

Changing user role

Change a role of a user.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "user.update",
    "params": {
        "userid": "12",
        "roleid": "6"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userids": [
            "12"
        ]
    },
    "id": 1
}
```

See also

- Authentication

Source

CUser::update() in *ui/include/classes/api/services/CUser.php*.

User directory

This class is designed to work with user directories.

Object references:

- User directory
  - Media type mappings
  - Provisioning groups mappings

Available methods:

- userdirectory.create - create new user directory
- userdirectory.delete - delete user directory
- userdirectory.get - retrieve user directory
- userdirectory.update - update user directory
- userdirectory.test - test user directory connection

User directory object

The following objects are directly related to the `userdirectory` API.

User directory

The user directory object has the following properties.

Property	Type	Description
userdirectoryid	ID	ID of the user directory. If a user directory is deleted, the value of the User object property `userdirectoryid` is set to ”0” for all users that are linked to the deleted user directory. Property behavior: - *read-only* - *required* for update operations
idp_type	integer	Type of the authentication protocol used by the identity provider for the user directory. Note that only one user directory of type SAML can exist. Possible values: 1 - User directory of type LDAP; 2 - User directory of type SAML. Property behavior: - *required* for create operations
group_name	string	LDAP/SAML user directory attribute that contains the group name used to map groups between the LDAP/SAML user directory and Zabbix. Example: *cn* Property behavior: - *required* if `provision_status` is set to ”Enabled” and `saml_jit_status` of Authentication object is set to ”Enabled for configured SAML IdPs”
user_username	string	LDAP/SAML user directory attribute (also SCIM attribute if `scim_status` is set to ”SCIM provisioning is enabled”) that contains the user’s name which is used as the value for the User object property `name` when the user is provisioned. Examples: *cn*, *commonName*, *displayName*, *name*
user_lastname	string	LDAP/SAML user directory attribute (also SCIM attribute if `scim_status` is set to ”SCIM provisioning is enabled”) that contains the user’s last name which is used as the value for the User object property `surname` when the user is provisioned. Examples: *sn*, *surname*, *lastName*
provision_status	integer	Provisioning status of the user directory. Possible values: 0 - *(default)* Disabled (provisioning of users created by this user directory is disabled); 1 - Enabled (provisioning of users created by this user directory is enabled; additionally, the status of LDAP or SAML provisioning (`ldap_jit_status` or `saml_jit_status` of Authentication object) must be enabled).
provision_groups	array	Array of provisioning groups mappings objects for mapping LDAP/SAML user group pattern to Zabbix user group and user role. Property behavior: - *required* if `provision_status` is set to ”Enabled”
provision_media	array	Array of media type mappings objects for mapping user’s LDAP/SAML media attributes (e.g., email) to Zabbix user media for sending notifications.

Property	Type	Description
LDAP-specific properties:		
name	string	Unique name of the user directory. Property behavior: - *required* if `idp_type` is set to "User directory of type LDAP"
host	string	Host name, IP or URI of the LDAP server. URI must contain schema (`ldap://` or `ldaps://`), host, and port (optional). Examples: *host.example.com* *127.0.0.1* *ldap://ldap.example.com:389* Property behavior: - *required* if `idp_type` is set to "User directory of type LDAP"
port	integer	Port of the LDAP server. Property behavior: - *required* if `idp_type` is set to "User directory of type LDAP"
base_dn	string	LDAP user directory base path to user accounts. Examples: *ou=Users,dc=example,dc=org* *ou=Users,ou=system* (for OpenLDAP) *DC=company,DC=com* (for Microsoft Active Directory) *uid=%{user},dc=example,dc=com* (for direct user binding; placeholder "*%{user}*" is mandatory) Property behavior: - *required* if `idp_type` is set to "User directory of type LDAP"
search_attribute	string	LDAP user directory attribute by which to identify the user account from the information provided in the login request. Examples: *uid* (for OpenLDAP) *sAMAccountName* (for Microsoft Active Directory) Property behavior: - *required* if `idp_type` is set to "User directory of type LDAP"
bind_dn	string	LDAP server account for binding and searching over the LDAP server. For direct user binding and anonymous binding, `bind_dn` must be empty. Examples: *uid=ldap_search,ou=system* (for OpenLDAP) *CN=ldap_search,OU=user_group,DC=company,DC=com* (for Microsoft Active Directory) *CN=Admin,OU=Users,OU=Zabbix,DC=zbx,DC=local* Property behavior: - *supported* if `idp_type` is set to "User directory of type LDAP"
bind_password	string	LDAP password of the account for binding and searching over the LDAP server. For direct user binding and anonymous binding, `bind_password` must be empty. Property behavior: - *supported* if `idp_type` is set to "User directory of type LDAP"

Property	Type	Description
description	string	Description of the user directory. Property behavior: - *supported* if `idp_type` is set to "User directory of type LDAP"
group_basedn	string	LDAP user directory base path to groups; used to configure a user membership check in the LDAP user directory. Ignored when provisioning a user if `group_membership` is set. Example: *ou=Groups,dc=example,dc=com* Property behavior: - *supported* if `idp_type` is set to "User directory of type LDAP"
group_filter	string	Filter string for retrieving LDAP user directory groups that the user is a member of; used to configure a user membership check in the LDAP user directory. Ignored when provisioning a user if `group_membership` is set. Supported `group_filter` placeholders: *%{attr}* - search attribute (replaced by the `search_attribute` property value); *%{groupattr}* - group attribute (replaced by the `group_member` property value); *%{host}* - host name, IP or URI of the LDAP server (replaced by the `host` property value); *%{user}* - Zabbix user username. Default: *(%{groupattr}=%{user})* Examples: - *(member=uid=%{ref},ou=Users,dc=example,dc=com)* will match "User1" if an LDAP group object contains the "*member*" attribute with the value "*uid=User1,ou=Users,dc=example,dc=com*", and will return the group that "User1" is a member of; - *(%{groupattr}=cn=%{ref},ou=Users,ou=Zabbix,DC=example,DC=com)* will match "User1" if an LDAP group object contains the attribute specified in the `group_member` property with the value "*cn=User1,ou=Users,ou=Zabbix,DC=example,DC=com*", and will return the group that "User1" is a member of. Property behavior: - *supported* if `idp_type` is set to "User directory of type LDAP"
group_member	string	LDAP user directory attribute that contains information about the group members; used to configure a user membership check in the LDAP user directory. Ignored when provisioning a user if `group_membership` is set. Property behavior: - *supported* if `idp_type` is set to "User directory of type LDAP"
group_membership	string	LDAP user directory attribute that contains information about the groups that a user belongs to. Example: *memberOf* Property behavior: - *supported* if `idp_type` is set to "User directory of type LDAP"

Property	Type	Description
search_filter	string	Custom filter string used to locate and authenticate a user in an LDAP user directory based on the information provided in the login request. Supported `search_filter` placeholders: *%{attr}* - search attribute name (e.g., *uid*, *sAMAccountName*); *%{user}* - Zabbix user username. Default: *(%{attr}=%{user})* Property behavior: - *supported* if `idp_type` is set to "User directory of type LDAP"
start_tls	integer	LDAP server configuration option that allows the communication with the LDAP server to be secured using Transport Layer Security (TLS). Note that `start_tls` must be set to "Disabled" for hosts using the `ldaps://` protocol. Possible values: 0 - *(default)* Disabled; 1 - Enabled. Property behavior: - *supported* if `idp_type` is set to "User directory of type LDAP"
user_ref_attr	string	LDAP user directory attribute used to reference a user object. The value of `user_ref_attr` is used to get values from the specified attribute in the user directory and place them instead of the *%{ref}* placeholder in the `group_filter` string. Examples: *cn*, *uid*, *member*, *uniqueMember* Property behavior: - *supported* if `idp_type` is set to "User directory of type LDAP"
SAML-specific properties:		
idp_entityid	string	URI that identifies the identity provider and is used to communicate with the identity provider in SAML messages. Example: *https://idp.example.com/idp* Property behavior: - *required* if `idp_type` is set to "User directory of type SAML"
sp_entityid	string	URL or any string that identifies the identity provider's service provider. Examples: *https://idp.example.com/sp* *zabbix* Property behavior: - *required* if `idp_type` is set to "User directory of type SAML"
username_attribute	string	SAML user directory attribute (also SCIM attribute if `scim_status` is set to "SCIM provisioning is enabled") that contains the user's username which is compared with the value of the User object property `username` when authenticating. Examples: *uid*, *userprincipalname*, *samaccountname*, *username*, *userusername*, *urn:oid:0.9.2342.19200300.100.1.1*, *urn:oid:1.3.6.1.4.1.5923.1.1.1.13*, *urn:oid:0.9.2342.19200300.100.1.44* Property behavior: - *required* if `idp_type` is set to "User directory of type SAML"

Property	Type	Description
sso_url	string	URL of the identity provider's SAML single sign-on service, to which Zabbix will send the SAML authentication requests. Example: *http://idp.example.com/idp/sso/saml* Property behavior: - *required* if `idp_type` is set to "User directory of type SAML"
slo_url	string	URL of the identity provider's SAML single log-out service, to which Zabbix will send the SAML logout requests. Example: *https://idp.example.com/idp/slo/saml* Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML"
encrypt_nameid	integer	Whether the SAML name ID should be encrypted. Possible values: 0 - *(default)* Do not encrypt name ID; 1 - Encrypt name ID. Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML"
encrypt_assertions	integer	Whether the SAML assertions should be encrypted. Possible values: 0 - *(default)* Do not encrypt assertions; 1 - Encrypt assertions. Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML"
nameid_format	string	Name ID format of the SAML identity provider's service provider. Examples: *urn:oasis:names:tc:SAML:2.0:nameid-format:persistent* *urn:oasis:names:tc:SAML:2.0:nameid-format:transient* *urn:oasis:names:tc:SAML:2.0:nameid-format:kerberos* *urn:oasis:names:tc:SAML:2.0:nameid-format:entity* Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML"
scim_status	integer	Whether SCIM provisioning for SAML is enabled or disabled. Possible values: 0 - *(default)* SCIM provisioning is disabled; 1 - SCIM provisioning is enabled. Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML"
sign_assertions	integer	Whether the SAML assertions should be signed with a SAML signature. Possible values: 0 - Do not sign assertions; 1 - *(default)* Sign assertions. Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML"

Property	Type	Description
sign_authn_requests	integer	Whether the SAML AuthN requests should be signed with a SAML signature. Possible values: 0 - Do not sign AuthN requests; 1 - *(default)* Sign AuthN requests. Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML"
sign_messages	integer	Whether the SAML messages should be signed with a SAML signature. Possible values: 0 - Do not sign messages; 1 - *(default)* Sign messages. Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML"
sign_logout_requests	integer	Whether the SAML logout requests should be signed with a SAML signature. Possible values: 0 - Do not sign logout requests; 1 - *(default)* Sign logout requests. Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML"
sign_logout_responses	integer	Whether the SAML logout responses should be signed with a SAML signature. Possible values: 0 - Do not sign logout responses; 1 - *(default)* Sign logout responses. Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML"
idp_certificate	string	Contents of the Service Provider (SP) certificate for SAML single sign-on (SSO) service setup. Examples: *-----BEGIN CERTIFICATE-----* * MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEA7... ...more encoded data... -----END CERTIFICATE-----* Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML" and `$SSO['CERT_STORAGE']` is set to `database` in the frontend configuration file - *write-only*

Property	Type	Description
sp_private_key	string	Contents of the Service Provider (SP) Private Key for SAML single sign-on (SSO) service setup. Provides secure authentication and data exchange with the Identity Provider (IdP). Examples: *-----BEGIN CERTIFICATE----- MIIEvQIBADANBgkqhkiG9w0BA... ...more encoded data... -----END CERTIFICATE-----* Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML" and `$SSO['CERT_STORAGE']` is set to `database` in the frontend configuration file - *write-only*
sp_certificate	string	Contents of the Service Provider (SP) certificate for SAML single sign-on (SSO) service setup. Examples: *-----BEGIN CERTIFICATE----- MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEA7... ...more encoded data... -----END CERTIFICATE-----* Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML" and `$SSO['CERT_STORAGE']` is set to `database` in the frontend configuration file - *write-only*
idp_certificate_hash	string	The md5 hash of idp_certificate value. Will return empty string for empty idp_certificate. Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML" and `$SSO['CERT_STORAGE']` is set to `database` in the frontend configuration file - *read-only*
sp_private_key_hash	string	The md5 hash of sp_private_key value. Will return empty string for empty sp_private_key. Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML" and `$SSO['CERT_STORAGE']` is set to `database` in the frontend configuration file - *read-only*
sp_certificate_hash	string	The md5 hash of sp_certificate value. Will return empty string for empty sp_certificate. Property behavior: - *supported* if `idp_type` is set to "User directory of type SAML" and `$SSO['CERT_STORAGE']` is set to `database` in the frontend configuration file - *read-only*

Media type mappings

The media type mappings object has the following properties.

Property	Type	Description
userdirectory_mediaid	ID	Media type mapping ID. Property behavior: - *read-only*
name	string	Visible name in the list of media type mappings. Property behavior: - *required*
mediatypeid	ID	ID of the media type to be created; used as the value for the Media object property `mediatypeid`. Property behavior: - *required*
attribute	string	LDAP/SAML user directory attribute (also SCIM attribute if `scim_status` is set to "SCIM provisioning is enabled") that contains the user's media (e.g., *user@example.com*) which is used as the value for the Media object property `sendto`. If present in data received from the LDAP/SAML identity provider, and the value is not empty, this will trigger media creation for the provisioned user. Property behavior: - *required*
active	integer	User media `active` property value when media is created for the provisioned user. Possible values: 0 - *(default)* enabled; 1 - disabled.
severity	integer	User media `severity` property value when media is created for the provisioned user. Default: 63.
period	string	User media `period` property value when media is created for the provisioned user. Default: 1-7,00:00-24:00.

Provisioning groups mappings

The provisioning groups mappings has the following properties.

Property	Type	Description
name	string	Full name of a group (e.g., *Zabbix administrators*) in LDAP/SAML user directory (also SCIM if `scim_status` is set to "SCIM provisioning is enabled"). Supports the wildcard character "*". Unique across all provisioning groups mappings. Property behavior: - *required*
roleid	ID	ID of the user role to assign to the user. If multiple provisioning groups mappings are matched, the role of the highest user type (*User*, *Admin*, or *Super admin*) is assigned to the user. If there are multiple roles with the same user type, the first role (sorted in alphabetical order) is assigned to the user. Property behavior: - *required*

Property	Type	Description
user_groups	array	Array of Zabbix user group ID objects. Each object has the following properties: `usrgrpid` - (`ID`) ID of Zabbix user group to assign to the user. If multiple provisioning groups mappings are matched, Zabbix user groups of all matched mappings is assigned to the user. Property behavior: *- required*

userdirectory.create

Description

`object userdirectory.create(object/array userDirectory)`

This method allows you to create new user directories.

> **Note:**
> This method is only available to *Super admin* user type.

Parameters

(`object/array`) User directories to create.

The method accepts user directories with the standard user directory properties.

Return values

(`object`) Returns an object containing the IDs of the created user directories under the `userdirectoryids` property. The order of the returned IDs matches the order of the passed user directories.

Examples

Creating a user directory

Create a user directory to authenticate users with StartTLS over LDAP. Note that to authenticate users over LDAP, LDAP authentication must be enabled.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "userdirectory.create",
    "params": {
        "idp_type": "1",
        "name": "LDAP API server #1",
        "host": "ldap://local.ldap",
        "port": "389",
        "base_dn": "ou=Users,dc=example,dc=org",
        "bind_dn": "cn=ldap_search,dc=example,dc=org",
        "bind_password": "ldapsecretpassword",
        "search_attribute": "uid",
        "start_tls": "1"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userdirectoryids": [
            "3"
        ]
    },
```

```
    "id": 1
}
```

Creating a user directory (JIT provisioning enabled)

Create a user directory to authenticate users over LDAP (with JIT provisioning enabled). Note that to authenticate users over LDAP, LDAP authentication must be enabled.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "userdirectory.create",
    "params": {
            "idp_type": "1",
            "name": "AD server",
            "provision_status": "1",
            "description": "",
            "host": "host.example.com",
            "port": "389",
            "base_dn": "DC=zbx,DC=local",
            "search_attribute": "sAMAccountName",
            "bind_dn": "CN=Admin,OU=Users,OU=Zabbix,DC=zbx,DC=local",
            "start_tls": "0",
            "search_filter": "",
            "group_basedn": "OU=Zabbix,DC=zbx,DC=local",
            "group_name": "CN",
            "group_member": "member",
            "group_filter": "(%{groupattr}=CN=%{ref},OU=Users,OU=Zabbix,DC=zbx,DC=local)",
            "group_membership": "",
            "user_username": "givenName",
            "user_lastname": "sn",
            "user_ref_attr": "CN",
            "provision_media": [
                {
                    "name": "example.com",
                    "mediatypeid": "1",
                    "attribute": "user@example.com"
                }
            ],
            "provision_groups": [
                {
                    "name": "*",
                    "roleid": "4",
                    "user_groups": [
                        {
                            "usrgrpid": "8"
                        }
                    ]
                },
                {
                    "name": "Zabbix administrators",
                    "roleid": "2",
                    "user_groups": [
                        {
                            "usrgrpid": "7"
                        },
                        {
                            "usrgrpid": "8"
                        }
                    ]
                }
            ]
        },
```

```
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userdirectoryids": [
            "2"
        ]
    },
    "id": 1
}
```

Source

CUserDirectory::create() in *ui/include/classes/api/services/CUserDirectory.php*.

userdirectory.delete

Description

`object userdirectory.delete(array userDirectoryIds)`

This method allows you to delete user directories. User directory cannot be deleted when it is directly used for at least one user group.
 Default LDAP user directory cannot be deleted when `authentication.ldap_configured` is set to 1 or when there are more user directories left.

> **Note:**
> This method is only available to *Super admin* user type.

Parameters

`(array)` IDs of the user directories to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted user directories under the `userdirectoryids` property.

Examples

Deleting multiple user directories

Delete two user directories.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "userdirectory.delete",
    "params": [
        "2",
        "12"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userdirectoryids": [
            "2",
            "12"
        ]
    },
```

```
    "id": 1
}
```

Source

CUserDirectory::delete() in *ui/include/classes/api/services/CUserDirectory.php*.

userdirectory.get

Description

`integer/array userdirectory.get(object parameters)`

The method allows you to retrieve user directories according to the given parameters.

> **Note:**
> This method is only available to *Super admin* user types.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
userdirectoryids	ID/array	Return only user directories with the given IDs.
selectUsrgrps	query	Return a `usrgrps` property with user groups associated with a user directory. Supports `count`.
selectProvisionMedia	query	Return a `provision_media` property with media type mappings associated with a user directory.
selectProvisionGroups	query	Return a `provision_groups` property with provisioning groups mappings associated with a user directory.
sortfield	string/array	Sort the result by the given properties. Possible values: `name`.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names, and the values are either a single value or an array of values. Supports properties: `userdirectoryid`, `idp_type`, `provision_status`.
search	object	Return results that match the given pattern (case-insensitive). Accepts an object, where the keys are property names, and the values are strings to search for. If no additional options are given, this will perform a `LIKE "%…%"` search. Supported properties: `name`, `description`. User directory of type SAML will have an empty value for both `name` and `description` properties. Both properties can be changed with the `userdirectory.update` operation.
countOutput	boolean	These parameters are described in the reference commentary.
excludeSearch	boolean	
limit	integer	
output	query	
preservekeys	boolean	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	

Parameter	Type	Description
startSearch	boolean	

Return values

(integer/array) Returns either:

- an array of objects
- the count of retrieved objects, if the countOutput parameter has been used

Examples

Retrieving user directories

Retrieve all user directories with additional properties that display media type mappings and provisioning groups mappings associated with each user directory.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "userdirectory.get",
    "params": {
        "output": "extend",
        "selectProvisionMedia": "extend",
        "selectProvisionGroups": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
       {
            "userdirectoryid": "1",
            "idp_type": "2",
            "name": "",
            "provision_status": "1",
            "description": "",
            "group_name": "groups",
            "user_username": "",
            "user_lastname": "",
            "idp_entityid": "http://example.com/simplesaml/saml2/idp/metadata.php",
            "sso_url": "http://example.com/simplesaml/saml2/idp/SSOService.php",
            "slo_url": "",
            "username_attribute": "uid",
            "sp_entityid": "zabbix",
            "nameid_format": "",
            "sign_messages": "0",
            "sign_assertions": "0",
            "sign_authn_requests": "0",
            "sign_logout_requests": "0",
            "sign_logout_responses": "0",
            "encrypt_nameid": "0",
            "encrypt_assertions": "0",
            "scim_status": "1",
            "provision_media": [
                {
                    "userdirectory_mediaid": "1",
                    "name": "example.com",
                    "mediatypeid": "1",
                    "attribute": "user@example.com",
                    "active": "0",
                    "severity": "63",
```

```
                "period": "1-7,00:00-24:00"
            }
        ],
        "provision_groups": [
            {
                "name": "*",
                "roleid": "1",
                "user_groups": [
                    {
                        "usrgrpid": "13"
                    }
                ]
            }
        ]
    },
    {
        "userdirectoryid": "2",
        "idp_type": "1",
        "name": "AD server",
        "provision_status": "1",
        "description": "",
        "host": "host.example.com",
        "port": "389",
        "base_dn": "DC=zbx,DC=local",
        "search_attribute": "sAMAccountName",
        "bind_dn": "CN=Admin,OU=Users,OU=Zabbix,DC=zbx,DC=local",
        "start_tls": "0",
        "search_filter": "",
        "group_basedn": "OU=Zabbix,DC=zbx,DC=local",
        "group_name": "CN",
        "group_member": "member",
        "group_filter": "(%{groupattr}=CN=%{ref},OU=Users,OU=Zabbix,DC=zbx,DC=local)",
        "group_membership": "",
        "user_username": "givenName",
        "user_lastname": "sn",
        "user_ref_attr": "CN",
        "provision_media": [
            {
                "userdirectory_mediaid": "2",
                "name": "example.com",
                "mediatypeid": "1",
                "attribute": "user@example.com",
                "active": "0",
                "severity": "63",
                "period": "1-7,00:00-24:00"
            }
        ],
        "provision_groups": [
            {
                "name": "*",
                "roleid": "4",
                "user_groups": [
                    {
                        "usrgrpid": "8"
                    }
                ]
            },
            {
                "name": "Zabbix administrators",
                "roleid": "2",
                "user_groups": [
                    {
```

```
                            "usrgrpid": "7"
                        },
                        {
                            "usrgrpid": "8"
                        }
                    ]
                }
            ]
        },
        {
            "userdirectoryid": "3",
            "idp_type": "1",
            "name": "LDAP API server #1",
            "provision_status": "0",
            "description": "",
            "host": "ldap://local.ldap",
            "port": "389",
            "base_dn": "ou=Users,dc=example,dc=org",
            "search_attribute": "uid",
            "bind_dn": "cn=ldap_search,dc=example,dc=org",
            "start_tls": "1",
            "search_filter": "",
            "group_basedn": "",
            "group_name": "",
            "group_member": "",
            "group_filter": "",
            "group_membership": "",
            "user_username": "",
            "user_lastname": "",
            "user_ref_attr": "",
            "provision_media": [],
            "provision_groups": []
        }
    ],
    "id": 1
}
```


See also

- User group

Source

CUserDirectory::get() in *ui/include/classes/api/services/CUserDirectory.php*.

userdirectory.test

Description

`object userdirectory.test(array userDirectory)`

This method allows you to test user directory connection settings.

Note:
This method also allows you to test what configured data matches the user directory settings for user provisioning (e.g., what user role, user groups, user medias will be assigned to the user). For this type of test the API request should be made for a user directory that has `provision_status` set to enabled.

Note:
This method is only available to *Super admin* user type.

Parameters

`(object)` User directory properties.

Since `userdirectory.get` API does not return `bind_password` field, `userdirectoryid` and/or `bind_password` should be supplied.
 Additionally to the standard user directory properties, the method accepts the following parameters.

Parameter	Type	Description
test_username	string	Username to test in user directory.
test_password	string	Username associated password to test in user directory.

Return values

`(bool)` Returns true on success.

Examples

Test user directory for existing user

Test user directory "3" for "user1".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "userdirectory.test",
    "params": {
        "userdirectoryid": "3",
        "host": "127.0.0.1",
        "port": "389",
        "base_dn": "ou=Users,dc=example,dc=org",
        "search_attribute": "uid",
        "bind_dn": "cn=ldap_search,dc=example,dc=org",
        "bind_password": "password",
        "test_username": "user1",
        "test_password": "password"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": true,
    "id": 1
}
```

Test user directory for non-existing user

Test user directory "3" for non-existing "user2".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "userdirectory.test",
    "params": {
        "userdirectoryid": "3",
        "host": "127.0.0.1",
        "port": "389",
        "base_dn": "ou=Users,dc=example,dc=org",
        "search_attribute": "uid",
        "bind_dn": "cn=ldap_search,dc=example,dc=org",
        "test_username": "user2",
        "test_password": "password"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "error": {
        "code": -32500,
        "message": "Application error.",
        "data": "Incorrect user name or password or account is temporarily blocked."
    },
    "id": 1
}
```

Test user directory for user provisioning

Test userdirectory "3" for what configured data matches the user directory settings for "user3" provisioning (e.g., what user role, user groups, user medias will be assigned to the user).

Request:

```
{
    "jsonrpc": "2.0",
    "method": "userdirectory.test",
    "params": {
        "userdirectoryid": "2",
        "host": "host.example.com",
        "port": "389",
        "base_dn": "DC=zbx,DC=local",
        "search_attribute": "sAMAccountName",
        "bind_dn": "CN=Admin,OU=Users,OU=Zabbix,DC=zbx,DC=local",
        "test_username": "user3",
        "test_password": "password"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "username": "user3",
        "name": "John",
        "surname": "Doe",
        "medias": [],
        "usrgrps": [
            {
                "usrgrpid": "8"
            },
            {
                "usrgrpid": "7"
            }
        ],
        "roleid": "2",
        "userdirectoryid": "2"
    },
    "id": 1
}
```

Source

CUserDirectory::test() in *ui/include/classes/api/services/CUserDirectory.php*.

userdirectory.update

Description

`object userdirectory.update(object/array userDirectory)`

This method allows you to update existing user directories.

> **Note:**
> This method is only available to *Super admin* user type.

Parameters

`(object/array)` User directory properties to be updated.

The `userdirectoryid` property must be defined for each user directory, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

`(object)` Returns an object containing the IDs of the updated user directories under the `userdirectoryids` property.

Examples

Update bind password for user directory

Set new bind password for a user directory.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "userdirectory.update",
    "params": {
        "userdirectoryid": "3",
        "bind_password": "newldappassword"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userdirectoryids": [
            "3"
        ]
    },
    "id": 1
}
```

Update mappings for user directory

Update provisioning groups mappings and media type mappings for user directory "2".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "userdirectory.update",
    "params": {
        "userdirectoryid": "2",
        "provision_media": [
                {
                    "userdirectory_mediaid": "2"
                }
            ],
            "provision_groups": [
                {
                    "name": "Zabbix administrators",
                    "roleid": "2",
                    "user_groups": [
                        {
                            "usrgrpid": "7"
                        },
```

```
                        {
                            "usrgrpid": "8"
                        },
                        {
                            "usrgrpid": "11"
                        }
                    ]
                }
            ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "userdirectoryids": [
            "2"
        ]
    },
    "id": 1
}
```

Source

CUserDirectory::update() in *ui/include/classes/api/services/CUserDirectory.php*.

User group

This class is designed to work with user groups.

Object references:

- User group
- Permission
- Tag-based permission
- APM-based permission

Available methods:

- usergroup.create - create new user groups
- usergroup.delete - delete user groups
- usergroup.get - retrieve user groups
- usergroup.update - update user groups

User group object

The following objects are directly related to the `usergroup` API.

User group

The user group object has the following properties.

Property	Type	Description
usrgrpid	ID	ID of the user group. Property behavior: - *read-only* - *required* for update operations
name	string	Name of the user group. Property behavior: - *required* for create operations

Property	Type	Description
debug_mode	integer	Whether debug mode is enabled or disabled. Possible values: 0 - *(default)* disabled; 1 - enabled.
gui_access	integer	Frontend authentication method of the users in the group. Possible values: 0 - *(default)* use the system default authentication method; 1 - use internal authentication; 2 - use LDAP authentication; 3 - disable access to the frontend.
mfa_status	integer	Whether MFA is enabled or disabled for the users in the group. Possible values: 0 - disabled (for all configured MFA methods); 1 - enabled (for all configured MFA methods).
mfaid	ID	MFA method used for the users in the group. Property behavior: - *supported* for *Super admin* type users - *supported* if `mfa_status` of Authentication object is set to ”Enabled”
proxy_mode	integer	Mode for handling access to proxies that are not included in proxy groups. Possible values: 0 - Deny list; 1 - *(default)* Allow list. Property behavior: - *supported* for *Super admin* type users.
proxy_group_mode	integer	Mode for handling access to proxy groups (and their proxies). Possible values: 0 - Deny list; 1 - *(default)* Allow list. Property behavior: - *supported* for *Super admin* type users.
users_status	integer	Whether the user group is enabled or disabled. For deprovisioned users, the user group cannot be enabled. Possible values: 0 - *(default)* enabled; 1 - disabled.
userdirectoryid	ID	ID of the user directory used for authentication. Property behavior: - *supported* for *Super admin* type users - *supported* if `gui_access` is set to ”use the system default authentication method” or ”use LDAP authentication”

Permission

The permission object has the following properties.

Property	Type	Description
id	ID	ID of the host group or template group to add permission to. Property behavior: - *required*
permission	integer	Access level to the host group or template group. Possible values: 0 - access denied; 2 - read-only access; 3 - read-write access. Property behavior: - *required* for create operations

Tag-based permission

The tag-based permission object has the following properties.

Property	Type	Description
groupid	ID	ID of the host group to add permission to. Property behavior: - *required*
tag	string	Tag name.
value	string	Tag value.

APM-based permission

The APM-based permission object has the following properties.

Property	Type	Description
key	string	APM resource attribute key. Property behavior: - *required*
value	string	APM resource attribute value.
signal_type	integer	APM resource attribute category. Possible values: -1 - *(default)* All; 0 - Traces; 1 - Metrics; 2 - Logs.

usergroup.create

Description

`object usergroup.create(object/array userGroups)`

This method allows you to create new user groups.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` User groups to create.

Additionally to the standard user group properties, the method accepts the following parameters.

Parameter	Type	Description
hostgroup_rights	object/array	Host group permissions to assign to the user group.
proxies	object/array	Proxies to be linked to the user group. The proxies must have only the `proxyid` property defined.
proxy_groups	object/array	Proxy groups to be linked to the user group. The proxy groups must have only the `proxy_groupid` property defined.
templategroup_rights	object/array	Template group permissions to assign to the user group.
tag_filters	array	Tag-based permissions to assign to the user group.
apm_filters	object/array	APM-based permissions to assign to the user group.
users	object/array	Users to add to the user group. The user must have only the `userid` property defined.

Return values

`(object)` Returns an object containing the IDs of the created user groups under the `usrgrpids` property. The order of the returned IDs matches the order of the passed user groups.

Examples

Creating a user group

Create a user group *Operation managers* with denied access to host group "2", and add a user to it.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usergroup.create",
    "params": {
        "name": "Operation managers",
        "hostgroup_rights": {
            "id": "2",
            "permission": 0
        },
        "users": [
            {
                "userid": "12"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "usrgrpids": [
            "20"
        ]
    },
    "id": 1
}
```

See also

- Permission

Source

CUserGroup::create() in *ui/include/classes/api/services/CUserGroup.php*.

usergroup.delete

Description

`object usergroup.delete(array userGroupIds)`

This method allows you to delete user groups.

> **Attention:**
> Deprovisioned users group (the user group specified for `disabled_usrgrpid` in Authentication) cannot be deleted.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the user groups to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted user groups under the `usrgrpids` property.

Examples

Deleting multiple user groups

Delete two user groups.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usergroup.delete",
    "params": [
        "20",
        "21"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "usrgrpids": [
            "20",
            "21"
        ]
    },
    "id": 1
}
```

Source

CUserGroup::delete() in *ui/include/classes/api/services/CUserGroup.php*.

usergroup.get

Description

`integer/array usergroup.get(object parameters)`

The method allows you to retrieve user groups according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
mfaids	ID/array	Return only user groups with the given MFA methods. Parameter behavior: - *supported* for *Super admin* type users
mfa_status	integer	Return only user groups with the given MFA status. Refer to the user group page for a list of supported statuses.
proxyids	ID/array	Return only user groups that contain the given proxies. Parameter behavior: - *supported* for *Super admin* type users.
proxy_groupids	ID/array	Return only user groups that contain the given proxy groups. Parameter behavior: - *supported* for *Super admin* type users.
selectProxies	ID/array	Return a `proxies` property with user group proxies. Parameter behavior: - *supported* for *Super admin* type users.
selectProxyGroups	ID/array	Return a `proxy_groups` property with user group proxy groups. Parameter behavior: - *supported* for *Super admin* type users.
status	integer	Return only user groups with the given status. Refer to the user group page for a list of supported statuses.
userids	ID/array	Return only user groups that contain the given users.
usrgrpids	ID/array	Return only user groups with the given IDs.
selectTagFilters	query	Return user group tag-based permissions in the `tag_filters` property.
selectApmFilters	query	Return user group APM-based permissions in the `apm_filters` property. Supported properties: `key`, `value`, `signal_type`. Default: `extend`.
selectUsers	query	Return the users from the user group in the `users` property. See `user.get` for restrictions based on user type.
selectHostGroupRights	query	Return user group host group permissions in the `hostgroup_rights` property. Refer to the user group page for a list of access levels to host groups.
selectTemplateGroupRights	query	Return user group template group permissions in the `templategroup_rights` property. Refer to the user group page for a list of access levels to template groups.
limitSelects	integer	Limits the number of records returned by subselects.

Parameter	Type	Description
output	query	User group object properties to be returned. *Admin* and *User* type users may retrieve only the following properties: `usrgrpid`, `name`, `gui_access`, `users_status`, `debug_mode`, `mfa_status`. Default: `extend`.
sortfield	string/array	Sort the result by the given properties. Possible values: `usrgrpid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving enabled user groups

Retrieve all enabled user groups.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usergroup.get",
    "params": {
        "output": "extend",
        "status": 0
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "usrgrpid": "7",
            "name": "Zabbix administrators",
            "gui_access": "0",
            "users_status": "0",
            "debug_mode": "1",
            "userdirectoryid": "0",
            "mfa_status": "0",
            "mfaid": "0"
        },
        {
            "usrgrpid": "8",
            "name": "Guests",
```

```
            "gui_access": "0",
            "users_status": "0",
            "debug_mode": "0",
            "userdirectoryid": "0",
            "mfa_status": "0",
            "mfaid": "0"
        },
        {
            "usrgrpid": "11",
            "name": "Enabled debug mode",
            "gui_access": "0",
            "users_status": "0",
            "debug_mode": "1",
            "userdirectoryid": "0",
            "mfa_status": "0",
            "mfaid": "0"
        },
        {
            "usrgrpid": "12",
            "name": "No access to the frontend",
            "gui_access": "2",
            "users_status": "0",
            "debug_mode": "0",
            "userdirectoryid": "0",
            "mfa_status": "0",
            "mfaid": "0"
        },
        {
            "usrgrpid": "14",
            "name": "Read only",
            "gui_access": "0",
            "users_status": "0",
            "debug_mode": "0",
            "userdirectoryid": "0",
            "mfa_status": "0",
            "mfaid": "0"
        },
        {
            "usrgrpid": "18",
            "name": "Deny",
            "gui_access": "0",
            "users_status": "0",
            "debug_mode": "0",
            "userdirectoryid": "0",
            "mfa_status": "0",
            "mfaid": "0"
        }
    ],
    "id": 1
}
```


See also

- User

Source

CUserGroup::get() in *ui/include/classes/api/services/CUserGroup.php*.

usergroup.update

Description

`object usergroup.update(object/array userGroups)`

This method allows you to update existing user groups.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` User group properties to be updated.

The `usrgrpid` property must be defined for each user group, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard user group properties, the method accepts the following parameters.

Parameter	Type	Description
hostgroup_rights	object/array	Host group permissions to replace the current permissions assigned to the user group.
proxies	object/array	Proxies to be linked to the user group. The proxies must have only the `proxyid` property defined.
proxy_groups	object/array	Proxy groups to be linked to the user group. The proxy groups must have only the `proxy_groupid` property defined.
templategroup_rights	object/array	Template group permissions to replace the current permissions assigned to the user group.
tag_filters	array	Tag-based permissions to replace the current permissions assigned to the user group.
apm_filters	object/array	APM-based permissions to replace the current permissions assigned to the user group.
users	object/array	Users to replace the current users assigned to the user group. The user must have only the `userid` property defined.

Return values

`(object)` Returns an object containing the IDs of the updated user groups under the `usrgrpids` property.

Examples

Enabling a user group and updating permissions

Enable a user group and provide read-write access for it to host groups "2" and "4".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usergroup.update",
    "params": {
        "usrgrpid": "17",
        "users_status": "0",
        "hostgroup_rights": [
            {
                "id": "2",
                "permission": 3
            },
            {
                "id": "4",
                "permission": 3
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "usrgrpids": [
            "17"
        ]
    },
    "id": 1
}
```

See also

- Permission

Source

CUserGroup::update() in *ui/include/classes/api/services/CUserGroup.php*.

User macro

This class is designed to work with host-level and global user macros.

Object references:

- Global macro
- Host macro
    - Macro configuration

Available methods:

- usermacro.create - create new host macros
- usermacro.createglobal - create new global macros
- usermacro.delete - delete host macros
- usermacro.deleteglobal - delete global macros
- usermacro.get - retrieve host and global macros
- usermacro.update - update host macros
- usermacro.updateglobal - update global macros

User macro object

The following objects are directly related to the `usermacro` API.

Global macro

The global macro object has the following properties.

Property	Type	Description
globalmacroid	ID	ID of the global macro. Property behavior: - *read-only* - *required* for update operations
description	string	Description of the macro.
macro	string	Macro string. Property behavior: - *required* for create operations
type	integer	Type of macro. Possible values: 0 - *(default)* Text macro; 1 - Secret macro; 2 - Vault secret.

Property	Type	Description
value	string	Value of the macro. Property behavior: - *write-only* if `type` is set to "Secret macro" - *required* for create operations

Host macro

The host macro object defines a macro available on a host, host prototype or template. It has the following properties.

Property	Type	Description
hostmacroid	ID	ID of the macro. Property behavior: - *read-only* - *required* for update operations
automatic	integer	Defines whether the macro is controlled by discovery rule. Possible values: 0 - *(default)* Macro is managed by user; 1 - Macro is managed by discovery rule. User is not allowed to create automatic macro. To update automatic macro, it must be converted to manual.
config	object/array	Macro configuration responsible for how the macro is displayed in the Host Wizard.
description	string	Description of the macro.
hostid	ID	ID of the host, host prototype, or template that the macro belongs to. Property behavior: - *constant* - *required* for create operations
macro	string	Macro string. Property behavior: - *required* for create operations
type	integer	Type of macro. Possible values: 0 - *(default)* Text macro; 1 - Secret macro; 2 - Vault secret.
value	string	Value of the macro. Property behavior: - *write-only* if `type` is set to "Secret macro" - *required* for create operations

Macro configuration

The macro configuration object defines how a macro is displayed in the Host Wizard.

Property	Type	Description
type	integer	Type of macro input field. Possible values: 0 - Macro is not used in Host Wizard; 1 - Textbox; 2 - List; 3 - Checkbox. Property behavior: - *required*
label	string	Label for macro input field. Property behavior: - *required* if `type` is set to ”Textbox”, ”List”, or ”Checkbox”
description	text	Help text displayed alongside macro input field. Supports Markdown formatting.
priority	integer	Macro input field position in the macro list. If `priority` is not set, the macro is added at the end of the ungrouped macro list.
required	integer	Marks the macro as mandatory. Possible values: 0 - Not mandatory; 1 - Mandatory. Property behavior: - *supported* if `type` is set to ”Textbox” or ”List”
regex	string	Regular expression to validate user input in a textbox field. Property behavior: - *supported* if `type` is set to ”Textbox”
section_name	string	Label of the collapsible section where the macro is grouped. If `section_name` is not set, the macro is ungrouped. Ungrouped macros are displayed first; grouped macros are displayed below, ordered by their `priority` within each collapsible section.
options	text	JSON string defining list items or checkbox values. For lists: an array of objects with `value` and `text` properties. Example: `[{"value": "http", "text": "HTTP"}, {"value": "https", "text": "HTTPS"}]` For checkboxes: an object with `checked` and `unchecked` properties. Example: `{"checked": true, "unchecked": false}` Property behavior: - *required* if `type` is set to ”List” or ”Checkbox”

usermacro.create

Description

`object usermacro.create(object/array hostMacros)`

This method allows you to create new host macros.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Host macros to create.

The method accepts host macros with the standard host macro properties.

Return values

`(object)` Returns an object containing the IDs of the created host macros under the `hostmacroids` property. The order of the returned IDs matches the order of the passed host macros.

Examples

Creating a host macro

Create a host macro `{$SNMP_COMMUNITY}` with the value "public" on host "10198".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usermacro.create",
    "params": {
        "hostid": "10198",
        "macro": "{$SNMP_COMMUNITY}",
        "value": "public"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostmacroids": [
            "11"
        ]
    },
    "id": 1
}
```

Source

CUserMacro::create() in *ui/include/classes/api/services/CUserMacro.php*.

usermacro.createglobal

Description

`object usermacro.createglobal(object/array globalMacros)`

This method allows you to create new global macros.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Global macros to create.

The method accepts global macros with the standard global macro properties.

Return values

`(object)` Returns an object containing the IDs of the created global macros under the `globalmacroids` property. The order of the returned IDs matches the order of the passed global macros.

Examples

Creating a global macro

Create a global macro {$SNMP_COMMUNITY} with value "public".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usermacro.createglobal",
    "params":  {
        "macro": "{$SNMP_COMMUNITY}",
        "value": "public"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "globalmacroids": [
            "6"
        ]
    },
    "id": 1
}
```

Source

CUserMacro::createGlobal() in *ui/include/classes/api/services/CUserMacro.php*.

usermacro.delete

Description

`object usermacro.delete(array hostMacroIds)`

This method allows you to delete host macros.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the host macros to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted host macros under the `hostmacroids` property.

Examples

Deleting multiple host macros

Delete two host macros.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usermacro.delete",
    "params": [
        "32",
        "11"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostmacroids": [
            "32",
            "11"
        ]
    },
    "id": 1
}
```

Source

CUserMacro::delete() in *ui/include/classes/api/services/CUserMacro.php*.

usermacro.deleteglobal

Description

`object usermacro.deleteglobal(array globalMacroIds)`

This method allows you to delete global macros.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the global macros to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted global macros under the `globalmacroids` property.

Examples

Deleting multiple global macros

Delete two global macros.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usermacro.deleteglobal",
    "params": [
        "32",
        "11"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "globalmacroids": [
            "32",
            "11"
        ]
    },
    "id": 1
}
```

Source

CUserMacro::deleteGlobal() in *ui/include/classes/api/services/CUserMacro.php*.

usermacro.get

Description

`integer/array usermacro.get(object parameters)`

The method allows you to retrieve host and global macros according to the given parameters.

> **Note:**
> This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
globalmacro	flag	Return global macros instead of host macros.
globalmacroids	ID/array	Return only global macros with the given IDs.
groupids	ID/array	Return only host macros that belong to hosts or templates from the given host groups or template groups.
hostids	ID/array	Return only macros that belong to the given hosts or templates.
hostmacroids	ID/array	Return only host macros with the given IDs.
inherited	boolean	If set to `true` return only host prototype user macros inherited from a template.
selectHostGroups	query	Return host groups that the host macro belongs to in the `hostgroups` property. Used only when retrieving host macros.
selectHosts	query	Return hosts that the host macro belongs to in the `hosts` property. Used only when retrieving host macros.
selectTemplateGroups	query	Return template groups that the template macro belongs to in the `templategroups` property. Used only when retrieving template macros.
selectTemplates	query	Return templates that the host macro belongs to in the `templates` property. Used only when retrieving host macros.
sortfield	string/array	Sort the result by the given properties. Possible values: `macro`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving host macros for a host

Retrieve all host macros defined for host ”10198”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usermacro.get",
    "params": {
        "output": "extend",
        "hostids": "10198"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "hostmacroid": "9",
            "hostid": "10198",
            "macro": "{$INTERFACE}",
            "value": "eth0",
            "description": "",
            "type": "0",
            "automatic": "0"
        },
        {
            "hostmacroid": "11",
            "hostid": "10198",
            "macro": "{$SNMP_COMMUNITY}",
            "value": "public",
            "description": "",
            "type": "0",
            "automatic": "0"
        }
    ],
    "id": 1
}
```

Retrieving host macros for a template

Retrieve host macros defined for template ”10265” that contain ”STATUS” in the macro name.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usermacro.get",
    "params": {
        "output": "extend",
        "hostids": "10265",
        "search": {
            "macro": "STATUS"
        }
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
```

```
    "result": [
        {
            "hostmacroid": "6709",
            "hostid": "10265",
            "macro": "{$APACHE.STATUS.HOST}",
            "value": "",
            "description": "The hostname or IP address of the Apache status page host.",
            "type": "0",
            "automatic": "0",
            "config": {
                "type": "1",
                "priority": "1",
                "section_name": "",
                "label": "Apache status host",
                "description": "Hostname or IP address of the Apache status page host.",
                "required": "1",
                "regex": "",
                "options": []
            }
        },
        {
            "hostmacroid": "814",
            "hostid": "10265",
            "macro": "{$APACHE.STATUS.PATH}",
            "value": "server-status?auto",
            "description": "The URL path of the Apache status page.",
            "type": "0",
            "automatic": "0",
            "config": {
                "type": "1",
                "priority": "3",
                "section_name": "",
                "label": "Apache status page path",
                "description": "URL path of the Apache status page.",
                "required": "1",
                "regex": "",
                "options": []
            }
        },
        {
            "hostmacroid": "815",
            "hostid": "10265",
            "macro": "{$APACHE.STATUS.PORT}",
            "value": "80",
            "description": "The port of the Apache status page.",
            "type": "0",
            "automatic": "0",
            "config": {
                "type": "1",
                "priority": "2",
                "section_name": "",
                "label": "Apache status page port",
                "description": "In the range from 1 to 65535 inclusive.",
                "required": "1",
                "regex": "^-?([0-9]+|(([0-9]+)\\.([0-9]+)))$",
                "options": []
            }
        },
        {
            "hostmacroid": "816",
            "hostid": "10265",
            "macro": "{$APACHE.STATUS.SCHEME}",
```

```
            "value": "http",
            "description": "The request scheme, which may be either HTTP or HTTPS.",
            "type": "0",
            "automatic": "0",
            "config": {
                "type": "2",
                "priority": "4",
                "section_name": "",
                "label": "Request scheme",
                "description": "Request scheme, which may be either HTTP or HTTPS.",
                "required": "0",
                "regex": "",
                "options": [
                    {
                        "value": "http",
                        "text": "HTTP"
                    },
                    {
                        "value": "https",
                        "text": "HTTPS"
                    }
                ]
            }
        }
    ],
    "id": 1
}
```


Retrieving global macros

Retrieve all global macros.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "usermacro.get",
    "params": {
        "output": "extend",
        "globalmacro": true
    },
    "id": 1
}
```


Response:


```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "globalmacroid": "6",
            "macro": "{$SNMP_COMMUNITY}",
            "value": "public",
            "description": "",
            "type": "0"
        }
    ],
    "id": 1
}
```


Source

CUserMacro::get() in *ui/include/classes/api/services/CUserMacro.php*.

usermacro.update

Description

`object usermacro.update(object/array hostMacros)`

This method allows you to update existing host macros.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Host macro properties to be updated.

The `hostmacroid` property must be defined for each host macro, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

`(object)` Returns an object containing the IDs of the updated host macros under the `hostmacroids` property.

Examples

Changing the value of a host macro

Change the value of a host macro to "public".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usermacro.update",
    "params": {
        "hostmacroid": "1",
        "value": "public"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostmacroids": [
            "1"
        ]
    },
    "id": 1
}
```

Change macro value that was created by discovery rule

Convert discovery rule created "automatic" macro to "manual" and change its value to "new-value".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usermacro.update",
    "params": {
        "hostmacroid": "1",
        "value": "new-value",
        "automatic": "0"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "hostmacroids": [
            "1"
        ]
    },
    "id": 1
}
```

Source

CUserMacro::update() in *ui/include/classes/api/services/CUserMacro.php*.

usermacro.updateglobal

Description

`object usermacro.updateglobal(object/array globalMacros)`

This method allows you to update existing global macros.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Global macro properties to be updated.

The `globalmacroid` property must be defined for each global macro, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

`(object)` Returns an object containing the IDs of the updated global macros under the `globalmacroids` property.

Examples

Changing the value of a global macro

Change the value of a global macro to "public".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "usermacro.updateglobal",
    "params": {
        "globalmacroid": "1",
        "value": "public"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "globalmacroids": [
            "1"
        ]
    },
    "id": 1
}
```

Source

CUserMacro::updateGlobal() in *ui/include/classes/api/services/CUserMacro.php*.

Value map

This class is designed to work with value maps.

Object references:

- Value map
  - Value mappings

Available methods:

- valuemap.create - create new value maps
- valuemap.delete - delete value maps
- valuemap.get - retrieve value maps
- valuemap.update - update value maps

Value map object

The following objects are directly related to the `valuemap` API.

Value map

The value map object has the following properties.

Property	Type	Description
valuemapid	ID	ID of the value map. Property behavior: - *read-only* - *required* for update operations
hostid	ID	ID of the host or template that the value map belongs to. Property behavior: - *constant* - *required* for create operations
name	string	Name of the value map. Property behavior: - *required* for create operations
mappings	array	Value mappings for current value map. The mapping object is described in detail below. Property behavior: - *required* for create operations
uuid	string	Universal unique identifier, used for linking imported value maps to already existing ones. Auto-generated, if not given. Property behavior: - *supported* if the value map belongs to a template

Value mappings

The value mappings object defines value mappings of the value map. It has the following properties.

Property	Type	Description
type	integer	Mapping match type. Possible values: 0 - *(default)* mapping will be applied if value is equal; 1 - mapping will be applied if value is greater or equal[1]; 2 - mapping will be applied if value is less or equal[1]; 3 - mapping will be applied if value is in range (ranges are inclusive; multiple ranges, separated by comma character, can be defined)[1]; 4 - mapping will be applied if value matches a regular expression[2]; 5 - if no matches are found, mapping will not be applied, and the default value will be used. If `type` is set to ”0”, ”1”, ”2”, ”3”, ”4”, then `value` cannot be empty. If `type` is set to ”5”, then `value` must be empty.
value	string	Original value. Property behavior: - *required* if `type` is set to ”1”, ”2”, ”3”, ”4” - *supported* if `type` is set to ”5”
newvalue	string	Value to which the original value is mapped to. Property behavior: - *required*

[1] supported only for items having value type ”numeric unsigned”, ”numeric float”.
 [2] supported only for items having value type ”character”.

valuemap.create

Description

`object valuemap.create(object/array valuemaps)`

This method allows you to create new value maps.

Note:
This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Value maps to create.

The method accepts value maps with the standard value map properties.

Return values

`(object)` Returns an object containing the IDs of the created value maps the `valuemapids` property. The order of the returned IDs matches the order of the passed value maps.

Examples

Creating a value map

Create one value map with two mappings.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "valuemap.create",
    "params": {
        "hostid": "50009",
        "name": "Service state",
```

```
        "mappings": [
            {
                "type": "1",
                "value": "1",
                "newvalue": "Up"
            },
            {
                "type": "5",
                "newvalue": "Down"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "valuemapids": [
            "1"
        ]
    },
    "id": 1
}
```

Source

CValueMap::create() in *ui/include/classes/api/services/CValueMap.php*.

valuemap.delete

Description

`object valuemap.delete(array valuemapids)`

This method allows you to delete value maps.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the value maps to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted value maps under the `valuemapids` property.

Examples

Deleting multiple value maps

Delete two value maps.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "valuemap.delete",
    "params": [
        "1",
        "2"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "valuemapids": [
            "1",
            "2"
        ]
    },
    "id": 1
}
```

Source

CValueMap::delete() in *ui/include/classes/api/services/CValueMap.php*.

valuemap.get

Description

`integer/array valuemap.get(object parameters)`

The method allows you to retrieve value maps according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
valuemapids	ID/array	Return only value maps with the given IDs.
selectMappings	query	Return the value mappings for current value map in the `mappings` property. Supports `count`.
sortfield	string/array	Sort the result by the given properties. Possible values: `valuemapid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving value maps

Retrieve all configured value maps.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "valuemap.get",
    "params": {
        "output": "extend"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "valuemapid": "4",
            "name": "APC Battery Replacement Status"
        },
        {
            "valuemapid": "5",
            "name": "APC Battery Status"
        },
        {
            "valuemapid": "7",
            "name": "Dell Open Manage System Status"
        }
    ],
    "id": 1
}
```

Retrieve one value map with its mappings.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "valuemap.get",
    "params": {
        "output": "extend",
        "selectMappings": "extend",
        "valuemapids": ["4"]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "valuemapid": "4",
            "name": "APC Battery Replacement Status",
            "mappings": [
                {
                    "type": "0",
                    "value": "1",
                    "newvalue": "unknown"
                },
                {
```

```
                    "type": "0",
                    "value": "2",
                    "newvalue": "notInstalled"
                },
                {
                    "type": "0",
                    "value": "3",
                    "newvalue": "ok"
                },
                {
                    "type": "0",
                    "value": "4",
                    "newvalue": "failed"
                },
                {
                    "type": "0",
                    "value": "5",
                    "newvalue": "highTemperature"
                },
                {
                    "type": "0",
                    "value": "6",
                    "newvalue": "replaceImmediately"
                },
                {
                    "type": "0",
                    "value": "7",
                    "newvalue": "lowCapacity"
                }
            ]
        }
    ],
    "id": 1
}
```


Source

CValueMap::get() in *ui/include/classes/api/services/CValueMap.php*.

valuemap.update

Description

`object valuemap.update(object/array valuemaps)`

This method allows you to update existing value maps.

> **Note:**
> This method is only available to *Super admin* user type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Value map properties to be updated.

The `valuemapid` property must be defined for each value map, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Return values

`(object)` Returns an object containing the IDs of the updated value maps under the `valuemapids` property.

Examples

Changing value map name

Change value map name to ”Device status”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "valuemap.update",
    "params": {
        "valuemapid": "2",
        "name": "Device status"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "valuemapids": [
            "2"
        ]
    },
    "id": 1
}
```

Changing mappings for one value map.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "valuemap.update",
    "params": {
        "valuemapid": "2",
        "mappings": [
            {
                "type": "0",
                "value": "0",
                "newvalue": "Online"
            },
            {
                "type": "0",
                "value": "1",
                "newvalue": "Offline"
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "valuemapids": [
            "2"
        ]
    },
    "id": 1
}
```

Source

CValueMap::update() in *ui/include/classes/api/services/CValueMap.php*.

Web scenario

This class is designed to work with web scenarios.

Object references:

- Web scenario
- Scenario step
  - HTTP field
- Web scenario tag

Available methods:

- httptest.create - create new web scenarios
- httptest.delete - delete web scenarios
- httptest.get - retrieve web scenarios
- httptest.update - update web scenarios

Web scenario object

The following objects are directly related to the `webcheck` API.

Web scenario

The web scenario object has the following properties.

Property	Type	Description
httptestid	ID	ID of the web scenario. Property behavior: - *read-only* - *required* for update operations
hostid	ID	ID of the host that the web scenario belongs to. Property behavior: - *constant* - *required* for create operations
name	string	Name of the web scenario. Property behavior: - *required* for create operations
agent	string	User agent string that will be used by the web scenario. Default: Zabbix
authentication	integer	Authentication method that will be used by the web scenario. Possible values: 0 - *(default)* none; 1 - basic HTTP authentication; 2 - NTLM authentication; 3 - Kerberos authentication; 4 - Digest authentication.
delay	string	Execution interval of the web scenario. Accepts a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., 30s, 1m, 2h, 1d), or a user macro. Default: 1m.
headers	array	HTTP headers that will be sent when performing a request.
http_password	string	Password used for basic HTTP, NTLM, Kerberos, or Digest authentication.
http_proxy	string	Proxy that will be used by the web scenario given as *http://[username[:password]@]proxy.example.com[:port]*.
http_user	string	User name used for basic HTTP, NTLM, Kerberos, or Digest authentication.

Property	Type	Description
retries	integer	Number of times a web scenario will try to execute each step before failing. Default: 1.
ssl_cert_file	string	Name of the SSL certificate file used for client authentication (must be in PEM format).
ssl_key_file	string	Name of the SSL private key file used for client authentication (must be in PEM format).
ssl_key_password	string	SSL private key password.
status	integer	Whether the web scenario is enabled. Possible values: 0 - *(default)* enabled; 1 - disabled.
templateid	ID	ID of the parent template web scenario. Property behavior: - *read-only*
variables	array	Web scenario variables.
verify_host	integer	Whether to validate that the host name for the connection matches the one in the host's certificate. Possible values: 0 - *(default)* skip host verification; 1 - verify host.
verify_peer	integer	Whether to validate that the host's certificate is authentic. Possible values: 0 - *(default)* skip peer verification; 1 - verify peer.
uuid	string	Global unique identifier, used for linking imported web scenarios to already existing ones. Auto-generated, if not given. Property behavior: - *supported* if the web scenario belongs to a template

Scenario step

The scenario step object defines a specific web scenario check. It has the following properties.

Property	Type	Description
name	string	Name of the scenario step. Property behavior: - *required*
no	integer	Sequence number of the step in a web scenario. Property behavior: - *required*
url	string	URL to be checked. Property behavior: - *required*
follow_redirects	integer	Whether to follow HTTP redirects. Possible values: 0 - don't follow redirects; 1 - *(default)* follow redirects.
headers	array	HTTP headers that will be sent when performing a request. Scenario step headers will overwrite headers specified for the web scenario.

Property	Type	Description
posts	string/array	HTTP POST variables as a string (raw post data) or as an array of HTTP fields (form field data).
required	string	Text that must be present in the response.
retrieve_mode	integer	Part of the HTTP response that the scenario step must retrieve. Possible values: 0 - *(default)* only body; 1 - only headers; 2 - headers and body.
status_codes	string	Ranges of required HTTP status codes, separated by commas.
timeout	string	Request timeout in seconds. Accepts a number of seconds (e.g., 1, 60, 300), a value with a time suffix (e.g., 30s, 1m, 1h), or a user macro. Default: 15s. Maximum: 1h. Minimum: 1s.
variables	array	Scenario step variables.
query_fields	array	Query fields - array of HTTP fields that will be added to URL when performing a request.

HTTP field

The HTTP field object defines the name and value that is used to specify the web scenario variables, HTTP headers, and POST fields or query fields. It has the following properties.

Property	Type	Description
name	string	Name of header/variable/POST or GET field. Property behavior: - *required*
value	string	Value of header/variable/POST or GET field. Property behavior: - *required*

Web scenario tag

The web scenario tag object has the following properties.

Property	Type	Description
tag	string	Web scenario tag name. Property behavior: - *required*
value	string	Web scenario tag value. Property behavior: - *read-only*.
object	integer	Type of object the tag was inherited from. Possible values: 0 - template; 1 - host. Property behavior: - *read-only*.
objectid	ID	ID of the object the tag was inherited from. Property behavior: - *read-only*.

httptest.create

Description

`object httptest.create(object/array webScenarios)`

This method allows you to create new web scenarios.

> **Note:**
> Creating a web scenario will automatically create a set of web monitoring items.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Web scenarios to create.

Additionally to the standard web scenario properties, the method accepts the following parameters.

Parameter	Type	Description
steps	array	Scenario steps. Parameter behavior: - *required*
tags	array	Web scenario tags.

Return values

`(object)` Returns an object containing the IDs of the created web scenarios under the `httptestids` property. The order of the returned IDs matches the order of the passed web scenarios.

Examples

Creating a web scenario

Create a web scenario to monitor the company home page. The scenario will have two steps, to check the home page and the "About" page and make sure they return the HTTP status code 200.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "httptest.create",
    "params": {
        "name": "Homepage check",
        "hostid": "10085",
        "steps": [
            {
                "name": "Homepage",
                "url": "http://example.com",
                "status_codes": "200",
                "no": 1
            },
            {
                "name": "Homepage / About",
                "url": "http://example.com/about",
                "status_codes": "200",
                "no": 2
            }
        ]
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "httptestids": [
            "5"
        ]
    },
    "id": 1
}
```

See also

- Scenario step

Source

CHttpTest::create() in *ui/include/classes/api/services/CHttpTest.php*.

httptest.delete

Description

`object httptest.delete(array webScenarioIds)`

This method allows you to delete web scenarios.

Note:
This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(array)` IDs of the web scenarios to delete.

Return values

`(object)` Returns an object containing the IDs of the deleted web scenarios under the `httptestids` property.

Examples

Deleting multiple web scenarios

Delete two web scenarios.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "httptest.delete",
    "params": [
        "2",
        "3"
    ],
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "httptestids": [
            "2",
            "3"
        ]
    },
    "id": 1
}
```

Source

CHttpTest::delete() in *ui/include/classes/api/services/CHttpTest.php*.

httptest.get

Description

`integer/array httptest.get(object parameters)`

The method allows you to retrieve web scenarios according to the given parameters.

Note:
This method is available to users of any type. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object)` Parameters defining the desired output.

The method supports the following parameters.

Parameter	Type	Description
groupids	ID/array	Return only web scenarios that belong to the given host groups.
hostids	ID/array	Return only web scenarios that belong to the given hosts.
httptestids	ID/array	Return only web scenarios with the given IDs.
inherited	boolean	If set to `true` return only web scenarios inherited from a template.
inheritedTags	boolean	Return web scenarios that have given `tags` also in template/host/linked templates. Possible values: `true` - template/host/linked templates must also have given tags; `false` - *(default)* tags from template/host/linked templates are ignored.
monitored	boolean	If set to `true` return only enabled web scenarios that belong to monitored hosts.
selectInheritedTags	query	Return an `inheritedTags` property with tags that are on template/host/linked templates.
templated	boolean	If set to `true` return only web scenarios that belong to templates.
templateids	ID/array	Return only web scenarios that belong to the given templates.
expandName	flag	Expand macros in the name of the web scenario.
expandStepName	flag	Expand macros in the names of scenario steps.
evaltype	integer	Tag evaluation method. Possible values: 0 - *(default)* And/Or; 2 - Or.
tags	array	Return only web scenarios with the given tags. Format: `[{"tag": "<tag>", "value": "<value>", "operator": "<operator>"}, ...]`. An empty array returns all web scenarios. Possible operator values: 0 - *(default)* Contains; 1 - Equals; 2 - Does not contain; 3 - Does not equal; 4 - Exists; 5 - Does not exist.
selectHosts	query	Return the hosts that the web scenario belongs to as an array in the `hosts` property.
selectSteps	query	Return web scenario steps in the `steps` property. Supports `count`.

Parameter	Type	Description
selectTags	query	Return web scenario tags in the `tags` property.
sortfield	string/array	Sort the result by the given properties. Possible values: `httptestid`, `name`.
countOutput	boolean	These parameters are described in the reference commentary.
editable	boolean	
excludeSearch	boolean	
filter	object	
limit	integer	
output	query	
preservekeys	boolean	
search	object	
searchByAny	boolean	
searchWildcardsEnabled	boolean	
sortorder	string/array	
startSearch	boolean	

Return values

`(integer/array)` Returns either:

- an array of objects;
- the count of retrieved objects, if the `countOutput` parameter has been used.

Examples

Retrieving a web scenario

Retrieve all data about web scenario "9".

Request:

```
{
    "jsonrpc": "2.0",
    "method": "httptest.get",
    "params": {
        "output": "extend",
        "selectSteps": "extend",
        "httptestids": "9"
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
            "httptestid": "9",
            "name": "Homepage check",
            "delay": "1m",
            "status": "0",
            "agent": "Zabbix",
            "authentication": "0",
            "http_user": "",
            "http_password": "",
            "hostid": "10084",
            "templateid": "0",
            "http_proxy": "",
            "retries": "1",
            "ssl_cert_file": "",
            "ssl_key_file": "",
            "ssl_key_password": "",
            "verify_peer": "0",
```

```
            "verify_host": "0",
            "uuid": "",
            "headers": [],
            "variables": [],
            "steps": [
                {
                    "httpstepid": "36",
                    "httptestid": "9",
                    "name": "Homepage",
                    "no": "1",
                    "url": "http://example.com",
                    "timeout": "15s",
                    "posts": "",
                    "required": "",
                    "status_codes": "200",
                    "follow_redirects": "1",
                    "retrieve_mode": "0",
                    "post_type": "1",
                    "headers": [],
                    "variables": [
                        {
                            "name":"{var}",
                            "value":"12"
                        }
                    ],
                    "query_fields": []
                },
                {
                    "httpstepid": "37",
                    "httptestid": "9",
                    "name": "Homepage / About",
                    "no": "2",
                    "url": "http://example.com/about",
                    "timeout": "15s",
                    "posts": "",
                    "required": "",
                    "status_codes": "200",
                    "follow_redirects": "1",
                    "retrieve_mode": "0",
                    "post_type": "1",
                    "headers": [],
                    "variables": [],
                    "query_fields": []
                }
            ]
        }
    ],
    "id": 1
}
```


See also

- Host
- Scenario step

Source

CHttpTest::get() in *ui/include/classes/api/services/CHttpTest.php*.

httptest.update

Description

`object httptest.update(object/array webScenarios)`

This method allows you to update existing web scenarios.

> **Note:**
> This method is only available to *Admin* and *Super admin* user types. Permissions to call the method can be revoked in user role settings. See User roles for more information.

Parameters

`(object/array)` Web scenario properties to be updated.

The `httptestid` property must be defined for each web scenario, all other properties are optional. Only the passed properties will be updated, all others will remain unchanged.

Additionally to the standard web scenario properties, the method accepts the following parameters.

Parameter	Type	Description
steps	array	Scenario steps to replace existing steps.
tags	array	Web scenario tags.

Return values

`(object)` Returns an object containing the IDs of the updated web scenarios under the `httptestid` property.

Examples

Enabling a web scenario

Enable a web scenario, that is, set its status to ”0”.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "httptest.update",
    "params": {
        "httptestid": "5",
        "status": 0
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "httptestids": [
            "5"
        ]
    },
    "id": 1
}
```

See also

- Scenario step

Source

CHttpTest::update() in *ui/include/classes/api/services/CHttpTest.php*.

Appendix 1. Reference commentary

Notation Method signature

Each API method page contains the method signature.

Example: `object action.create(object/array actions)`

Signature	Description
`object`	The type of data the method returns.
`action.create`	The full name of the method.
`object/array`	The type of data the method accepts.
`actions`	A short summary of what the method accepts.

Data types

The Zabbix API supports the following data types as input:

Type	Description
ID	A unique identifier used to reference an entity.
boolean	A boolean value (either `true` or `false`).
flag	A value that is considered to be `true` if passed and not equal to `null`; otherwise, the value is considered to be `false`.
integer	A whole number.
float	A floating point number.
string	A text string.
text	A longer text string.
timestamp	A Unix timestamp.
array	An ordered sequence of values (a plain array).
object	An associative array.
query	A value that defines the data to be returned. The value can be defined as an array of property names (to return only specific properties), or as one of the predefined values: `extend` - returns all object properties; `count` - returns the number of retrieved records, supported only by certain subselects.

Attention:
Zabbix API always returns values as strings or arrays only.

Property behavior

Some of the object properties are marked with short labels to describe their behavior. The following labels are used:

- ***read-only*** - the value of the property is set automatically and cannot be defined or changed by the user, even in some specific conditions (e.g., *read-only* for inherited objects or discovered objects);
- ***write-only*** - the value of the property can be set, but cannot be accessed after;
- ***constant*** - the value of the property can be set when creating an object, but cannot be changed after;
- ***supported*** - the value of the property is not required to be set, but is allowed to be set in some specific conditions (e.g., *supported* if `type` is set to "Simple check", "External check", "SSH agent", "TELNET agent", or "HTTP agent"); note, however, that *supported* properties may still be set to their default values regardless of conditions;
- ***required*** - the value of the property is required to be set for all operations (except get operations) or in some specific conditions (e.g., *required* for create operations; *required* if `operationtype` is set to "global script" and `opcommand_hst` is not set).

Note:
For update operations, a property is considered as "set" when setting it during the update operation.

Properties that are not marked with labels are optional.

Parameter behavior

Some of the operation parameters are marked with short labels to describe their behavior for the operation. The following labels are used:

- ***read-only*** - the value of the parameter is set automatically and cannot be defined or changed by the user, even in some specific conditions (e.g., *read-only* for inherited objects or discovered objects);

- ***write-only*** - the value of the parameter can be set, but cannot be accessed after;
- ***supported*** - the value of the parameter is not required to be set, but is allowed to be set in some specific conditions (e.g., *supported* if `operating_mode` of Proxy object is set to "passive proxy"); note, however, that *supported* parameters may still be set to their default values regardless of conditions;
- ***required*** - the value of the parameter is required to be set.

Parameters that are not marked with labels are optional.

Reserved ID value "0" Reserved ID value "0" can be used to filter elements and to remove referenced objects. For example, to remove a referenced proxy from a host, proxyid should be set to 0 ("proxyid": "0") or to filter hosts monitored by server option proxyids should be set to 0 ("proxyids": "0").

Common "get" method parameters The following parameters are supported by all `get` methods:

Parameter	Type	Description
countOutput	boolean	Return the number of records in the result instead of the actual data.
editable	boolean	If set to `true`, return only objects that the user has write permissions to. Default: `false`.
excludeSearch	boolean	Return results that do not match the criteria given in the `search` parameter.
filter	object	Return only those results that exactly match the given filter. Accepts an object, where the keys are property names (e.g., Host object properties in `host.get`, Item object properties in `item.get`, etc.), and the values are either a single value or an array of values to match against. Does not support properties of `text` data type. Note that some methods have specific functionality for this parameter, which is described on the method page (e.g., the `filter` parameter in host.get also supports Host interface properties).
limit	integer	Limit the number of records returned.
output	query	Object properties to be returned. Note that the object ID (i.e., `hostid`, `itemid`, etc.) is always included in the response, even if it is not specified in the `output` parameter. Default: `extend`.
preservekeys	boolean	Use IDs as keys in the resulting array.
search	object	Return results that match the given pattern (case-insensitive). Accepts an object, where the keys are property names (e.g., Host object properties in `host.get`, Item object properties in `item.get`, etc.), and the values are strings to search for. If no additional options are given, this will perform a `LIKE "%…%"` search. Supports only properties of `string` and `text` data type. Note that some methods have specific functionality for this parameter, which is described on the method page (e.g., the `search` parameter in host.get also supports Host interface properties).
searchByAny	boolean	If set to `true`, return results that match any of the criteria given in the `filter` or `search` parameter instead of all of them. Default: `false`.
searchWildcardsEnabled	boolean	If set to `true`, enables the use of "*" as a wildcard character in the `search` parameter. Default: `false`.
sortfield	string/array	Sort the result by the given properties. Refer to a specific API get method description for a list of properties that can be used for sorting. Macros are not expanded before sorting. If no value is specified, data will be returned unsorted.

Parameter	Type	Description
sortorder	string/array	Order of sorting. If an array is passed, each value will be matched to the corresponding property given in the `sortfield` parameter. Possible values: `ASC` - *(default)* ascending; `DESC` - descending.
startSearch	boolean	The `search` parameter will compare the beginning of fields, that is, perform a `LIKE "…%"` search instead. Ignored if `searchWildcardsEnabled` is set to `true`.

Entity origin flags Get methods return a `flags` property for entities related to low-level discovery (LLD rule/LLD rule prototype, item/item prototype, etc). This property is useful to denote if the entity has been discovered or not, since editing for discovered entities is limited.

The `flags` property returns a result based on a combination (”+” operation) of these values:

Value	Description
0	Base entity (item, trigger, graph, host)
1	Low-level discovery rule
2	Any prototype (item prototype, trigger prototype, LLD rule prototype, etc)
4	Discovered entity (discovered item, trigger, graph, host, LLD rule)

The **combined** value returned by the `flags` property may be:

Value	Combination of	Description
0	0	Plain entity (item, trigger, graph, host).
2	2	Entity prototype (item prototype, trigger prototype, etc).
6	2+4	Discovered item, trigger, graph, host (converted from prototype).
1	1	Low-level discovery rule.
3	1+2	Low-level discovery rule prototype.
5	1+4	Discovered low-level discovery rule (converted from prototype).
7	1+2+4	Discovered low-level discovery rule prototype.

Examples User permission check

Does the user have permission to write to hosts whose names begin with ”MySQL” or ”Linux” ?

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "countOutput": true,
        "search": {
            "host": ["MySQL", "Linux"]
        },
        "editable": true,
        "startSearch": true,
        "searchByAny": true
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": "0",
```

```
    "id": 1
}
```

Note:
Zero result means no hosts with read/write permissions.

Mismatch counting

Count the number of hosts whose names do not contain the substring ”ubuntu”

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "countOutput": true,
        "search": {
            "host": "ubuntu"
        },
        "excludeSearch": true
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": "44",
    "id": 1
}
```

Searching for hosts using wildcards

Find hosts whose name contains word ”server” and have interface ports ”10050” or ”10071”. Sort the result by host name in descending order and limit it to 5 hosts.

Request:

```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["hostid", "host"],
        "selectInterfaces": ["port"],
        "filter": {
            "port": ["10050", "10071"]
        },
        "search": {
            "host": "*server*"
        },
        "searchWildcardsEnabled": true,
        "searchByAny": true,
        "sortfield": "host",
        "sortorder": "DESC",
        "limit": 5
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": [
        {
```

```
            "hostid": "50003",
            "host": "WebServer-Tomcat02",
            "interfaces": [
                {
                    "port": "10071"
                }
            ]
        },
        {
            "hostid": "50005",
            "host": "WebServer-Tomcat01",
            "interfaces": [
                {
                    "port": "10071"
                }
            ]
        },
        {
            "hostid": "50004",
            "host": "WebServer-Nginx",
            "interfaces": [
                {
                    "port": "10071"
                }
            ]
        },
        {
            "hostid": "99032",
            "host": "MySQL server 01",
            "interfaces": [
                {
                    "port": "10050"
                }
            ]
        },
        {
            "hostid": "99061",
            "host": "Linux server 01",
            "interfaces": [
                {
                    "port": "10050"
                }
            ]
        }
    ],
    "id": 1
}
```


Searching for hosts using wildcards with ”preservekeys”

If you add the parameter ”preservekeys” to the previous request, the result is returned as an associative array, where the keys are the id of the objects.

Request:


```
{
    "jsonrpc": "2.0",
    "method": "host.get",
    "params": {
        "output": ["hostid", "host"],
        "selectInterfaces": ["port"],
        "filter": {
            "port": ["10050", "10071"]
        },
```

```
        "search": {
            "host": "*server*"
        },
        "searchWildcardsEnabled": true,
        "searchByAny": true,
        "sortfield": "host",
        "sortorder": "DESC",
        "limit": 5,
        "preservekeys": true
    },
    "id": 1
}
```

Response:

```
{
    "jsonrpc": "2.0",
    "result": {
        "50003": {
            "hostid": "50003",
            "host": "WebServer-Tomcat02",
            "interfaces": [
                {
                    "port": "10071"
                }
            ]
        },
        "50005": {
            "hostid": "50005",
            "host": "WebServer-Tomcat01",
            "interfaces": [
                {
                    "port": "10071"
                }
            ]
        },
        "50004": {
            "hostid": "50004",
            "host": "WebServer-Nginx",
            "interfaces": [
                {
                    "port": "10071"
                }
            ]
        },
        "99032": {
            "hostid": "99032",
            "host": "MySQL server 01",
            "interfaces": [
                {
                    "port": "10050"
                }
            ]
        },
        "99061": {
            "hostid": "99061",
            "host": "Linux server 01",
            "interfaces": [
                {
                    "port": "10050"
                }
            ]
        }
```

```
    },
    "id": 1
}
```

Appendix 2. Changes from 8.0 to 8.2

Breaking changes

Other changes and bug fixes

Appendix 3. Changes in 8.2

8.2.1 This version is not released yet.

Extensions

Overview Although Zabbix offers a multiplicity of features, there is always room for additional functionality. Extensions are a convenient way of modifying and enhancing the monitoring capabilities of Zabbix without changing its source code.

You can extend Zabbix functionality either by using built-in extension options (trapper items, user parameters, etc.) or by using or creating custom extensions (loadable modules, plugins, etc.).

This section provides an overview with references to all the options for extending Zabbix.

Data collection with custom commands Trapper items

Trapper items are items that accept incoming data instead of querying for it. Trapper items are useful for sending specific data to Zabbix server or proxy, for example, periodic availability and performance data in the case of long-running user scripts. Sending data to Zabbix server or proxy is possible using the Zabbix sender utility or Zabbix sender protocol. Sending data to Zabbix server is also possible using the `history.push` API method.

External checks

An external check is an item for executing checks by running an executable, for example, a shell script or a binary.

External checks are executed by Zabbix server or proxy (when host is monitored by proxy), and do not require an agent running on the host being monitored.

User parameters

A user parameter is a user-defined command (associated with a user-defined key) that, when executed, can retrieve the data you need from the host where Zabbix agent is running. User parameters are useful for configuring agent or agent 2 items that are not predefined in Zabbix.

`system.run[]` Zabbix agent items

`system.run[]` Zabbix agent item is an item for a user-defined command (associated with a predefined key `system.run[]`, for example, `system.run[myscript.sh]`) that can be executed on the host where Zabbix agent is running.

Note:
`system.run[]` items are disabled by default and, if used, must be enabled (allowed) and defined in the Zabbix agent or Zabbix agent 2 configuration file (`AllowKey` or `AllowKeyRegexp` configuration parameter).

Attention:
User-defined commands in items such as external checks, user parameters and `system.run[]` Zabbix agent items are executed from the OS user that is used to run Zabbix components. To execute these commands, this user must have the necessary permissions.

HTTP agent items

HTTP agent item is an item for executing data requests over HTTP/HTTPS. HTTP agent items are useful for sending requests to HTTP endpoints to retrieve data from services such as *Elasticsearch* and *OpenWeatherMap*, for checking the status of Zabbix API or the status of Apache or NGINX web server, etc. HTTP agent items (with trapping enabled) can also function as trapper items.

Script items

A script item is an item for executing user-defined JavaScript code that retrieves data over HTTP/HTTPS. Script items are useful when the functionality provided by HTTP agent items is not enough. For example, in demanding data collection scenarios that require multiple steps or complex logic, a script item can be configured to make an HTTP call, then process the data received, and then pass the transformed value to a second HTTP call.

> **Note:**
> HTTP agent items and script items are supported by Zabbix server and proxy, and do not require an agent running on the host being monitored.

Advanced extensions Loadable modules

Loadable modules, written in C, are a versatile and performance-minded option for extending the functionality of Zabbix components (server, proxy, agent) on UNIX platforms. A loadable module is basically a shared library used by Zabbix daemon and loaded on startup. The library should contain certain functions, so that a Zabbix process may detect that the file is indeed a module it can load and work with.

Loadable modules have a number of benefits, including the ability to add new metrics or implement any other logic (for example, Zabbix history data export), great performance, and the option to develop, use and share the functionality they provide. It contributes to trouble-free maintenance and helps to deliver new functionality easier and independently of the Zabbix core code base.

Loadable modules are especially useful in a complex monitoring setup. When monitoring embedded systems, having a large number of monitored parameters or heavy scripts with complex logic or long startup time, extensions such as user parameters, `system.run[]` Zabbix agent items, and external checks will have an impact on performance. Loadable modules offer a way of extending Zabbix functionality without sacrificing performance.

Plugins

Plugins extend the monitoring capabilities of Zabbix agent 2. Written in the Go programming language, they provide an alternative to loadable modules (written in C).

It is possible to develop your own plugin (independent data collector) for Zabbix agent 2. See Developer center.

Alert customization Webhooks

A webhook is a Zabbix media type that provides an option to extend Zabbix alerting capabilities to external software such as helpdesk systems, chats, or messengers. Similarly to script items, webhooks are useful for making HTTP calls using custom JavaScript code, for example, to push notifications to different platforms such as Microsoft Teams, Discord, and Jira. It is also possible to return some data (for example, about created helpdesk tickets) that is then displayed in Zabbix.

Existing webhooks are available in the Zabbix Git repository. For custom webhook development, see Webhook development guidelines.

Alert scripts

An alert script is a Zabbix media type that provides an option to create an alternative way (script) to handle Zabbix alerts. Alert scripts are useful if you are not satisfied with the existing media types for sending alerts in Zabbix.

Frontend customization Custom themes

It is possible to change Zabbix frontend visual appearance by using custom themes. See the instructions on creating and applying your own themes.

Frontend modules

Frontend modules provide an option to extend Zabbix frontend functionality by adding third-party modules or by developing your own. With frontend modules you can add new menu items, their respective views, actions, etc.

Global scripts A global script is a user-defined set of commands that can be executed on a monitoring target (by shell (/bin/sh) interpreter), depending on the configured scope and user permissions. Global scripts can be configured for the following actions:

- Action operation
- Manual host action
- Manual event action

Global scripts are useful in many cases. For example, if configured for action operations or manual host actions, you can use global scripts to automatically or manually execute remote commands such as restarting an application (web server, middleware, CRM, etc.) or freeing disk space (removing older files, cleaning `/tmp`, etc). Or, another example, if configured for manual event actions, you can use global scripts to manage problem tickets in external systems.

Global scripts can be executed by Zabbix server, proxy or agent.

> **Attention:**
> User-defined commands are executed from the OS user that is used to run Zabbix components. To execute these commands, this user must have the necessary permissions.

Zabbix API Zabbix API is an HTTP-based API that is part of Zabbix frontend. With Zabbix API, you can do any of the following operations:

- Programmatically retrieve and modify the configuration of Zabbix.
- Import and export Zabbix configuration.
- Access Zabbix historical and trend data.
- Configure applications to work with Zabbix.
- Integrate Zabbix with third-party software.
- Automate routine tasks.

Zabbix API consists of a multiplicity of methods that are nominally grouped into separate APIs. Each method performs a specific task. For the available methods, as well as an overview of the functions provided by Zabbix API, see Zabbix API Method reference.

Loadable modules

Overview

Loadable modules offer a performance-minded option for extending Zabbix functionality.

You can extend Zabbix functionality in many ways, for example, with user parameters, external checks, and system.run Zabbix agent items. These work very well, but have one major drawback, namely fork(). Zabbix has to fork a new process every time it handles a user metric, which is not good for performance. It is not a big deal normally, however it could be a serious issue when monitoring embedded systems, having a large number of monitored parameters or heavy scripts with complex logic or long startup time.

Support of loadable modules offers ways for extending Zabbix agent, server and proxy without sacrificing performance.

A loadable module is basically a shared library used by Zabbix daemon and loaded on startup. The library should contain certain functions, so that a Zabbix process may detect that the file is indeed a module it can load and work with.

Loadable modules have a number of benefits. Great performance and ability to implement any logic are very important, but perhaps the most important advantage is the ability to develop, use and share Zabbix modules. It contributes to trouble-free maintenance and helps to deliver new functionality easier and independently of the Zabbix core code base.

Module licensing and distribution in binary form is governed by the AGPL-3.0 license (modules are linking with Zabbix in runtime and are using Zabbix headers; the whole Zabbix code is licensed under AGPL-3.0 license since Zabbix 7.0). Binary compatibility is not guaranteed by Zabbix.

Module API stability is guaranteed during one Zabbix LTS (Long Term Support) release cycle. Stability of Zabbix API is not guaranteed (technically it is possible to call Zabbix internal functions from a module, but there is no guarantee that such modules will work).

Module API

In order for a shared library to be treated as a Zabbix module, it should implement and export several functions. There are currently six functions in the Zabbix module API, only one of which is mandatory and the other five are optional.

Mandatory interface

The only mandatory function is `zbx_module_api_version()`:

```
int zbx_module_api_version(void);
```

This function should return the API version implemented by this module and in order for the module to be loaded this version must match module API version supported by Zabbix. Version of module API supported by Zabbix is ZBX_MODULE_API_VERSION. So this function should return this constant. Old constant ZBX_MODULE_API_VERSION_ONE used for this purpose is now defined to equal ZBX_MODULE_API_VERSION to preserve source compatibility, but it's usage is not recommended.

Optional interface

The optional functions are zbx_module_init(), zbx_module_item_list(), zbx_module_item_timeout(), zbx_module_history_w and zbx_module_uninit():

```
int zbx_module_init(void);
```

This function should perform the necessary initialization for the module (if any). If successful, it should return ZBX_MODULE_OK. Otherwise, it should return ZBX_MODULE_FAIL. In the latter case Zabbix will not start.

```
ZBX_METRIC  *zbx_module_item_list(void);
```

This function should return a list of items supported by the module. Each item is defined in a ZBX_METRIC structure, see the section below for details. The list is terminated by a ZBX_METRIC structure with "key" field of NULL.

```
void    zbx_module_item_timeout(int timeout);
```

If module exports zbx_module_item_list() then this function is used by Zabbix to specify the timeout settings in Zabbix configuration file that the item checks implemented by the module should obey. Here, the "timeout" parameter is in seconds.

```
ZBX_HISTORY_WRITE_CBS   zbx_module_history_write_cbs(void);
```

This function should return callback functions Zabbix server will use to export history of different data types. Callback functions are provided as fields of ZBX_HISTORY_WRITE_CBS structure, fields can be NULL if module is not interested in the history of certain type.

```
int zbx_module_uninit(void);
```

This function should perform the necessary uninitialization (if any) like freeing allocated resources, closing file descriptors, etc.

All functions are called once on Zabbix startup when the module is loaded, with the exception of zbx_module_uninit(), which is called once on Zabbix shutdown when the module is unloaded.

Defining items

Each item is defined in a ZBX_METRIC structure:

```
typedef struct
{
    char        *key;
    unsigned    flags;
    int     (*function)();
    char        *test_param;
}
ZBX_METRIC;
```

Here, **key** is the item key (e.g., "dummy.random"), **flags** is either CF_HAVEPARAMS or 0 (depending on whether the item accepts parameters or not), **function** is a C function that implements the item (e.g., zbx_module_dummy_random), and **test_param** is the parameter list to be used when Zabbix agent is started with the -p flag (e.g., "1,1000", can be NULL). An example definition may look like this:

```
static ZBX_METRIC keys[] =
{
    { "dummy.random", CF_HAVEPARAMS, zbx_module_dummy_random, "1,1000" },
    { NULL }
}
```

Each function that implements an item should accept two pointer parameters, the first one of type AGENT_REQUEST and the second one of type AGENT_RESULT:

```
int zbx_module_dummy_random(AGENT_REQUEST *request, AGENT_RESULT *result)
{
    ...
```

```
    SET_UI64_RESULT(result, from + rand() % (to - from + 1));

    return SYSINFO_RET_OK;
}
```

These functions should return SYSINFO_RET_OK, if the item value was successfully obtained. Otherwise, they should return SYSINFO_RET_FAIL. See example "dummy" module below for details on how to obtain information from AGENT_REQUEST and how to set information in AGENT_RESULT.

Providing history export callbacks

Attention:
History export via module is no longer supported by Zabbix proxy.

Module can specify functions to export history data by type: Numeric (float), Numeric (unsigned), Character, Text and Log:

```
typedef struct
{
    void    (*history_float_cb)(const ZBX_HISTORY_FLOAT *history, int history_num);
    void    (*history_integer_cb)(const ZBX_HISTORY_INTEGER *history, int history_num);
    void    (*history_string_cb)(const ZBX_HISTORY_STRING *history, int history_num);
    void    (*history_text_cb)(const ZBX_HISTORY_TEXT *history, int history_num);
    void    (*history_log_cb)(const ZBX_HISTORY_LOG *history, int history_num);
}
ZBX_HISTORY_WRITE_CBS;
```

Each of them should take history array of history_num elements as arguments. Depending on history data type to be exported, history is an array of the following structures, respectively:

```
typedef struct
{
    zbx_uint64_t    itemid;
    int     clock;
    int     ns;
    double      value;
}
ZBX_HISTORY_FLOAT;

typedef struct
{
    zbx_uint64_t    itemid;
    int     clock;
    int     ns;
    zbx_uint64_t    value;
}
ZBX_HISTORY_INTEGER;

typedef struct
{
    zbx_uint64_t    itemid;
    int     clock;
    int     ns;
    const char  *value;
}
ZBX_HISTORY_STRING;

typedef struct
{
    zbx_uint64_t    itemid;
    int     clock;
    int     ns;
    const char  *value;
}
ZBX_HISTORY_TEXT;
```

```
typedef struct
{
    zbx_uint64_t    itemid;
    int     clock;
    int     ns;
    const char  *value;
    const char  *source;
    int     timestamp;
    int     logeventid;
    int     severity;
}
ZBX_HISTORY_LOG;
```

Callbacks will be used by Zabbix server history syncer processes in the end of history sync procedure after data is written into Zabbix database and saved in value cache.

Attention:
In case of internal error in history export module, it is recommended that module is written in such a way that it does not block whole monitoring until it recovers but discards data instead and allows Zabbix server to continue running.

Building modules

Modules are currently meant to be built inside Zabbix source tree, because the module API depends on some data structures that are defined in Zabbix headers.

The most important header for loadable modules is `include/module.h`, which defines these data structures. Other necessary system headers that help `include/module.h` to work properly are `stdlib.h` and `stdint.h`.

With this information in mind, everything is ready for the module to be built. The module should include `stdlib.h`, `stdint.h` and `module.h`, and the build script should make sure that these files are in the include path. See example "dummy" module below for details.

Another useful header is `include/zbxcommon.h`, which defines `zabbix_log()` function, which can be used for logging and debugging purposes.

Configuration parameters

Zabbix agent, server and proxy support two parameters to deal with modules:

- `LoadModulePath` - full path to the location of loadable modules
- `LoadModule` - module(s) to load at startup.

The modules must be located in a directory specified by `LoadModulePath` or the path must precede the module name. If the preceding path is absolute (starts with /) then `LoadModulePath` is ignored. Multiple `LoadModule` parameters are allowed.

For example, to extend Zabbix agent we could add the following parameters:

```
LoadModulePath=/usr/local/lib/zabbix/agent/
LoadModule=mariadb.so
LoadModule=apache.so
LoadModule=kernel.so
LoadModule=/usr/local/lib/zabbix/dummy.so
```

Upon agent startup it will load the mariadb.so, apache.so and kernel.so modules from the `/usr/local/lib/zabbix/agent` directory while dummy.so will be loaded from `/usr/local/lib/zabbix`. The agent will fail to start if a module is missing, in case of bad permissions or if a shared library is not a Zabbix module.

Frontend configuration

Loadable modules are supported by Zabbix agent, server and proxy. Therefore, item type in Zabbix frontend depends on where the module is loaded. If the module is loaded into the agent, then the item type should be "Zabbix agent" or "Zabbix agent (active)". If the module is loaded into server or proxy, then the item type should be "Simple check".

History export through Zabbix modules does not need any frontend configuration. If the module is successfully loaded by server and provides `zbx_module_history_write_cbs()` function which returns at least one non-NULL callback function then history export will be enabled automatically.

Dummy module

Zabbix includes a sample module written in C language. The module is located under `src/modules/dummy`:

```
alex@alex:~trunk/src/modules/dummy$ ls -l
-rw-rw-r-- 1 alex alex 9019 Apr 24 17:54 dummy.c
-rw-rw-r-- 1 alex alex   67 Apr 24 17:54 Makefile
-rw-rw-r-- 1 alex alex  245 Apr 24 17:54 README
```

The module is well documented, it can be used as a template for your own modules.

After `./configure` has been run in the root of Zabbix source tree as described above, just run `make` in order to build `dummy.so`.

```
/*
** Zabbix
** Copyright (C) 2001-2020 Zabbix SIA
**
** This program is free software; you can redistribute it and/or modify
** it under the terms of the GNU General Public License as published by
** the Free Software Foundation; either version 2 of the License, or
** (at your option) any later version.
**
** This program is distributed in the hope that it will be useful,
** but WITHOUT ANY WARRANTY; without even the implied warranty of
** MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
** GNU General Public License for more details.
**
** You should have received a copy of the GNU General Public License
** along with this program; if not, write to the Free Software
** Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301, USA.
**/

###include <stdlib.h>
###include <string.h>
###include <time.h>
###include <stdint.h>

###include "module.h"

/* the variable keeps timeout setting for item processing */
static int  item_timeout = 0;

/* module SHOULD define internal functions as static and use a naming pattern different from Zabbix intern
/* symbols (zbx_*) and loadable module API functions (zbx_module_*) to avoid conflicts
static int  dummy_ping(AGENT_REQUEST *request, AGENT_RESULT *result);
static int  dummy_echo(AGENT_REQUEST *request, AGENT_RESULT *result);
static int  dummy_random(AGENT_REQUEST *request, AGENT_RESULT *result);

static ZBX_METRIC keys[] =
/*  KEY         FLAG        FUNCTION    TEST PARAMETERS */
{
    {"dummy.ping",      0,      dummy_ping, NULL},
    {"dummy.echo",      CF_HAVEPARAMS,  dummy_echo, "a message"},
    {"dummy.random",    CF_HAVEPARAMS,  dummy_random,   "1,1000"},
    {NULL}
};

/******************************************************************************
 *                                                                            *
 * Function: zbx_module_api_version                                           *
 *                                                                            *
 * Purpose: returns version number of the module interface                    *
 *                                                                            *
 * Return value: ZBX_MODULE_API_VERSION - version of module.h module is       *
 *               compiled with, in order to load module successfully Zabbix   *
 *               MUST be compiled with the same version of this header file   *
 *                                                                            *
 ******************************************************************************/
```

```c
int zbx_module_api_version(void)
{
    return ZBX_MODULE_API_VERSION;
}

/******************************************************************************
 *                                                                            *
 * Function: zbx_module_item_timeout                                          *
 *                                                                            *
 * Purpose: set timeout value for processing of items                         *
 *                                                                            *
 * Parameters: timeout - timeout in seconds, 0 - no timeout set               *
 *                                                                            *
 ******************************************************************************/
void    zbx_module_item_timeout(int timeout)
{
    item_timeout = timeout;
}

/******************************************************************************
 *                                                                            *
 * Function: zbx_module_item_list                                             *
 *                                                                            *
 * Purpose: returns list of item keys supported by the module                 *
 *                                                                            *
 * Return value: list of item keys                                            *
 *                                                                            *
 ******************************************************************************/
ZBX_METRIC  *zbx_module_item_list(void)
{
    return keys;
}

static int  dummy_ping(AGENT_REQUEST *request, AGENT_RESULT *result)
{
    SET_UI64_RESULT(result, 1);

    return SYSINFO_RET_OK;
}

static int  dummy_echo(AGENT_REQUEST *request, AGENT_RESULT *result)
{
    char    *param;

    if (1 != request->nparam)
    {
        /* set optional error message */
        SET_MSG_RESULT(result, strdup("Invalid number of parameters."));
        return SYSINFO_RET_FAIL;
    }

    param = get_rparam(request, 0);

    SET_STR_RESULT(result, strdup(param));

    return SYSINFO_RET_OK;
}

/******************************************************************************
 *                                                                            *
 * Function: dummy_random                                                     *
 *                                                                            *
```

```
 * Purpose: a main entry point for processing of an item                      *
 *                                                                            *
 * Parameters: request - structure that contains item key and parameters      *
 *              request->key - item key without parameters                    *
 *              request->nparam - number of parameters                        *
 *              request->params[N-1] - pointers to item key parameters        *
 *              request->types[N-1] - item key parameters types:              *
 *                  REQUEST_PARAMETER_TYPE_UNDEFINED (key parameter is empty) *
 *                  REQUEST_PARAMETER_TYPE_ARRAY (array)                      *
 *                  REQUEST_PARAMETER_TYPE_STRING (quoted or unquoted string) *
 *                                                                            *
 *             result - structure that will contain result                    *
 *                                                                            *
 * Return value: SYSINFO_RET_FAIL - function failed, item will be marked      *
 *                                  as not supported by zabbix                *
 *               SYSINFO_RET_OK - success                                     *
 *                                                                            *
 * Comment: get_rparam(request, N-1) can be used to get a pointer to the Nth  *
 *          parameter starting from 0 (first parameter). Make sure it exists  *
 *          by checking value of request->nparam.                             *
 *          In the same manner get_rparam_type(request, N-1) can be used to   *
 *          get a parameter type.                                             *
 *                                                                            *
 ******************************************************************************/
static int  dummy_random(AGENT_REQUEST *request, AGENT_RESULT *result)
{
    char    *param1, *param2;
    int from, to;

    if (2 != request->nparam)
    {
        /* set optional error message */
        SET_MSG_RESULT(result, strdup("Invalid number of parameters."));
        return SYSINFO_RET_FAIL;
    }

    param1 = get_rparam(request, 0);
    param2 = get_rparam(request, 1);

    /* there is no strict validation of parameters and types for simplicity sake */
    from = atoi(param1);
    to = atoi(param2);

    if (from > to)
    {
        SET_MSG_RESULT(result, strdup("Invalid range specified."));
        return SYSINFO_RET_FAIL;
    }

    SET_UI64_RESULT(result, from + rand() % (to - from + 1));

    return SYSINFO_RET_OK;
}

/******************************************************************************
 *                                                                            *
 * Function: zbx_module_init                                                  *
 *                                                                            *
 * Purpose: the function is called on agent startup                           *
 *          It should be used to call any initialization routines             *
 *                                                                            *
 * Return value: ZBX_MODULE_OK - success                                      *
```

```
 *               ZBX_MODULE_FAIL - module initialization failed                *
 *                                                                            *
 * Comment: the module won't be loaded in case of ZBX_MODULE_FAIL             *
 *                                                                            *
 ******************************************************************************/
int zbx_module_init(void)
{
    /* initialization for dummy.random */
    srand(time(NULL));

    return ZBX_MODULE_OK;
}

/******************************************************************************
 *                                                                            *
 * Function: zbx_module_uninit                                                *
 *                                                                            *
 * Purpose: the function is called on agent shutdown                          *
 *          It should be used to cleanup used resources if there are any      *
 *                                                                            *
 * Return value: ZBX_MODULE_OK - success                                      *
 *               ZBX_MODULE_FAIL - function failed                            *
 *                                                                            *
 ******************************************************************************/
int zbx_module_uninit(void)
{
    return ZBX_MODULE_OK;
}

/******************************************************************************
 *                                                                            *
 * Functions: dummy_history_float_cb                                          *
 *            dummy_history_integer_cb                                        *
 *            dummy_history_string_cb                                         *
 *            dummy_history_text_cb                                           *
 *            dummy_history_log_cb                                            *
 *                                                                            *
 * Purpose: callback functions for storing historical data of types float,    *
 *          integer, string, text and log respectively in external storage    *
 *                                                                            *
 * Parameters: history     - array of historical data                         *
 *             history_num - number of elements in history array              *
 *                                                                            *
 ******************************************************************************/
static void dummy_history_float_cb(const ZBX_HISTORY_FLOAT *history, int history_num)
{
    int i;

    for (i = 0; i < history_num; i++)
    {
        /* do something with history[i].itemid, history[i].clock, history[i].ns, history[i].value, ... */
    }
}

static void dummy_history_integer_cb(const ZBX_HISTORY_INTEGER *history, int history_num)
{
    int i;

    for (i = 0; i < history_num; i++)
    {
        /* do something with history[i].itemid, history[i].clock, history[i].ns, history[i].value, ... */
    }
```

```
}

static void dummy_history_string_cb(const ZBX_HISTORY_STRING *history, int history_num)
{
    int i;

    for (i = 0; i < history_num; i++)
    {
        /* do something with history[i].itemid, history[i].clock, history[i].ns, history[i].value, ... */
    }
}

static void dummy_history_text_cb(const ZBX_HISTORY_TEXT *history, int history_num)
{
    int i;

    for (i = 0; i < history_num; i++)
    {
        /* do something with history[i].itemid, history[i].clock, history[i].ns, history[i].value, ... */
    }
}

static void dummy_history_log_cb(const ZBX_HISTORY_LOG *history, int history_num)
{
    int i;

    for (i = 0; i < history_num; i++)
    {
        /* do something with history[i].itemid, history[i].clock, history[i].ns, history[i].value, ... */
    }
}

/******************************************************************************
 *                                                                            *
 * Function: zbx_module_history_write_cbs                                     *
 *                                                                            *
 * Purpose: returns a set of module functions Zabbix will call to export      *
 *          different types of historical data                                *
 *                                                                            *
 * Return value: structure with callback function pointers (can be NULL if    *
 *               module is not interested in data of certain types)           *
 *                                                                            *
 ******************************************************************************/
ZBX_HISTORY_WRITE_CBS   zbx_module_history_write_cbs(void)
{
    static ZBX_HISTORY_WRITE_CBS    dummy_callbacks =
    {
        dummy_history_float_cb,
        dummy_history_integer_cb,
        dummy_history_string_cb,
        dummy_history_text_cb,
        dummy_history_log_cb,
    };

    return dummy_callbacks;
}
```

The module exports three new items:

- `dummy.ping` - always returns '1'
- `dummy.echo[param1]` - returns the first parameter as it is, for example, `dummy.echo[ABC]` will return ABC
- `dummy.random[param1, param2]` - returns a random number within the range of param1-param2, for example, `dummy.random[1,1000000]`

Limitations

Support of loadable modules is implemented for the Unix platform only. It means that it does not work for Windows agents.

In some cases, a module may need to read module-related configuration parameters from `zabbix_agentd.conf`. It is not supported currently. If you need your module to use some configuration parameters you should probably implement parsing of a module-specific configuration file.

Frontend modules

Overview

It is possible to enhance Zabbix frontend functionality by adding third-party modules or by developing your own modules without the need to change the source code of Zabbix.

Note that the module code will run with the same privileges as Zabbix source code. This means:

- Third-party modules can be harmful. You must trust the modules you are installing.
- Errors in a third-party module code may crash the frontend. If this happens, just remove the module code from the frontend. As soon as you reload Zabbix frontend, you'll see a note saying that some modules are absent. Go to Module administration (in *Administration* > *General* > *Modules*) and click *Scan directory* again to remove non-existent modules from the database.

Installation

Please always read the installation manual for a particular module. It is recommended to install new modules one by one to catch failures easily.

Just before you install a module:

- Make sure you have downloaded the module from a trusted source. Installation of harmful code may lead to consequences, such as data loss.
- Different versions of the same module (same ID) can be installed in parallel, but only a single version can be enabled at once.

Steps to install a module:

- Unpack your module into its own folder in the `modules` directory of the Zabbix frontend installation (for example, `zabbix/ui/modules`).
- Ensure that your module folder contains at least the `manifest.json` file.
- Navigate to Module administration and click the *Scan directory* button.
- New module will appear in the list along with its version, author, description and status.
- Enable module by clicking its status.

Troubleshooting:

Problem	Solution
Module did not appear in the list	If your module didn't appear in the list, make sure `manifest.json` exists in `modules/your-module/`. If it doesn't, you might have unpacked the module to the wrong directory. If it does, the module might not be compatible with your Zabbix version. Also, check that the web server user has at least read and search access (`r-x`) to the module folder and all subdirectories, and read access (`r--`) to all files inside.
Frontend crashed	The module code is not compatible with the current Zabbix version or server configuration. Please delete module files and reload the frontend. You'll see a notice that some modules are absent. Go to Module administration and click *Scan directory* again to remove non-existent modules from the database.
Error message about identical namespace, ID or actions appears	New module tried to register a namespace, ID or actions which are already registered by other enabled modules. Disable the conflicting module (mentioned in error message) prior to enabling the new one.
Technical error messages appear	Report errors to the developer of the module.

Developing modules

For information about developing custom modules, see Developer center.

Appendixes

Overview This section contains reference materials, configuration parameters, technical specifications, and procedural guides that support the main manual.

Essential reference

- Installation and setup - database creation, server and agent installation and configuration.
- Process configuration - configuration parameters for Zabbix server, proxy, agent, and more.
- Items - item key reference and supported data types.
- Macros - built-in and user macro syntax and usage.

Installation and setup

This section covers database creation, security configuration, platform-specific setup procedures, and feature configuration.

Please use the sidebar to find the procedure relevant to your environment.

Database creation

Overview

A Zabbix database must be created during the installation of Zabbix server or proxy.

This section provides instructions for creating a Zabbix database. A separate set of instructions is available for each supported database.

> **Note:**
> To improve database security by creating database roles/users with minimal privileges, see database creation best practices for each supported database:

> - MySQL/MariaDB
> - PostgreSQL/TimescaleDB
>
> For configuring secure TLS connections, see Secure connection to the database.

UTF-8 is the only encoding supported by Zabbix. It is known to work without any security flaws. Users should be aware that there are known security issues if using some of the other encodings. For switching to UTF-8, see Repairing Zabbix database character set and collation. See also Limits of filtering with utf8mb4 collations.

> **Note:**
> If installing from Zabbix Git repository, you need to run the following command prior to proceeding to the next steps:

 `make dbschema`

MySQL/MariaDB

The utf8mb4 character set with the utf8mb4_bin collation is supported for Zabbix server/proxy to work properly with MySQL/MariaDB databases.
 To avoid future problems, it is strongly recommended to use the utf8mb4 character set with the utf8mb4_bin collation for both new and existing installations. To convert databases using the utf8mb3 character set and utf8mb3_bin collation to utf8mb4 character set and utf8mb4_bin collation, follow the instructions in Repairing Zabbix database character set and collation.

Deterministic triggers need to be created during the import of schema. On MySQL and MariaDB, this requires `GLOBAL log_bin_trust_function_creators = 1` to be set if binary logging is enabled and there is no superuser privileges and `log_bin_trust_function_creators = 1` is not set in MySQL configuration file.

If you are installing Zabbix from packages or containers, see the respective instructions.

If you are installing Zabbix from sources:

- Create and configure a database and a user. For security reasons, use a strong, unique password for the Zabbix database user.

```
mysql -uroot -p<password>

mysql> create database zabbix character set utf8mb4 collate utf8mb4_bin;
mysql> create user 'zabbix'@'localhost' identified by '<password>';
mysql> grant all privileges on zabbix.* to 'zabbix'@'localhost';
mysql> SET GLOBAL log_bin_trust_function_creators = 1;
mysql> quit;
```

- Import the data into the database and set utf8mb4 character set as default (assuming you are in the root directory of Zabbix sources). For a Zabbix proxy database, only `schema.sql` should be imported (no images.sql nor data.sql).

```
cd database/mysql
mysql -uzabbix -p<password> zabbix < schema.sql
#### stop here if you are creating database for Zabbix proxy
mysql -uzabbix -p<password> zabbix < images.sql
mysql -uzabbix -p<password> --default-character-set=utf8mb4 zabbix < data.sql
```

`log_bin_trust_function_creators` can be disabled after the schema has been successfully imported:

```
mysql -uroot -p<password>

mysql> SET GLOBAL log_bin_trust_function_creators = 0;
mysql> quit;
```

PostgreSQL

You need to have database user with permissions to create database objects.

If you are installing Zabbix from packages or containers, see the respective instructions.

If you are installing Zabbix from sources:

- Create a database user.

The following shell command will create user `zabbix`. Specify a password when prompted and repeat the password (note, you may first be asked for `sudo` password):

```
sudo -u postgres createuser --pwprompt zabbix
```

- Create a database.

The following shell command will create the database `zabbix` (last parameter) with the previously created user as the owner (`-O zabbix`).

```
sudo -u postgres createdb -O zabbix -E Unicode -T template0 zabbix
```

- Import the initial schema and data (assuming you are in the root directory of Zabbix sources). For a Zabbix proxy database, only `schema.sql` should be imported (no images.sql nor data.sql).

```
cd database/postgresql
cat schema.sql | sudo -u zabbix psql zabbix
#### stop here if you are creating database for Zabbix proxy
cat images.sql | sudo -u zabbix psql zabbix
cat data.sql | sudo -u zabbix psql zabbix
```

Attention:
The above commands are provided as an example that will work in most of GNU/Linux installations. You can use different commands depending on how your system/database is configured, for example:

 `psql -U <username>`

 If you have any trouble setting up the database, please consult your Database administrator.

TimescaleDB

Instructions for creating and configuring TimescaleDB are provided in a separate section.

Elasticsearch and ClickHouse

Zabbix can also store history data in Elasticsearch or ClickHouse. For setup instructions, see the linked pages.

SQLite

Using SQLite is supported for **Zabbix proxy** only!

The database will be automatically created if it does not exist.

Return to the installation section.

Repairing Zabbix database character set and collation

MySQL/MariaDB

Historically, MySQL and derivatives used 'utf8' as an alias for utf8mb3 - MySQL's own 3-byte implementation of the standard UTF8, which is 4-byte. Starting from MySQL 8.0.28 and MariaDB 10.6.1, the 'utf8mb3' character set is deprecated and at some point its support will be dropped while 'utf8' will become a reference to 'utf8mb4'. Since Zabbix 6.0, 'utf8mb4' is supported. To avoid future problems, it is highly recommended to use 'utf8mb4'. Another advantage of switching to 'utf8mb4' is support of supplementary Unicode characters.

Warning:
As versions before Zabbix 6.0 are not aware of utf8mb4, make sure to first upgrade Zabbix server and DB schema to 6.0.x or later before executing utf8mb4 conversion.

1. Check the database character set and collation.

For example:

```
mysql> SELECT @@character_set_database, @@collation_database;
+--------------------------+----------------------+
| @@character_set_database | @@collation_database |
+--------------------------+----------------------+
| latin2                   | latin2_general_ci    |
+--------------------------+----------------------+
```

Or:

```
mysql> SELECT @@character_set_database, @@collation_database;
+--------------------------+----------------------+
| @@character_set_database | @@collation_database |
+--------------------------+----------------------+
| utf8                     | utf8_bin             |
+--------------------------+----------------------+
```

As we see, the character set here is not 'utf8mb4' and collation is not 'utf8mb4_bin', so we need to fix them.

2. Stop Zabbix.

3. Create a backup copy of the database!

4. Fix the character set and collation on database level:

```
alter database <your DB name> character set utf8mb4 collate utf8mb4_bin;
```

Fixed values:

```
mysql> SELECT @@character_set_database, @@collation_database;
+--------------------------+----------------------+
| @@character_set_database | @@collation_database |
+--------------------------+----------------------+
| utf8mb4                  | utf8mb4_bin          |
+--------------------------+----------------------+
```

5. Load the script to fix character set and collation on table and column level:

```
mysql <your DB name> < utf8mb4_convert.sql
```

6. Execute the script:

```
                SET @ZABBIX_DATABASE = '<your DB name>';
If MariaDB →    set innodb_strict_mode = OFF;
                CALL zbx_convert_utf8();
If MariaDB →    set innodb_strict_mode = ON;
                drop procedure zbx_convert_utf8;
```

Please note that 'utf8mb4' is expected to consume slightly more disk space.

7. If no errors - you may want to create a database backup copy with the fixed database.

8. Start Zabbix.

Database upgrade to primary keys

Overview

This section provides instructions for manually upgrading tables in existing installations to primary keys.

Upgrading to primary keys optimizes how data is indexed and accessed, which may speed up queries and save space. It also improves data management and synchronization in clustered setups, helping with scaling and ensuring the system remains reliable even if some servers fail.

> **Attention:**
> The instructions provided on this page are designed for advanced users and may need to be adjusted for your specific configuration. Upgrading to primary keys can be time-consuming and resource-intensive. Ensure that enough free disk space is available; depending on your database size and stored data, the process may require up to 2.5 times the space currently used by history tables.

Primary keys are used for all tables in new installations since Zabbix 6.0.

There is no automatic database upgrade to primary keys; however, existing installations may be upgraded manually **after** upgrading Zabbix server to 6.0 or newer.

> **Attention:**
> Since Zabbix 7.0, upgrading tables to primary keys also upgrades tables to use double precision data types.

 If you have Zabbix 7.0 (or later), tables already use double precision. However, the instructions on this page can still be followed to upgrade tables to primary keys without affecting tables that already use double precision.

 If you have Zabbix 6.4 (or earlier), consider upgrading tables to double precision first. For more information, see Upgrading to numeric values of extended range in Zabbix 7.0 documentation.

Instructions are available for:

- MariaDB/MySQL
- PostgreSQL
- TimescaleDB

Important notes

To perform the database upgrade:

1. Stop Zabbix server.

Stopping Zabbix server for the time of the upgrade is strongly recommended. However, if absolutely necessary, you can perform the upgrade while the server is running (only for MySQL, MariaDB, and PostgreSQL without TimescaleDB).

2. Back up your database.
3. Install the latest zabbix-sql-scripts package compatible with your Zabbix version (e.g., for RHEL: `dnf install zabbix-sql-scripts`).
4. Run the scripts for your database.
5. Start Zabbix server.

> **Warning:**
> Run the scripts for the server database only. The proxy will not benefit from this upgrade.

If the database uses partitions, contact the DB administrator or Zabbix Support for help.

CSV files can be removed after a successful upgrade to primary keys.

Optionally, Zabbix frontend may be switched to maintenance mode.

MariaDB/MySQL

Export and import must be performed in tmux/screen to ensure that the session isn't dropped.

See also: Important notes

MySQL 8.4+ with mysqlsh

This method can be used with a running Zabbix server, but it is recommended to stop the server for the time of the upgrade. The MySQL Shell (*mysqlsh*) must be installed and able to connect to the DB.

- Log in to MySQL console as root (recommended) or as any user with FILE privileges.
- Start MySQL with the local_infile variable enabled.
- Rename old tables and create new tables by running `history_upgrade_prepare.sql`.

```
mysql -uzabbix -p<password> zabbix < /usr/share/zabbix/sql-scripts/mysql/option-patches/history_upgrade_pr
```

- Export and import data.

Connect via mysqlsh. If using a socket connection, specifying the path might be required.

```
sudo mysqlsh -uroot -S /run/mysqld/mysqld.sock --no-password -Dzabbix
```

Switch into JavaScript mode using:

```
\js
```

Then run the code below (CSVPATH can be changed as needed):

```
CSVPATH="/var/lib/mysql-files";

util.exportTable("history_old", CSVPATH + "/history.csv", { dialect: "csv" });
util.importTable(CSVPATH + "/history.csv", {"dialect": "csv", "table": "history" });

util.exportTable("history_uint_old", CSVPATH + "/history_uint.csv", { dialect: "csv" });
util.importTable(CSVPATH + "/history_uint.csv", {"dialect": "csv", "table": "history_uint" });

util.exportTable("history_str_old", CSVPATH + "/history_str.csv", { dialect: "csv" });
util.importTable(CSVPATH + "/history_str.csv", {"dialect": "csv", "table": "history_str" });

util.exportTable("history_log_old", CSVPATH + "/history_log.csv", { dialect: "csv" });
util.importTable(CSVPATH + "/history_log.csv", {"dialect": "csv", "table": "history_log" });

util.exportTable("history_text_old", CSVPATH + "/history_text.csv", { dialect: "csv" });
util.importTable(CSVPATH + "/history_text.csv", {"dialect": "csv", "table": "history_text" });
```

If you receive the message "JavaScript is not supported", your MySQL Shell installation lacks JS support. In that case, install the official MySQL Shell package from Oracle (or build it from source) so that JavaScript mode is enabled.

- Follow post-migration instructions to drop the old tables.

MariaDB/MySQL 8.4+ without mysqlsh

This upgrade method takes more time and should be used only if an upgrade with *mysqlsh* is not possible.

Table upgrade

- Log in to MySQL console as root (recommended) or any user with FILE privileges.
- If performing migration with a running Zabbix server, start MySQL with the local_infile variable enabled.
- Rename old tables and create new tables by running `history_upgrade_prepare.sql`:

```
mysql -uzabbix -p<password> zabbix < /usr/share/zabbix/sql-scripts/mysql/option-patches/history_upgrade_pr
```

Migration with stopped server

`max_execution_time` (in MySQL) or `max_statement_time` (in MariaDB) must be disabled before migrating data to avoid time-out during migration.

For MySQL:

```
SET @@max_execution_time=0;
```

For MariaDB:

```
SET @@max_statement_time=0;
```

```
INSERT IGNORE INTO history SELECT * FROM history_old;
INSERT IGNORE INTO history_uint SELECT * FROM history_uint_old;
INSERT IGNORE INTO history_str SELECT * FROM history_str_old;
INSERT IGNORE INTO history_log SELECT * FROM history_log_old;
INSERT IGNORE INTO history_text SELECT * FROM history_text_old;
```

Follow post-migration instructions to drop the old tables.

Migration with running server

Check for which paths import/export is enabled:

```
mysql> SELECT @@secure_file_priv;
+-----------------------+
| @@secure_file_priv    |
+-----------------------+
| /var/lib/mysql-files/ |
+-----------------------+
```

If *secure_file_priv* value is a path to a directory, export/import will be performed for files in that directory. In this case, edit paths to files in queries accordingly or set the *secure_file_priv* value to an empty string for the upgrade time.

If *secure_file_priv* value is empty, export/import can be performed from any location.

If *secure_file_priv* value is NULL, set it to the path that contains exported table data ('/var/lib/mysql-files/' in the example above).

For more information, see MySQL documentation or MariaDB documentation.

`max_execution_time` (in MySQL) or `max_statement_time` (in MariaDB) must be disabled before exporting data to avoid time-out during export.

For MySQL:

```
SET @@max_execution_time=0;
```

For MariaDB:

```
SET @@max_statement_time=0;
```

```
SELECT * INTO OUTFILE '/var/lib/mysql-files/history.csv' FIELDS TERMINATED BY ',' ESCAPED BY '"' LINES TER
LOAD DATA INFILE '/var/lib/mysql-files/history.csv' IGNORE INTO TABLE history FIELDS TERMINATED BY ',' ESC

SELECT * INTO OUTFILE '/var/lib/mysql-files/history_uint.csv' FIELDS TERMINATED BY ',' ESCAPED BY '"' LINE
LOAD DATA INFILE '/var/lib/mysql-files/history_uint.csv' IGNORE INTO TABLE history_uint FIELDS TERMINATED

SELECT * INTO OUTFILE '/var/lib/mysql-files/history_str.csv' FIELDS TERMINATED BY ',' ESCAPED BY '"' LINES
LOAD DATA INFILE '/var/lib/mysql-files/history_str.csv' IGNORE INTO TABLE history_str FIELDS TERMINATED BY

SELECT * INTO OUTFILE '/var/lib/mysql-files/history_log.csv' FIELDS TERMINATED BY ',' ESCAPED BY '"' LINES
LOAD DATA INFILE '/var/lib/mysql-files/history_log.csv' IGNORE INTO TABLE history_log FIELDS TERMINATED BY

SELECT * INTO OUTFILE '/var/lib/mysql-files/history_text.csv' FIELDS TERMINATED BY ',' ESCAPED BY '"' LINE
LOAD DATA INFILE '/var/lib/mysql-files/history_text.csv' IGNORE INTO TABLE history_text FIELDS TERMINATED
```

Follow post-migration instructions to drop the old tables.

PostgreSQL

Export and import must be performed in tmux/screen to ensure that the session isn't dropped. For installations with TimescaleDB, skip this section and proceed to PostgreSQL + TimescaleDB.

See also: Important notes

Table upgrade

- Rename tables using `history_upgrade_prepare.sql`:

```
sudo -u zabbix psql zabbix < /usr/share/zabbix/sql-scripts/postgresql/option-patches/history_upgrade_prepa
```

Migration with stopped server

- Export current history, import it to the temp table, then insert the data into new tables while ignoring duplicates:

```
INSERT INTO history SELECT * FROM history_old ON CONFLICT (itemid,clock,ns) DO NOTHING;

INSERT INTO history_uint SELECT * FROM history_uint_old ON CONFLICT (itemid,clock,ns) DO NOTHING;

INSERT INTO history_str SELECT * FROM history_str_old ON CONFLICT (itemid,clock,ns) DO NOTHING;

INSERT INTO history_log SELECT * FROM history_log_old ON CONFLICT (itemid,clock,ns) DO NOTHING;

INSERT INTO history_text SELECT * FROM history_text_old ON CONFLICT (itemid,clock,ns) DO NOTHING;
```

See tips for improving INSERT performance: PostgreSQL: Bulk Loading Huge Amounts of Data, Checkpoint Distance and Amount of WAL.

- Follow post-migration instructions to drop the old tables.

Migration with running server

- Export current history, import it to the temp table, then insert the data into new tables while ignoring duplicates:

```
\copy history_old TO '/tmp/history.csv' DELIMITER ',' CSV
CREATE TEMP TABLE temp_history (
    itemid                   bigint                                    NOT NULL,
    clock                    integer         DEFAULT '0'               NOT NULL,
    value                    DOUBLE PRECISION DEFAULT '0.0000'          NOT NULL,
    ns                       integer         DEFAULT '0'               NOT NULL
);
\copy temp_history FROM '/tmp/history.csv' DELIMITER ',' CSV
INSERT INTO history SELECT * FROM temp_history ON CONFLICT (itemid,clock,ns) DO NOTHING;

\copy history_uint_old TO '/tmp/history_uint.csv' DELIMITER ',' CSV
CREATE TEMP TABLE temp_history_uint (
    itemid                   bigint                                    NOT NULL,
    clock                    integer         DEFAULT '0'               NOT NULL,
    value                    numeric(20)     DEFAULT '0'               NOT NULL,
    ns                       integer         DEFAULT '0'               NOT NULL
);
\copy temp_history_uint FROM '/tmp/history_uint.csv' DELIMITER ',' CSV
INSERT INTO history_uint SELECT * FROM temp_history_uint ON CONFLICT (itemid,clock,ns) DO NOTHING;

\copy history_str_old TO '/tmp/history_str.csv' DELIMITER ',' CSV
CREATE TEMP TABLE temp_history_str (
    itemid                   bigint                                    NOT NULL,
    clock                    integer         DEFAULT '0'               NOT NULL,
    value                    varchar(255)    DEFAULT ''                NOT NULL,
    ns                       integer         DEFAULT '0'               NOT NULL
);
\copy temp_history_str FROM '/tmp/history_str.csv' DELIMITER ',' CSV
INSERT INTO history_str (itemid,clock,value,ns) SELECT * FROM temp_history_str ON CONFLICT (itemid,clock,n

\copy history_log_old TO '/tmp/history_log.csv' DELIMITER ',' CSV
CREATE TEMP TABLE temp_history_log (
    itemid                   bigint                                    NOT NULL,
    clock                    integer         DEFAULT '0'               NOT NULL,
    timestamp                integer         DEFAULT '0'               NOT NULL,
    source                   varchar(64)     DEFAULT ''                NOT NULL,
    severity                 integer         DEFAULT '0'               NOT NULL,
    value                    text            DEFAULT ''                NOT NULL,
    logeventid               integer         DEFAULT '0'               NOT NULL,
    ns                       integer         DEFAULT '0'               NOT NULL
);
\copy temp_history_log FROM '/tmp/history_log.csv' DELIMITER ',' CSV
INSERT INTO history_log SELECT * FROM temp_history_log ON CONFLICT (itemid,clock,ns) DO NOTHING;

\copy history_text_old TO '/tmp/history_text.csv' DELIMITER ',' CSV
CREATE TEMP TABLE temp_history_text (
```

```
    itemid                  bigint                                  NOT NULL,
    clock                   integer         DEFAULT '0'             NOT NULL,
    value                   text            DEFAULT ''              NOT NULL,
    ns                      integer         DEFAULT '0'             NOT NULL
);
\copy temp_history_text FROM '/tmp/history_text.csv' DELIMITER ',' CSV
INSERT INTO history_text SELECT * FROM temp_history_text ON CONFLICT (itemid,clock,ns) DO NOTHING;
```

- Follow post-migration instructions to drop the old tables.

PostgreSQL + TimescaleDB

Export and import must be performed in tmux/screen to ensure that the session isn't dropped. Zabbix server should be down during the upgrade.

See also: Important notes

- Rename tables using `history_upgrade_prepare.sql`.
  - If compression is enabled (on default installation), run the script from `/usr/share/zabbix/sql-scripts/postgresql/times`
    ```
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/with-compression/history_
    ```
  - If compression is disabled, run the script from `/usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-pat`
    ```
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/without-compression/histo
    ```
- Run TimescaleDB hypertable migration scripts based on compression settings:
  - If compression is enabled (on default installation), run scripts from `/usr/share/zabbix/sql-scripts/postgresql/timesca`
    ```
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/with-compression/history_
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/with-compression/history_
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/with-compression/history_
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/with-compression/history_
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/with-compression/history_
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/with-compression/trends_u
    ```
  - If compression is disabled, run scripts from `/usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patch`
    ```
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/without-compression/histo
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/without-compression/histo
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/without-compression/histo
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/without-compression/histo
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/without-compression/histo
    cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/option-patches/without-compression/trend
    ```

See also: Tips for improving INSERT performance.

- Follow post-migration instructions to drop the old tables.

Post-migration

For all databases, once the migration is completed, do the following:

- Verify that everything works as expected.
- Drop old tables:

```
DROP TABLE history_old;
DROP TABLE history_uint_old;
DROP TABLE history_str_old;
DROP TABLE history_log_old;
DROP TABLE history_text_old;
```

- For TimescaleDB, also drop the following old table:

```
DROP TABLE trends_old;
```

See also

- Preparing auditlog table for partitioning

Preparing auditlog table for partitioning

Overview

Some databases (for example, MySQL) require the partitioning column to be part of the table's unique constraint. Therefore, to partition the `auditlog` table by time, the primary key must be changed from `auditid` to a composite key `auditid` + `clock`.

This section provides instructions for altering the primary key of the `auditlog` table.

Attention:
The instructions provided on this page are designed for advanced users. Note that these instructions might need to be adjusted for your specific configuration. Altering the primary key can also be incompatible with future upgrade patches, so manually handling future upgrades may be necessary.

 Altering the primary key can be a resource-intensive operation that takes a lot of time depending on the `auditlog` table size. Stopping Zabbix server and switching Zabbix frontend to maintenance mode for the time of the alteration is recommended. However, if absolutely necessary, there is a way to alter the primary key without downtime (see below).

Partitioning the `auditlog` table can improve, for example, housekeeping in large setups. Although Zabbix housekeeping currently cannot take advantage of partitioned tables (except for TimescaleDB), you can disable Zabbix housekeeping and delete partitions using scripts.

Since Zabbix 7.0, the `auditlog` table for TimescaleDB has been converted to a hypertable, which allows the housekeeper to drop data by chunks. To upgrade the existing `auditlog` table to a hypertable, see Upgrading TimescaleDB schema.

MySQL

Important notes on rebuilding indexes

MySQL automatically rebuilds indexes for the primary key during the `ALTER TABLE` operation. However, it is highly recommended to also manually rebuild indexes with the `OPTIMIZE TABLE` statement to ensure optimal database performance.

Rebuilding indexes may temporarily require as much additional disk space as the table itself uses. To obtain the current size of data and indexes, you can execute the following statements:

```
ANALYZE TABLE auditlog;
SHOW TABLE STATUS LIKE 'auditlog';
```

If the available disk space is a concern, follow the Altering primary key without downtime instructions. Other options are also available:

- Increasing the `sort_buffer_size` MySQL parameter may help to reduce disk space usage when manually rebuilding indexes. However, modifying this variable may impact overall memory usage of the database.
- Consider freeing up space by deleting potentially unnecessary data.
- Consider decreasing the *Data storage period* housekeeper parameter before executing the housekeeper.

Altering primary key with downtime

1. Drop the current `auditlog` table primary key and add the new primary key.

```
ALTER TABLE auditlog DROP PRIMARY KEY, ADD PRIMARY KEY (auditid, clock);
```

2. Rebuild indexes (optional but highly recommended, see Important notes on rebuilding indexes).

```
OPTIMIZE TABLE auditlog;
```

Altering primary key without downtime

Manual method of altering the primary key is described here. Alternatively, you can use the pt-online-schema-change toolkit from Percona. This toolkit performs the following actions automatically, while also minimizing the space used for altering the `auditlog` table.

1. Create a new table with the new primary key and create indexes.

```
CREATE TABLE `auditlog_new` (
  `auditid`            varchar(25)                         NOT NULL,
  `userid`             bigint unsigned                     NULL,
  `username`           varchar(100)    DEFAULT ''          NOT NULL,
  `clock`              integer         DEFAULT '0'         NOT NULL,
  `ip`                 varchar(39)     DEFAULT ''          NOT NULL,
  `action`             integer         DEFAULT '0'         NOT NULL,
  `resourcetype`       integer         DEFAULT '0'         NOT NULL,
  `resourceid`         bigint unsigned                     NULL,
  `resource_cuid`      varchar(25)                         NULL,
```

```
  `resourcename`       varchar(255)    DEFAULT ''                NOT NULL,
  `recordsetid`        varchar(25)                               NOT NULL,
  `details`            longtext                                  NOT NULL,
  PRIMARY KEY (auditid,clock)
) ENGINE=InnoDB;
CREATE INDEX `auditlog_1` ON `auditlog_new` (`userid`,`clock`);
CREATE INDEX `auditlog_2` ON `auditlog_new` (`clock`);
CREATE INDEX `auditlog_3` ON `auditlog_new` (`resourcetype`,`resourceid`);
```

2. Swap tables.

```
RENAME TABLE auditlog TO auditlog_old, auditlog_new TO auditlog;
```

3. Copy data from the old table to the new table.

```
INSERT INTO auditlog SELECT * FROM auditlog_old;
```

This can be done in chunks (multiple `INSERT INTO` statements with `WHERE clock` clauses as needed) to avoid excessive resource usage.

4. Drop the old table.

```
DROP TABLE auditlog_old;
```

PostgreSQL

Important notes on rebuilding indexes

PostgreSQL automatically rebuilds indexes for the primary key during the `ALTER TABLE` operation. However, it is highly recommended to also manually rebuild indexes with the `REINDEX TABLE CONCURRENTLY` statement to ensure optimal database performance.

Rebuilding indexes may temporarily require up to three times of disk space currently used by indexes. To obtain the current size of indexes, you can execute the following query:

```
SELECT pg_size_pretty(pg_indexes_size('auditlog'));
```

If the available disk space is a concern, follow the Altering primary key without downtime instructions. Other options are also available:

- Increasing the `maintenance_work_mem` PostgreSQL parameter may help to reduce disk space usage when manually rebuilding indexes. However, modifying this variable may impact overall memory usage of the database.
- If you have another disk or tablespace with more available space, you might consider changing the temporary storage location for the index rebuild. You can set the `temp_tablespaces` PostgreSQL parameter to specify a different tablespace for temporary objects.
- Consider freeing up space by deleting potentially unnecessary data.
- Consider decreasing the *Data storage period* housekeeper parameter before executing the housekeeper.

Altering primary key with downtime

1. Drop the current `auditlog` table primary key and add the new primary key.

```
ALTER TABLE auditlog DROP CONSTRAINT auditlog_pkey;
ALTER TABLE auditlog ADD PRIMARY KEY (auditid,clock);
```

2. Rebuild indexes (optional but highly recommended, see Important notes on rebuilding indexes).

```
REINDEX TABLE CONCURRENTLY auditlog;
```

Altering primary key without downtime

Manual method of altering the primary key is described here. Alternatively, the `pg_repack` extension can be considered for creating a new table, copying data, and swapping tables.

1. Create a new table with the new primary key and create indexes.

```
CREATE TABLE auditlog_new (
  auditid              varchar(25)                               NOT NULL,
  userid               bigint                                    NULL,
  username             varchar(100)    DEFAULT ''                NOT NULL,
  clock                integer         DEFAULT '0'               NOT NULL,
  ip                   varchar(39)     DEFAULT ''                NOT NULL,
```

```
  action                integer         DEFAULT '0'               NOT NULL,
  resourcetype          integer         DEFAULT '0'               NOT NULL,
  resourceid            bigint                                    NULL,
  resource_cuid         varchar(25)                               NULL,
  resourcename          varchar(255)    DEFAULT ''                NOT NULL,
  recordsetid           varchar(25)                               NOT NULL,
  details               text            DEFAULT ''                NOT NULL,
  PRIMARY KEY (auditid,clock)
);
CREATE INDEX auditlog_new_1 ON auditlog_new (userid,clock);
CREATE INDEX auditlog_new_2 ON auditlog_new (clock);
CREATE INDEX auditlog_new_3 ON auditlog_new (resourcetype,resourceid);
```

2. Swap tables.

```
ALTER TABLE auditlog RENAME TO auditlog_old;
ALTER TABLE auditlog_new RENAME TO auditlog;
```

3. Copy data from the old table to the new table.

```
INSERT INTO auditlog SELECT * FROM auditlog_old;
```

This can be done in chunks (multiple `INSERT INTO` statements with `WHERE clock` clauses as needed) to avoid excessive resource usage.

4. Drop the old table.

```
DROP TABLE auditlog_old;
```

See also

- Database upgrade to primary keys

TimescaleDB setup

Overview

Zabbix supports TimescaleDB, a PostgreSQL-based database solution of automatically partitioning data into time-based chunks to support faster performance at scale.

> **Warning:**
> Currently, TimescaleDB is not supported by Zabbix proxy.

Instructions on this page can be used for the following scenarios:

- Creating a TimescaleDB database or migrating from existing PostgreSQL tables to TimescaleDB (see Configuration).
- Upgrading an existing TimescaleDB database schema when upgrading Zabbix (see Upgrading TimescaleDB schema).

Configuration

Pre-requisites: TimescaleDB extension of a supported version installed on the database server. For installation instructions, see TimescaleDB documentation.

> **Warning:**
> Before installing TimescaleDB, install a supported PostgreSQL release from the official PostgreSQL repository.

Enable TimescaleDB extension for the specific database by executing:

```
echo "CREATE EXTENSION IF NOT EXISTS timescaledb CASCADE;" | sudo -u postgres psql zabbix
```

Running this command requires database administrator privileges.

> **Note:**
> If you use a database schema other than 'public' you need to add a SCHEMA clause to the command above. E.g.:
 echo "CREATE EXTENSION IF NOT EXISTS timescaledb SCHEMA yourschema CASCADE;" | sudo -u postgres psql zabbix

Then run the `postgresql/timescaledb/schema.sql` script. For new installations, the script must be run after the regular PostgreSQL database has been created with initial schema/data (see database creation).

```
cat /usr/share/zabbix/sql-scripts/postgresql/timescaledb/schema.sql | sudo -u zabbix psql zabbix
```

Attention:
Please ignore warning messages informing that the best practices are not followed while running `schema.sql` script on TimescaleDB version 2.9.0 and higher. Regardless of this warning, the configuration will be completed successfully.

The migration of existing history, trends and audit log data may take a lot of time. Zabbix server and frontend must be down for the period of migration.

The `schema.sql` script sets the following housekeeping parameters:

- Override item history period
- Override item trend period

In order to use partitioned housekeeping for history and trends, both these options must be enabled. It is also possible to enable override individually either for history only or trends only.

For PostgreSQL and TimescaleDB, the `postgresql/timescaledb/schema.sql` script sets two additional parameters:

- Enable compression
- Compress records older than 7 days

To successfully remove compressed data by housekeeper, both *Override item history period* and *Override item trend period* options must be enabled. If override is disabled and tables have compressed chunks, the housekeeper will not remove data from these tables, and warnings about incorrect configuration will be displayed in the *Housekeeping* and *System information* sections.

All of these parameters can be changed in *Administration* > *Housekeeping* after the installation.

Note:
You may want to run the timescaledb-tune tool provided by TimescaleDB to optimize PostgreSQL configuration parameters in your `postgresql.conf`.

Upgrading TimescaleDB schema

When upgrading Zabbix to a version that contains new TimescaleDB hypertables, Zabbix server does not automatically configure those hypertables (for example, when upgrading from Zabbix 6.4 to 8.2, as versions 7.0.0 and 7.0.2 have introduced new hypertables).

To configure new TimescaleDB hypertables, follow these steps:

1. Start Zabbix server; this upgrades the existing database.
2. Check the server log file that the database upgrade is completed; when completed, stop Zabbix server. Note that the server logs a warning if it attempts to enable compression for a table that is not a hypertable.
3. Run the `postgresql/timescaledb/schema.sql` script; this configures the new TimescaleDB hypertables. Note that, since Zabbix 7.0.0, the script's location and name have changed from `postgresql/timescaledb.sql` to `postgresql/timescaledb/schema.sql`.

Attention:
Please ignore warning messages informing that the best practices are not followed while running `schema.sql` script on TimescaleDB version 2.9.0 and higher. Regardless of this warning, the configuration will be completed successfully.

TimescaleDB compression

Native TimescaleDB compression is supported for all Zabbix tables that are TimescaleDB hypertables. During the upgrade or migration to TimescaleDB, initial compression of the large tables may take a lot of time.

Note that compression is supported under the "timescale" Timescale Community license and it is not supported under "apache" Apache 2.0 license. If Zabbix detects that compression is not supported a warning message is written into the Zabbix server log and users cannot enable compression in the frontend.

Note:
Users are encouraged to get familiar with compression in TimescaleDB documentation before using compression.

Note that there are certain limitations imposed by compression, specifically:

- Compressed chunk modifications (inserts, deletes, updates) are not allowed
- Schema changes for compressed tables are not allowed.

Compression settings can be changed in the *History, trends and audit log compression* block in *Administration* > *Housekeeping* section of Zabbix frontend.

Parameter	Default	Comments
Enable compression	Enabled	Checking or unchecking the checkbox does not activate/deactivate compression immediately. Because compression is handled by the Housekeeper, the changes will take effect in up to 2 times `HousekeepingFrequency` hours (set in `zabbix_server.conf`) After disabling compression, new chunks that fall into the compression period will not be compressed. However, all previously compressed data will stay compressed. To uncompress previously compressed chunks, follow the instructions in TimescaleDB documentation. When upgrading from older versions of Zabbix with TimescaleDB support, compression will not be enabled by default.
Compress records older than	7d	This parameter cannot be less than 7 days. Due to immutability of compressed chunks all late data (e.g., data delayed by a proxy) that is older than this value will be discarded.

For better trend update performance, you may want to lower the "chunk_time_interval" for `trends` and `trends_uint` tables from 30 days to 7 days or less depending on how many items use trends. The purpose of this setting is to adhere to TimescaleDB's best practices and to ensure that the chunk size remains within the system's available resources.

Elasticsearch setup

Zabbix can store history data in Elasticsearch as an alternative to a relational database.

Attention:
The support of Elasticsearch is currently experimental.

This guide covers the setup for the supported versions of Elasticsearch. If you're using a different version, some functionality may not work as intended.

Elasticsearch can store the following value types:

Item value type	Database table	Elasticsearch type
Numeric (unsigned)	history_uint	uint
Numeric (float)	history	dbl
Character	history_str	str
Log	history_log	log
Text	history_text	text
Binary	history_bin	not supported by Zabbix
JSON	history_json	json

Note:
Elasticsearch does not accept JSON arrays. A JSON value must be either a single object, or a set of objects.

Important notes

- Elasticsearch requires libcurl. See requirements for details.
- The housekeeper does not delete data from Elasticsearch. To control how long data is kept, see Index Lifecycle Management.
- Zabbix does **not** calculate or store trends in Elasticsearch. Consider extending the history storage period to preserve older data.
- When Elasticsearch is used, range queries retrieving values from the database are limited by the timestamp of the data storage period.

- Elasticsearch is not supported for Zabbix proxies.

If Elasticsearch isn't installed yet, refer to the official installation guide before proceeding.

Configuring Elasticsearch

To store history data in Elasticsearch, you need to:

- Create an index for each value type you want to store—this is where Elasticsearch stores the data, similar to a table in a relational database.
- Define a mapping for each index—this defines the structure of the data, similar to a table schema.
- Set up an ingest pipeline to process values before storage (required for JSON values and date-based indices).

Elasticsearch can store data in a single index per value type, or across multiple date-based indices. Both approaches are described below.

Storing history in a single index

In this approach, all history data for a given value type is written to a single index (e.g., `uint` or `text`).

To create an index for the *Numeric (unsigned)* value type, send the following request (with `/uint` in the URL) to your Elasticsearch instance:

```
curl -X PUT \
 http://localhost:9200/uint \
 -H 'content-type:application/json' \
 -d '{
     "settings": {
      "index": {
         "number_of_replicas": 1,
         "number_of_shards": 5
      }
   },
   "mappings": {
      "properties": {
         "itemid": { "type": "long" },
         "clock": { "format": "epoch_second", "type": "date" },
         "value": { "type": "long" }
      }
   }
}'
```

Elasticsearch will respond with a confirmation that the index was created:

```
{"acknowledged": true, "shards_acknowledged": true, "index": "uint"}
```

Similar requests must be sent for each additional value type you want to store in Elasticsearch.

> **Note:**
> Mappings for all value types are available in the Zabbix source repository.

For example, to create an index for the *Text* value type:

```
curl -X PUT \
 http://localhost:9200/text \
 -H 'content-type:application/json' \
 -d '{
   "settings": {
      "index": {
         "number_of_replicas": 1,
         "number_of_shards": 5
      }
   },
   "mappings": {
      "properties": {
         "itemid": { "type": "long" },
         "clock": { "format": "epoch_second", "type": "date" },
         "value": {
            "fields": {
```

```
               "analyzed": { "index": true, "type": "text", "analyzer": "standard" }
            },
            "index": false,
            "type": "text"
         }
      }
   }
}'
```

JSON value type

Unlike other value types, JSON values require additional processing before storage.

The index below uses separate fields for parsed and raw values, so an ingest pipeline is needed to parse each value as JSON and store it in the correct field.

To create an index for the *JSON* value type, send the following request (with `/json` in the URL) to your Elasticsearch instance.

```
curl -X PUT \
 http://localhost:9200/json \
 -H 'content-type:application/json' \
 -d '{
   "settings": {
      "number_of_shards": 5,
      "number_of_replicas": 1
   },
   "mappings": {
      "dynamic": false,
      "properties": {
         "itemid": { "type": "long" },
         "clock": { "type": "date", "format": "epoch_second" },
         "ns": { "type": "long" },
         "value_parsed": { "type": "flattened" },
         "value_raw": { "type": "keyword", "ignore_above": 1000000 }
    }
  }
}'
```

Then, create the ingest pipeline:

```
curl -X PUT \
 http://localhost:9200/_ingest/pipeline/json \
 -H 'content-type:application/json' \
 -d '{
   "processors": [
      {
         "json": {
            "field": "value",
            "target_field": "value_parsed",
            "ignore_failure": true
         }
      },
      {
         "set": {
            "if": "ctx.value_parsed == null",
            "field": "value_raw",
            "value": "{{{ value }}}"
         }
      }
   ],
   "on_failure": [
      {
         "set": {
            "field": "value_raw",
            "value": "{{{ value }}}"
         }
```

```
      }
   ]
}'
```

Elasticsearch will respond with a confirmation that the ingest pipeline was created:

```
{"acknowledged": true}
```

Storing history in date-based indices

Instead of writing all history data to a single index (e.g., `uint`), Elasticsearch can distribute this data across multiple date-based indices (e.g., `uint-2026-01-01`, `uint-2026-01-02`). This makes it easier to manage data volume and retention over time.

To enable this, you need to:

- Create an index template for each value type you want to store—this tells Elasticsearch what mapping to apply when it automatically creates a new date-based index.
- Create an ingest pipeline for each value type—it processes each incoming value and routes it to the correct date-based index.

Index templates

To create a template for the `text` index, send a request with the following details:

- Use `_template/text_template` in the URL of your Elasticsearch instance.
- Use `"text*"` in the `"index_patterns"` field to match the index name.
- Use a mapping for the `text` value type (see mappings in the Zabbix source repository).

```
curl -X PUT \
 http://localhost:9200/_template/text_template \
 -H 'content-type:application/json' \
 -d '{
   "index_patterns": [ "text*" ],
   "settings": {
      "index": {
         "number_of_replicas": 1,
         "number_of_shards": 5
      }
   },
   "mappings": {
      "properties": {
         "itemid": { "type": "long" },
         "clock": { "format": "epoch_second", "type": "date" },
         "value": {
            "fields": {
               "analyzed": { "index": true, "type": "text", "analyzer": "standard" }
            },
            "index": false,
            "type": "text"
         }
      }
   }
}'
```

Template for the `json` index:

```
curl -X PUT \
 http://localhost:9200/_template/json_template \
 -H 'content-type:application/json' \
 -d '{
   "index_patterns": [ "json*" ],
   "settings": {
      "number_of_shards": 5,
      "number_of_replicas": 1
   },
   "mappings": {
      "dynamic": false,
      "properties": {
```

```
         "itemid": { "type": "long" },
         "clock": { "type": "date", "format": "epoch_second" },
         "ns": { "type": "long" },
         "value_parsed": { "type": "flattened" },
         "value_raw": { "type": "keyword", "ignore_above": 1000000 }
      }
   }
}'
```

Ingest pipelines

To create an ingest pipeline for the `text` index:

- Use `_ingest/pipeline/text-pipeline` in the URL of your Elasticsearch instance.
- Include a `date_index_name` processor to route each value to the correct date-based index based on its timestamp.

```
curl -X PUT \
 http://localhost:9200/_ingest/pipeline/text-pipeline \
 -H 'content-type:application/json' \
 -d '{
   "description": "daily text index naming",
   "processors": [
      {
         "date_index_name": {
            "field": "clock",
            "date_formats": ["UNIX"],
            "index_name_prefix": "text-",
            "date_rounding": "d"
         }
      }
   ]
}'
```

For the `json` index, the pipeline also needs to parse the JSON value before routing it to the correct index:

```
curl -X PUT \
 http://localhost:9200/_ingest/pipeline/json-pipeline \
 -H 'content-type:application/json' \
 -d '{
   "description": "daily json index naming",
   "processors": [
      {
         "json": {
            "field": "value",
            "target_field": "value_parsed",
            "ignore_failure": true
         }
      },
      {
         "script": {
            "source": "if (ctx.value_parsed == null || !(ctx.value_parsed instanceof Map)) { ctx.value_raw
         }
      },
      {
         "date_index_name": {
            "field": "clock",
            "date_formats": [ "UNIX" ],
            "index_name_prefix": "json-",
            "date_rounding": "d"
         }
      }
   ]
}'
```

Configuring Zabbix server

In your Zabbix server configuration file (`zabbix_server.conf`), set the `HistoryProvider` parameter.

For example, to store *Character*, *Log*, *Text*, and *JSON* type values in Elasticsearch (while keeping *Numeric* values in a database):

```
HistoryProvider=elasticsearch;value_types="str,log,text,json",url=http://localhost:9200
```

If you're using date-based indices, add `date_index=1` to the parameter:

```
HistoryProvider=elasticsearch;value_types="str,log,text,json",url=http://localhost:9200,date_index=1
```

After making changes, restart Zabbix server:

```
systemctl restart zabbix-server
```

Configuring Zabbix frontend

In your Zabbix frontend configuration file (`zabbix.conf.php`), set the `$HISTORY_PROVIDERS` variable to match the server configuration:

```
$HISTORY_PROVIDERS[] = [
  'types' => ['str','log','text','json'],
  'provider' => 'elasticsearch',
  'url' => 'http://localhost:9200'
];
```

Troubleshooting

The following steps may help you troubleshoot problems with your Elasticsearch setup:

1. Check Elasticsearch or Zabbix server logs for errors.
2. To identify slow queries, use the `log_slow_queries` option in the `HistoryProvider` Zabbix server configuration parameter.
3. Verify that Elasticsearch allows access from Zabbix server and Zabbix frontend.
4. Verify that `auto_create_index` is enabled:

```
curl -X GET \
 "http://localhost:9200/_cluster/settings?include_defaults=true&filter_path=**.auto_create_index"

#### {"defaults": {"action": {"auto_create_index": "false"} } }
```

To enable it, send the following request:

```
curl -X PUT \
 http://localhost:9200/_cluster/settings \
 -H 'content-type:application/json' \
 -d '{
   "persistent": {
      "action.auto_create_index": "true"
   }
}'

#### {"acknowledged": true, "persistent": {"action": {"auto_create_index": "true"} }, "transient": {} }
```

5. Verify that mappings, templates, and ingest pipelines are correct by sending `GET` requests to their respective URLs:

```
curl -X GET http://localhost:9200/json
curl -X GET http://localhost:9200/_template/json*
curl -X GET http://localhost:9200/_ingest/pipeline/json*
```

You can compare received responses with expected responses in the Elasticsearch API documentation.

6. Check if any shards are in a failed state; restarting Elasticsearch may resolve this.
7. Query Elasticsearch to see if the data collected by Zabbix is stored, for example:

```
curl 'http://localhost:9200/json/_search' \
 -H 'Content-Type: application/json' \
 -d '{
   "query": {
     "term": {
```

```
      "itemid": 42269
    }
  }
}'
```

8. If you need to reset your Elasticsearch setup and start over, you can delete all indices, templates, and ingest pipelines:

```
curl -X DELETE "http://localhost:9200/_all"
curl -X DELETE "http://localhost:9200/_template/*"
curl -X DELETE "http://localhost:9200/_ingest/pipeline/*"
```

ClickHouse setup

Zabbix can store history data in ClickHouse as an alternative to a relational database.

Note:
ClickHouse can also be used as a telemetry data source for application performance monitoring (APM); for details, see *Administration* > *Data source*.

This guide covers the setup for the supported versions of ClickHouse. If you're using a different version, some functionality may not work as intended.

ClickHouse can store the following value types:

Item value type	Database table	ClickHouse type
Numeric (unsigned)	history_uint	uint
Numeric (float)	history	dbl
Character	history_str	str
Log	history_log	log
Text	history_text	text
Binary	history_bin	not supported by Zabbix
JSON	history_json	json

Note:
ClickHouse does not accept JSON arrays. A JSON value must be either a single object, or a set of objects. Additionally, ClickHouse handles JSON keys with NULL the same as missing keys.

Important notes

- The housekeeper does not delete data from ClickHouse. To control how long data is kept, configure ClickHouse data storage period in ClickHouse.
- Zabbix does **not** calculate or store trends in ClickHouse. Consider extending the history storage period to preserve older data.
- If you want to migrate your history data from an existing Zabbix database (MySQL or PostgreSQL) to ClickHouse, see the ClickHouse schema and history migration scripts.
- ClickHouse is not supported as history storage for Zabbix proxies, but can be used as a telemetry data source for Zabbix proxy with APM enabled.
- If you have a ClickHouse cluster, see Configuring ClickHouse clusters.

Configuring ClickHouse

You need to create and configure a Zabbix database and user, and import the database schema.

This guide gives instructions for Docker or package installations of ClickHouse.

Docker

1. Create and configure the Zabbix database and user when running the ClickHouse container:

```
sudo docker run -d \
  --name clickhouse \
  -e CLICKHOUSE_DB=zabbix \
  -e CLICKHOUSE_USER=zabbix \
  -e CLICKHOUSE_DEFAULT_ACCESS_MANAGEMENT=1 \
```

```
  -e CLICKHOUSE_PASSWORD=<password> \
  -p 8123:8123/tcp \
  -p 9000:9000/tcp \
  --cap-add=SYS_NICE \
  --cap-add=NET_ADMIN \
  --cap-add=IPC_LOCK \
  --ulimit nofile=262144:262144 \
  -v clickhouse_data:/var/lib/clickhouse \
  clickhouse/clickhouse-server:26.4
```

2. Confirm that ClickHouse is working and you can connect to it:

```
sudo docker exec -it clickhouse \
  clickhouse-client \
  --query "SELECT version()"

#### 26.4.4.38
```

3. Import the database schema using the history_all.sh script from your Zabbix directory:

```
./usr/share/zabbix/sql-scripts/clickhouse/history_all.sh \
  --user zabbix \
  --password <password> \
  --db zabbix \
  --server http://localhost:8123
```

Note:
The data storage period (Time-To-Live, or TTL) for ClickHouse is, by default, 31 days. To change it, use the `--ttl` option when importing the database schema (e.g., `--ttl 604800` for 7 days), or configure it later.

4. Verify that tables were created:

```
sudo docker exec -it clickhouse \
  clickhouse-client \
  --user zabbix \
  --password <password> \
  --query "SHOW TABLES FROM zabbix"

#### history
#### history_json
#### history_log
#### history_str
#### history_text
#### history_uint
```

Packages

1. Start ClickHouse:

```
sudo service clickhouse-server start
```

2. Create and configure the Zabbix database and user:

```
clickhouse-client --user default

localhost :) CREATE DATABASE IF NOT EXISTS zabbix
localhost :) CREATE USER IF NOT EXISTS zabbix IDENTIFIED WITH sha256_password BY '<password>'
localhost :) GRANT CREATE, ALTER, DROP, INSERT, SELECT, UPDATE, OPTIMIZE ON zabbix.* TO zabbix
localhost :) quit
```

3. Import the database schema using the history_all.sh script from your Zabbix directory:

```
./usr/share/zabbix/sql-scripts/clickhouse/history_all.sh \
  --user zabbix \
  --password <password> \
  --db zabbix \
  --server http://localhost:8123
```

Note:
The data storage period (Time-To-Live, or TTL) for ClickHouse is, by default, 31 days. To change it, use the `--ttl` option when importing the database schema (e.g., `--ttl 604800` for 7 days), or configure it later.

4. Verify that tables were created:

```
clickhouse-client \
  --user zabbix \
  --password <password> \
  --query "SHOW TABLES FROM zabbix"

#### history
#### history_json
#### history_log
#### history_str
#### history_text
#### history_uint
```

Configuring Zabbix server

In your Zabbix server configuration file (`zabbix_server.conf`), set the `HistoryProvider` parameter.

For example, to store all supported types of values in ClickHouse:

```
HistoryProvider=clickhouse;value_types="uint,dbl,str,log,text,json",url=http://localhost:8123,db=zabbix,us
```

After making changes, restart Zabbix server:

```
systemctl restart zabbix-server
```

Configuring Zabbix frontend

In your Zabbix frontend configuration file (`zabbix.conf.php`), set the `$HISTORY_PROVIDERS` variable to match the server configuration:

```
$HISTORY_PROVIDERS[] = [
  'types' => ['uint','dbl','str','log','text','json'],
  'provider' => 'clickhouse',
  'url' => 'http://localhost:8123',
  'db' => 'zabbix',
  'username' => 'zabbix',
  'password' => '<password>'
];
```

Additional configuration

The steps below are optional. You do not need them for a basic setup.

Configuring ClickHouse data storage period

The data storage period (Time-To-Live, or TTL) for ClickHouse is, by default, 31 days. To change it, run the commands below.

The examples below use Docker. If you installed ClickHouse using packages, run the queries in the ClickHouse client directly.

1. Alter the table (replace `history_json` and 3600 with the values you need):

```
sudo docker exec -it clickhouse \
  clickhouse-client \
  --user zabbix \
  --password <password> \
  --query "ALTER TABLE zabbix.history_json MODIFY TTL clock_ns + toIntervalSecond(3600)"
```

2. Apply the changes immediately:

```
sudo docker exec -it clickhouse \
  clickhouse-client \
  --user zabbix \
  --password <password> \
  --query "OPTIMIZE TABLE zabbix.history_json MATERIALIZE TTL"
```

3. Verify that the data storage period changed:

```
sudo docker exec -it clickhouse \
  clickhouse-client \
  --user zabbix \
  --password <password> \
  --query "SHOW CREATE TABLE zabbix.history_json"

#### CREATE TABLE zabbix.history_json
#### (
####     `itemid` UInt64,
####     `clock_ns` DateTime64(9),
####     `value` JSON,
####     `value_str` String
#### )
#### ENGINE = MergeTree
#### PARTITION BY toDate(clock_ns)
#### PRIMARY KEY (itemid, clock_ns)
#### ORDER BY (itemid, clock_ns)
#### TTL clock_ns + toIntervalSecond(3600)
#### SETTINGS index_granularity = 8192
```

4. Restart Zabbix server to refresh the data storage period in *Administration* > *Housekeeping*:

```
systemctl restart zabbix-server
```

Configuring ClickHouse clusters

This section provides configuration steps for ClickHouse clusters with replication and sharding.

The steps below use Docker and the ClickHouse `cluster_2S_2R` from ClickHouse documentation (running ClickHouse 26.4.4.38 with ClickHouse Keeper).

Before you start, it is recommended to put a load balancer in front of your ClickHouse cluster for Zabbix. Zabbix server and frontend configuration supports only a single ClickHouse URL, so it can only point to one node. If that node goes down, a load balancer will redirect Zabbix to a healthy node instead, through that same URL.

1. Create the Zabbix database:

```
sudo docker exec -it clickhouse-01 \
  clickhouse-client \
  --user default \
  --query "CREATE DATABASE IF NOT EXISTS zabbix ON CLUSTER cluster_2S_2R ENGINE = Replicated('/clickhouse/
```

Attention:
The `ON CLUSTER` clause reaches only nodes that are part of the cluster when you run the command. If you add nodes to your cluster later, run this command again.

2. Create and configure the Zabbix user:

```
sudo docker exec -it clickhouse-01 \
  clickhouse-client \
  --user default \
  --query "CREATE USER IF NOT EXISTS zabbix ON CLUSTER cluster_2S_2R IDENTIFIED WITH sha256_password BY '<

sudo docker exec -it clickhouse-01 \
  clickhouse-client \
  --user default \
  --query "GRANT ON CLUSTER cluster_2S_2R CREATE, ALTER, DROP, INSERT, SELECT, UPDATE, OPTIMIZE ON zabbix.
```

3. Import the database schema using the `history_all.sh` script from your Zabbix directory:

```
./usr/share/zabbix/sql-scripts/clickhouse/history_all.sh \
  --user zabbix \
  --password <password> \
  --db zabbix \
  --server http://localhost:8123 \
  --engine "ReplicatedMergeTree()"
```

Note:
The data storage period (Time-To-Live, or TTL) for ClickHouse is, by default, 31 days. To change it, use the `--ttl` option when importing the database schema (e.g., `--ttl 604800` for 7 days), or configure it later.

Note:
If your ClickHouse cluster requires explicit replica paths (for example, if your ClickHouse server configuration does not set default paths for replicated tables), run the schema scripts (`history_schema.sh`, `history_uint_schema.sh`, etc.) individually and set a different path for each table:

```
./usr/share/zabbix/sql-scripts/clickhouse/history_schema.sh \
  --user zabbix \
  --password <password> \
  --db zabbix \
  --server http://localhost:8123 \
  --engine "ReplicatedMergeTree('/clickhouse/tables/{shard}/history', '{replica}')"
```

For more details on ClickHouse replication, see Replicating data and Replicated* table engines in ClickHouse documentation.

4. Configure sharding for each table type. Rename the tables and create a `Distributed` table with the original name, pointing to the renamed table. This `Distributed` table is what Zabbix actually reads from and writes to; the table splits data across shards automatically. For example, for the `history` table:

```
sudo docker exec -it clickhouse-01 \
  clickhouse-client \
  --user default \
  --query "RENAME TABLE zabbix.history TO zabbix.history_local"

sudo docker exec -it clickhouse-01 \
  clickhouse-client \
  --user default \
  --query "CREATE TABLE zabbix.history AS zabbix.history_local ENGINE = Distributed('cluster_2S_2R', 'zabb
```

Note:
For more details on ClickHouse sharding, see Distributed table engine in ClickHouse documentation.

5. Configure Zabbix server.

6. Configure Zabbix frontend.

7. Verify that Zabbix data is spread across shards. The row count from `zabbix.history_local` may change between repeated checks, since Zabbix server continuously inserts new data. Run this command on each node:

```
sudo docker exec -it clickhouse-01 \
  clickhouse-client \
  --user default \
  --query "SELECT count(*) FROM zabbix.history_local"

#### 2446 (clickhouse-01)
#### 2471 (clickhouse-02)
#### 2494 (clickhouse-03)
```

8. Check that replication is in sync (`active_replicas` should match `total_replicas`):

```
sudo docker exec -it clickhouse-01 \
  clickhouse-client \
  --user default \
  --query "SELECT database, table, is_leader, total_replicas, active_replicas FROM system.replicas WHERE d

#### zabbix  history_json_local  1 2 2
#### zabbix  history_local       1 2 2
#### zabbix  history_log_local   1 2 2
#### zabbix  history_str_local   1 2 2
#### zabbix  history_text_local  1 2 2
#### zabbix  history_uint_local  1 2 2
```

Troubleshooting

The following steps may help you troubleshoot problems with your ClickHouse setup:

1. Check ClickHouse or Zabbix server logs for errors.
2. To identify slow queries, use the `log_slow_queries` option in the `HistoryProvider` Zabbix server configuration parameter.
3. Verify that ClickHouse allows access from Zabbix server and Zabbix frontend.
4. Query ClickHouse to see if the data collected by Zabbix is stored, for example:

```
sudo docker exec -it clickhouse \
  clickhouse-client \
  --user zabbix \
  --password <password> \
  --query "SELECT * FROM zabbix.history_uint WHERE itemid = 42269"
```

Distribution-specific notes on setting up NGINX for Zabbix

RHEL

On RHEL 8, 9, and 10, NGINX is included in AppStream—no extra repos needed. Simply enable and install with:

```
dnf module list nginx
dnf -y install nginx
```

On RHEL 7, NGINX can be installed from the CentOS 7 Extras repository (if using CentOS) via:

```
yum -y install nginx
```

If you prefer the very latest upstream builds, configure the official NGINX repository at `/etc/yum.repos.d/nginx.repo`:

```
cat > /etc/yum.repos.d/nginx.repo << 'EOF'
[nginx-stable]
name=nginx stable repo
baseurl=https://nginx.org/packages/rhel/$releasever/$basearch/
gpgcheck=1
enabled=1
gpgkey=https://nginx.org/keys/nginx_signing.key
EOF
dnf -y install nginx
```

SLES 15

In SUSE Linux Enterprise Server 15 you need to configure `php-fpm` (the path to configuration file may vary slightly depending on the service pack):

```
cp /etc/php8/fpm/php-fpm.conf{.default,}
cp /etc/php8/fpm/php-fpm.d/www.conf{.default,}
sed -i 's/user = nobody/user = wwwrun/; s/group = nobody/group = www/' /etc/php8/fpm/php-fpm.d/www.conf
```

SAML setup with Microsoft Entra ID

Overview

This section provides guidelines for configuring single sign-on and user provisioning into Zabbix from Microsoft Entra ID (formerly Microsoft Azure Active Directory) using SAML 2.0 authentication.

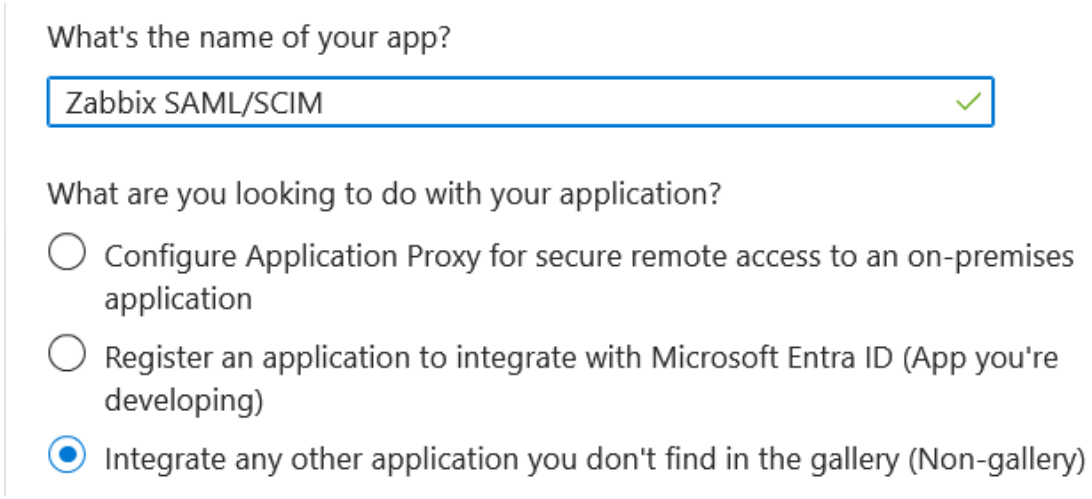
Microsoft Entra ID configuration

Creating application

1. Log into Microsoft Entra admin center at Microsoft Entra ID. For testing purposes, you may create a free trial account in Microsoft Entra ID.

2. In Microsoft Entra admin center select *Applications* > *Enterprise applications* > *New application* > *Create your own application*.

3. Add the name of your app and select the *Integrate any other application...* option. After that, click *Create*.

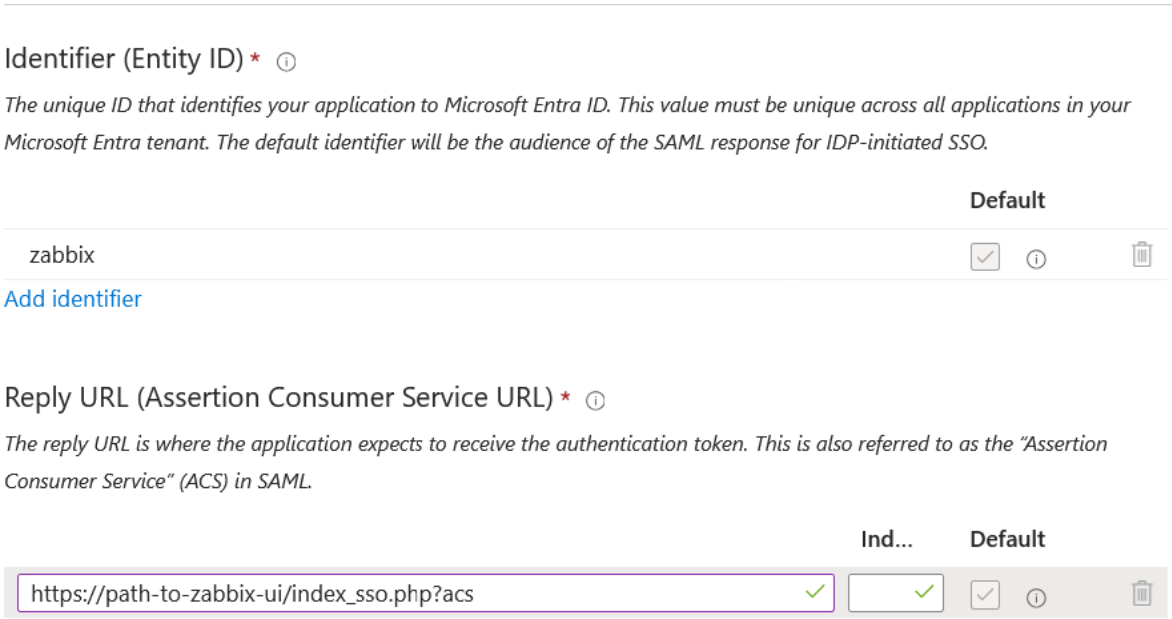


Setting up single sign-on

1. In your application page, go to *Set up single sign on* and click *Get started*. Then select *SAML*.

2. Edit *Basic SAML Configuration*:

- In *Identifier (Entity ID)* set a unique name to identify your app to Microsoft Entra ID, for example, `zabbix`;
- In *Reply URL (Assertion Consumer Service URL)* set the Zabbix single sign-on endpoint: `https://<path-to-zabbix-ui>/index_sso`



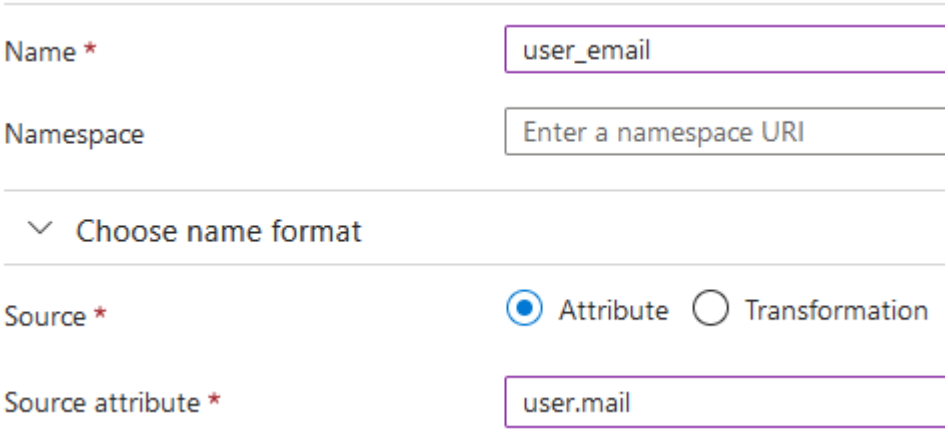
Note that `https` is required. To use this with Zabbix, add the following line to the frontend configuration file:

```
$SSO['SETTINGS'] = ['use_proxy_headers' => true];
```

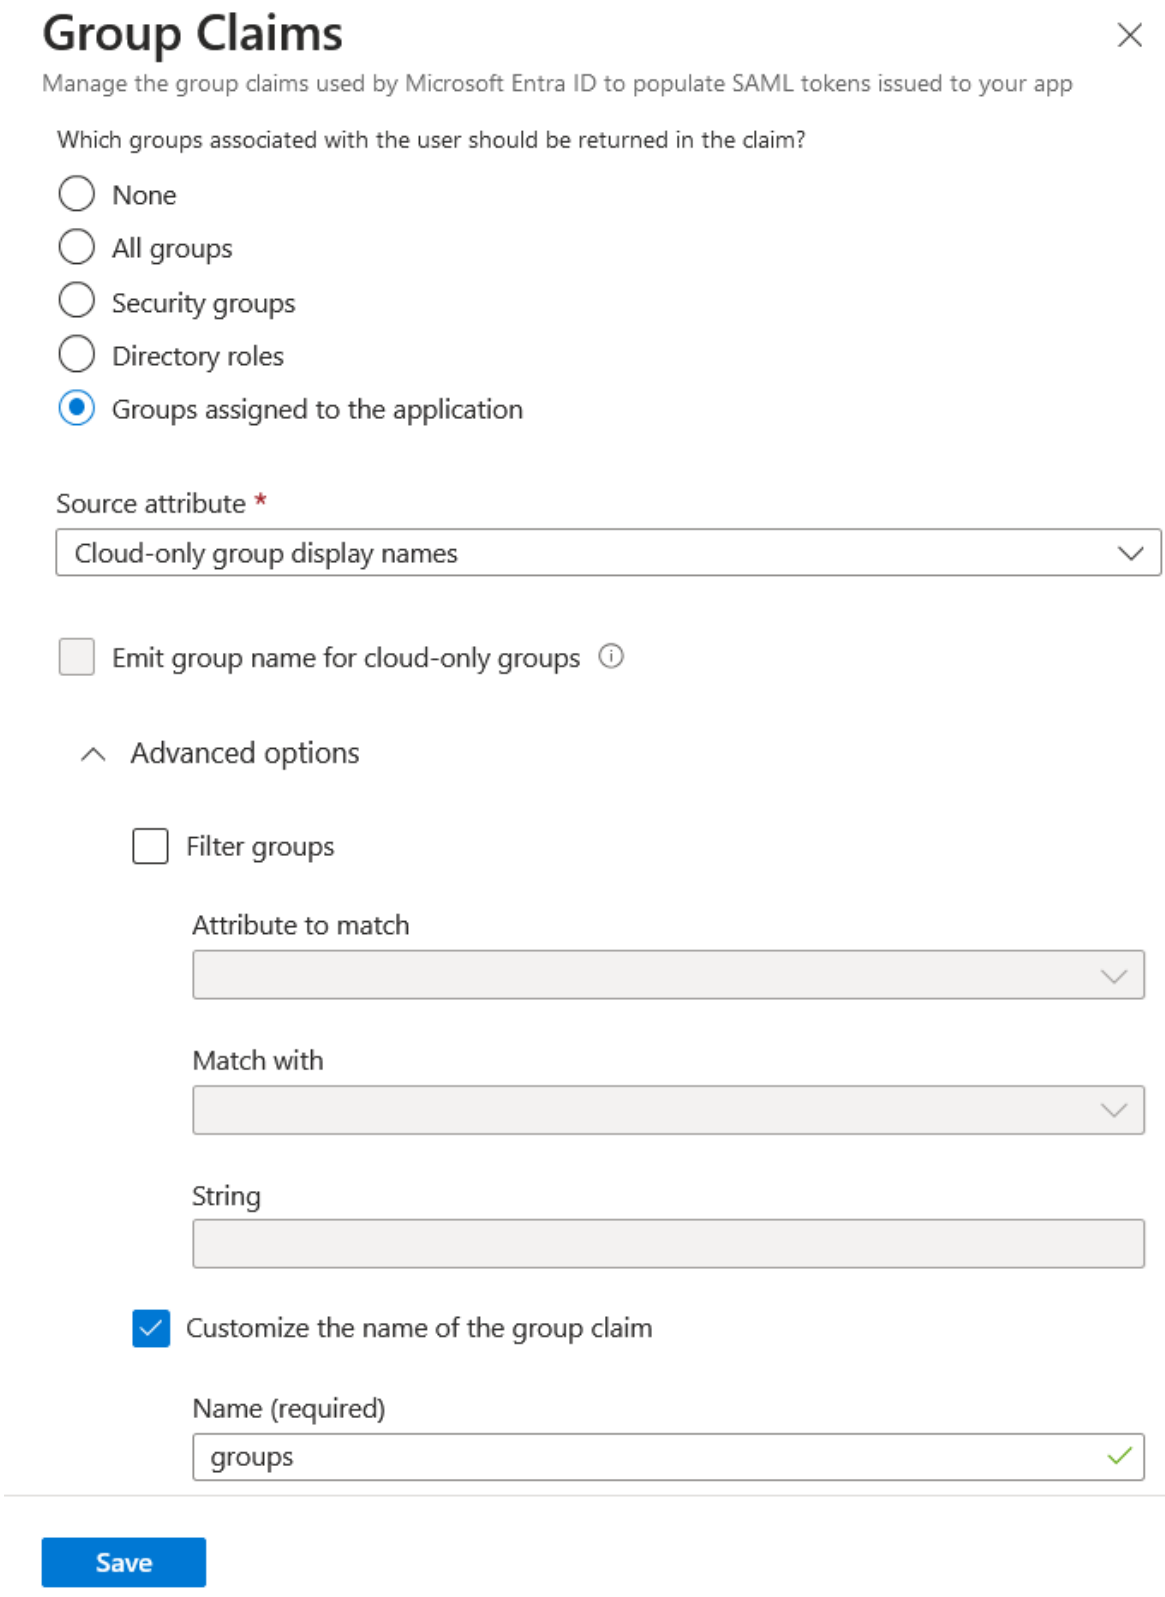
3. Edit *Attributes & Claims*. You must add all attributes that you want to pass to Zabbix (user_name, user_lastname, user_email, user_mobile, groups).

The attribute names are arbitrary. Different attribute names may be used, however, it is required that they match the respective field value in Zabbix SAML settings.

- Click *Add new claim* to add an attribute:



- Click *Add a group claim* to add an attribute for passing groups to Zabbix:



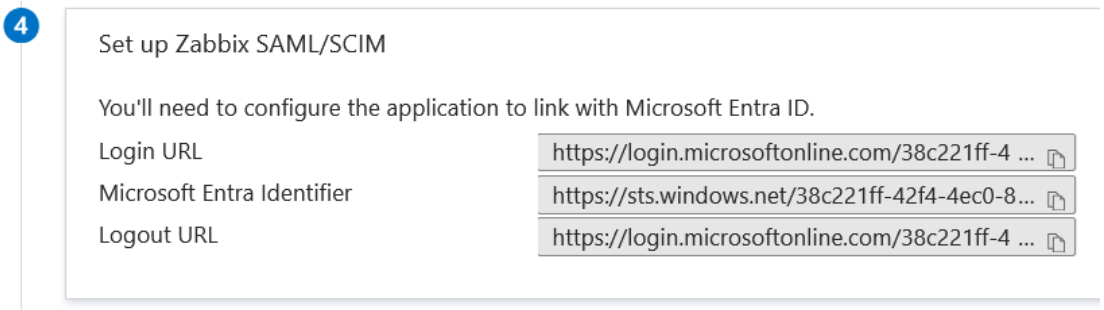
It is important in this claim that the group names (rather than group IDs) are passed to Zabbix by the selected *Source attribute*. Otherwise, JIT user provisioning will not work properly.

4. In *SAML Certificates* download the Base64 certificate provided by Entra ID and place it into `conf/certs` of the Zabbix frontend installation (the location may differ based on OS/source of your Zabbix installation).

Set 644 permissions to it by running:

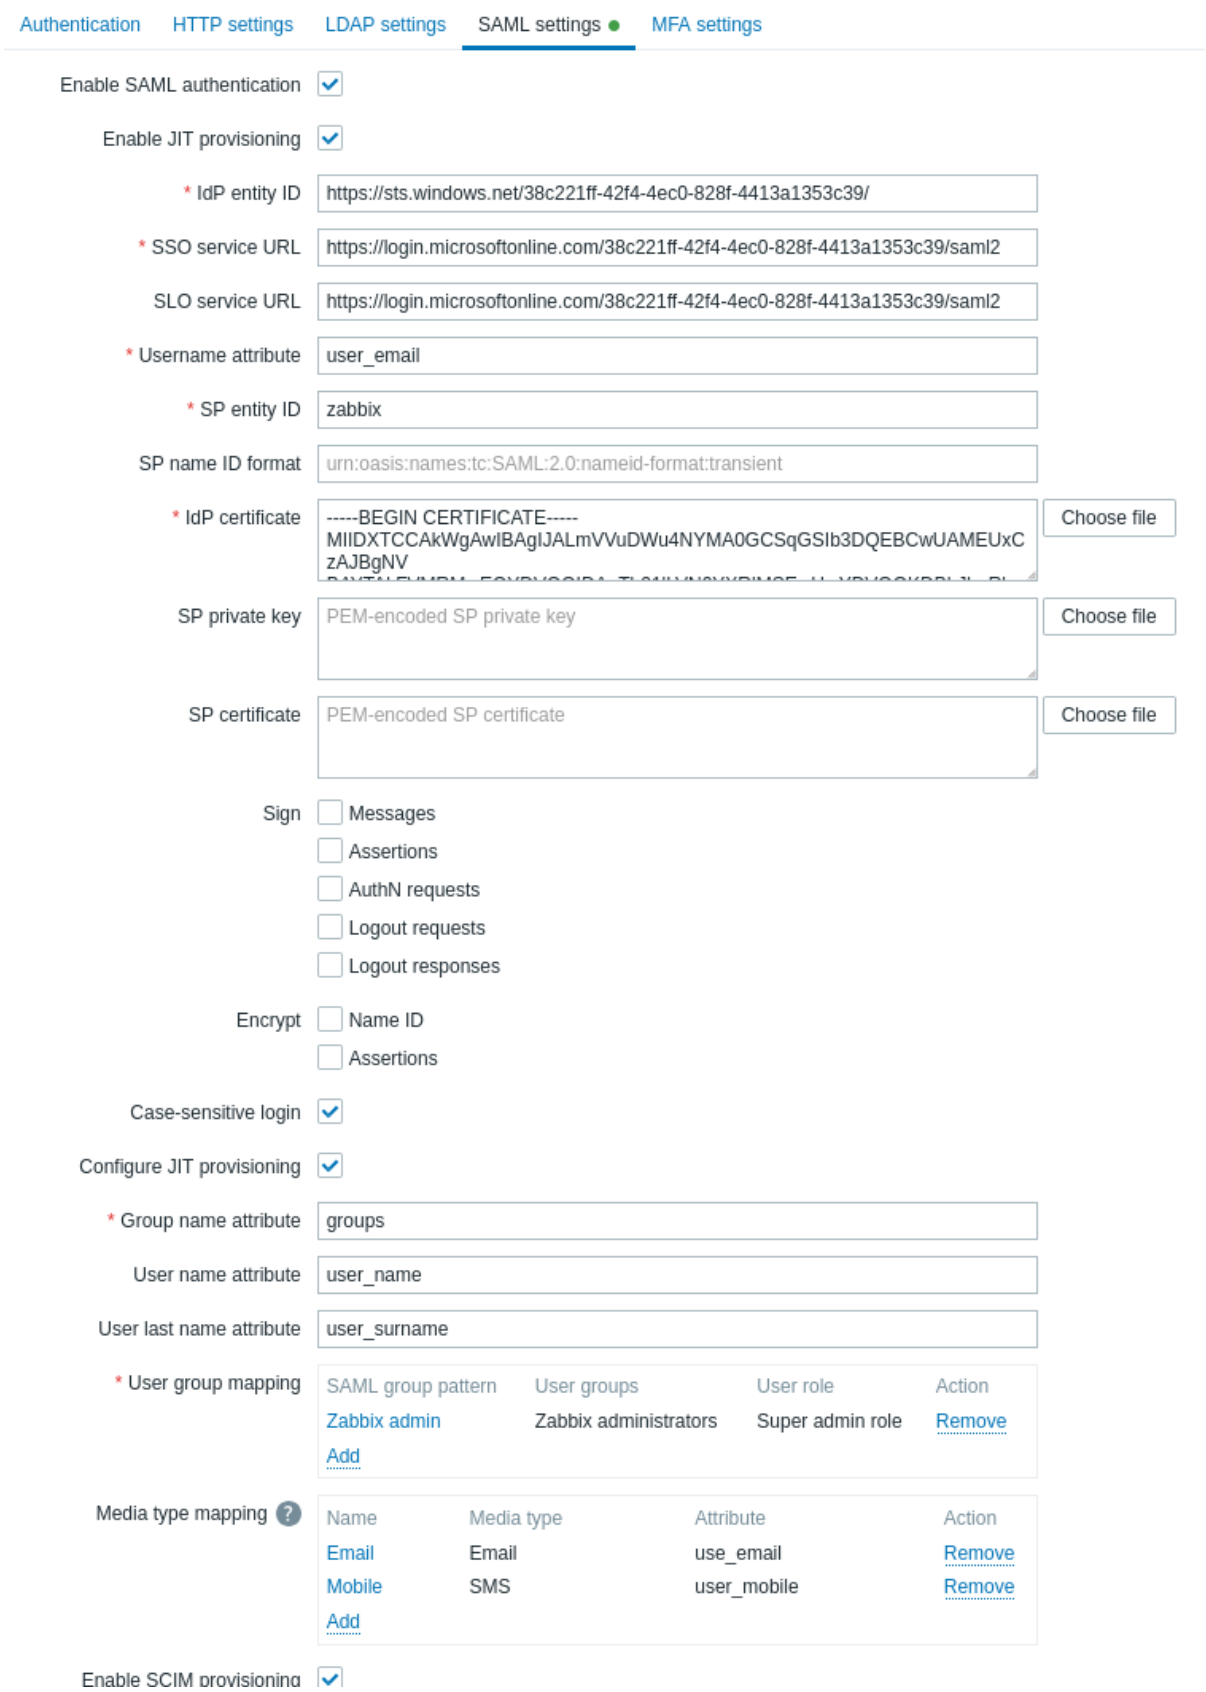
```
chmod 644 entra.cer
```

5. Use the values from *Set up <your app name>* in Entra ID to configure Zabbix SAML authentication (see next section):



Zabbix configuration

1. In Zabbix, go to the SAML settings and fill the configuration options based on the Entra ID configuration:



Zabbix field	Setup field in Entra ID	Sample value
IdP entity ID	Microsoft Entra identifier	
SSO service URL	Login URL	
SLO service URL	Logout URL	
SP entity ID	Identifier (Entity ID)	
Username attribute	Custom attribute (claim)	`user_email`
Group name attribute	Custom attribute (claim)	`groups`
User name attribute	Custom attribute (claim)	`user_name`

Zabbix field	Setup field in Entra ID	Sample value
User last name attribute	Custom attribute (claim)	`user_lastname`

2.Add the Base64 certificate provided by Entra ID.

If `$SSO['CERT_STORAGE'] = 'database'` is set in the frontend configuration file (`zabbix.conf.php`), you can paste the certificate text or upload the certificate file in the frontend during SAML configuration—no files are required on the filesystem.

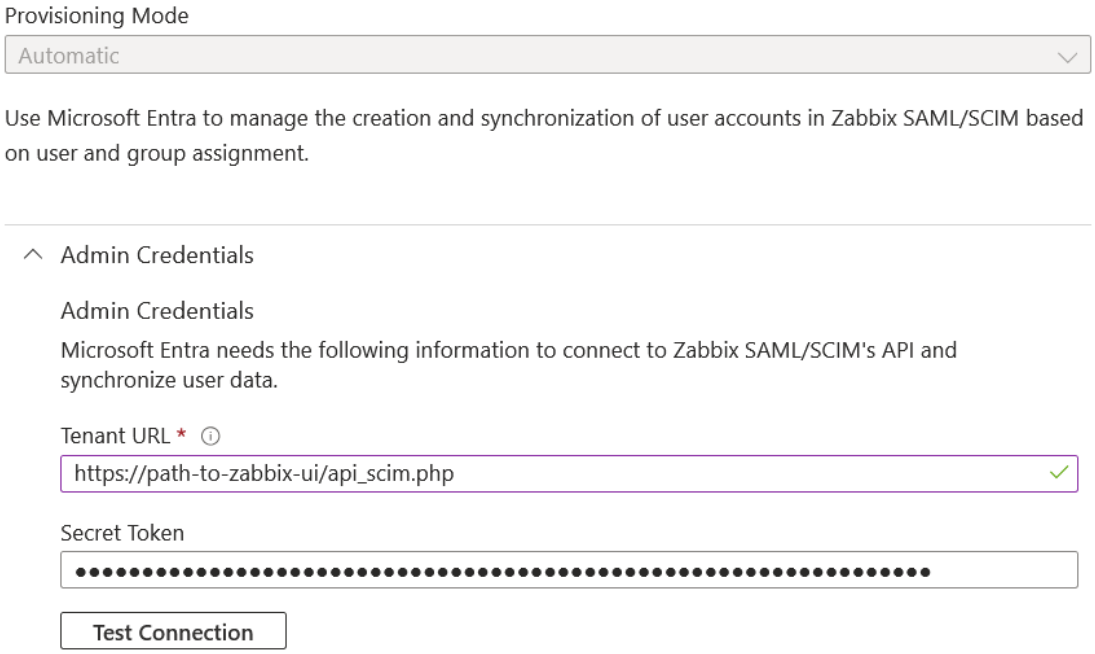
If `$SSO['CERT_STORAGE'] = 'file'` is set in the configuration file, the certificate must be made available on the filesystem (by default in `ui/conf/certs` or the path configured in the configuration file) and the frontend will not store certificates in the database. Note that if `$SSO['CERT_STORAGE']` is not set or is commented out, file storage is assumed and certificates are read from `ui/conf/certs`.

It is also required to configure user group mapping. Media mapping is optional.

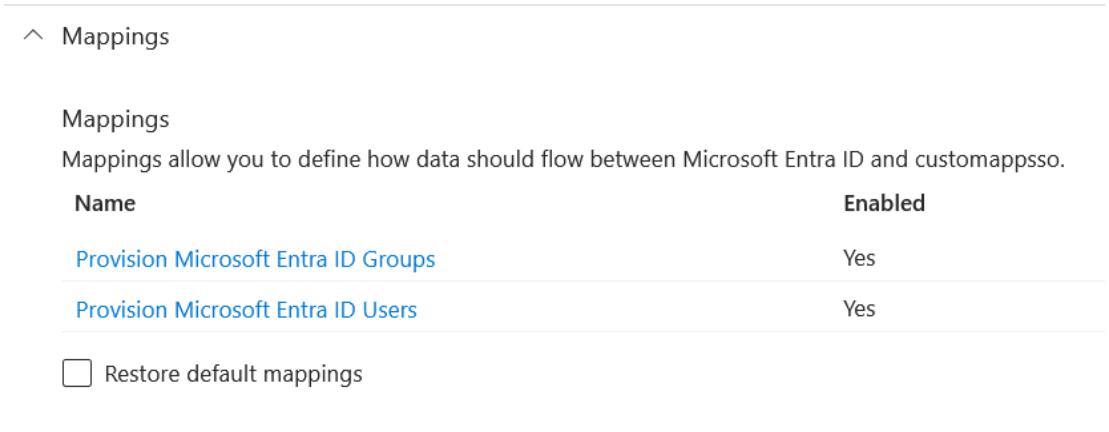
3. Click the *Update* button to save these settings.

SCIM user provisioning

1. In your Entra ID application page, from the main menu open the Provisioning page. Click *Get started* and then select Automatic provisioning mode: - In *Tenant URL*, set the following value: `https://<path-to-zabbix-ui>/api_scim.php` - In *Secret token*, enter a Zabbix API token with Super admin permissions. - Click *Test connection* to see if the connection is established.

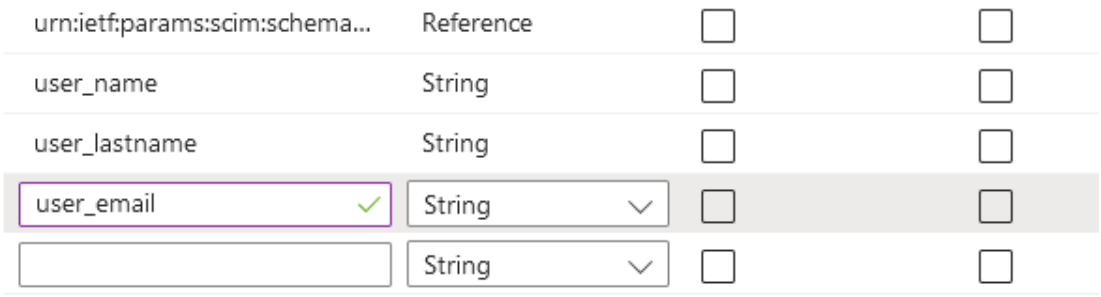


2. Now you can add all the attributes that will be passed with SCIM to Zabbix. To do that, click *Mappings* and then on *Provision Microsoft Entra ID Users*.



At the bottom of the Attribute Mapping list, enable *Show advanced options*, and then click *Edit attribute list for customappsso*.

At the bottom of the attribute list, add your own attributes with type 'String':



Save the list.

3. Now you can add mappings for the added attributes. At the bottom of the Attribute Mapping list, click *Add New Mapping* and create mappings as shown below:



When all mappings are added, save the list of mappings.



4. As a prerequisite of user provisioning into Zabbix, you must have users and groups configured in Entra ID.

To do that, go to *Microsoft Entra admin center* and then add users/groups in the respective Users and Groups pages.

5. When users and groups have been created in Entra ID, you can go to the *Users and groups* menu of your application and add them to the app.

6. Go to the *Provisioning* menu of your app, and click *Start provisioning* to have users provisioned to Zabbix.

Note that the Users PATCH request in Entra ID does not support changes in media.

Authentication request signing

It is possible to configure Entra ID to validate the signature of signed authentication requests.

To make this work, create public/private keys:

```
openssl req -x509 -newkey rsa:4096 -keyout /usr/share/zabbix/conf/certs/request-sign.key -out /usr/share/z
```

Assign permissions:

```
chown apache /usr/share/zabbix/conf/certs/request-sign.key
chmod 400 /usr/share/zabbix/conf/certs/request-sign.key
```

Update the Zabbix frontend configuration by adding:

```
$SSO['SP_KEY'] = 'conf/certs/request-sign.key';
$SSO['SP_CERT'] = 'conf/certs/request-sign.crt';
```

Troubleshooting

Authentication issues may occur with MicroSoft Edge browsers when a user, trying to login to Zabbix via SAML, is already logged in with the MicroSoft Edge profile. As a sign of such issue the user may be able to log in to Zabbix using MicroSoft Edge in private mode.

To avoid authentication issues in this case it may be necessary to set `requestedAuthnContext` to `false` in the Zabbix frontend configuration file.

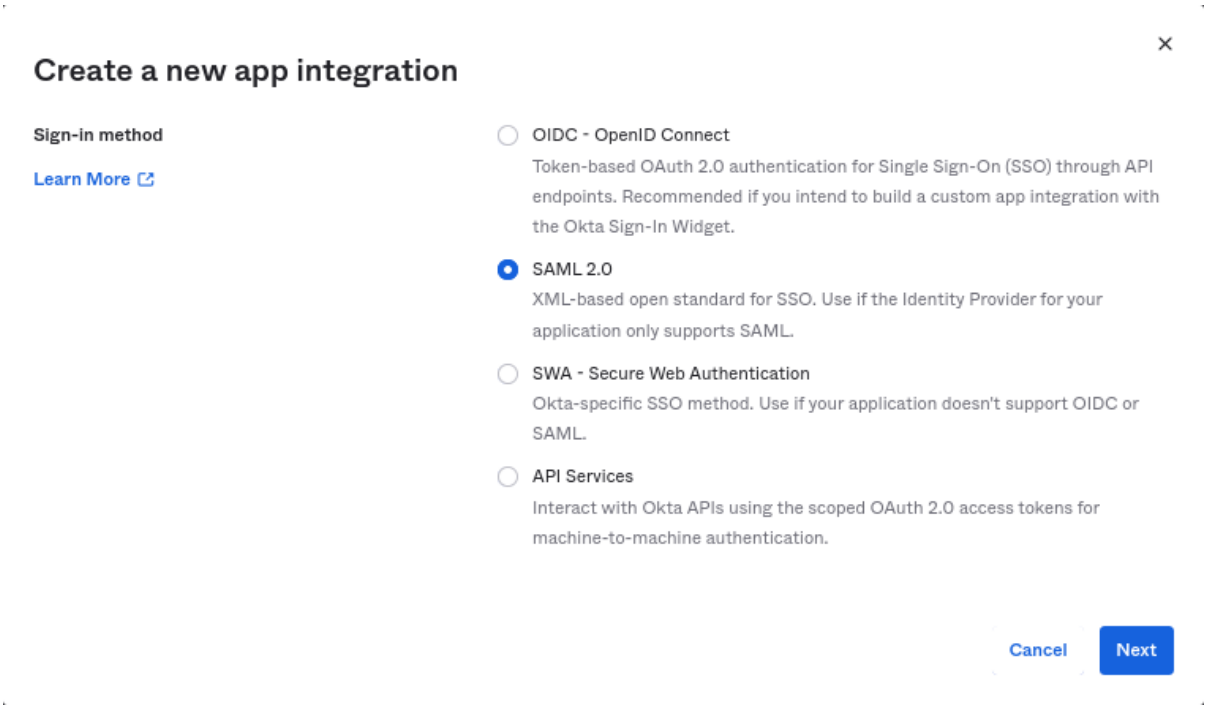
```
$SSO['SETTINGS'] = [
    'security' => [
        'requestedAuthnContext' => false
    ]
];
```

SAML setup with Okta

This section provides guidelines for configuring Okta to enable SAML 2.0 authentication and user provisioning for Zabbix.

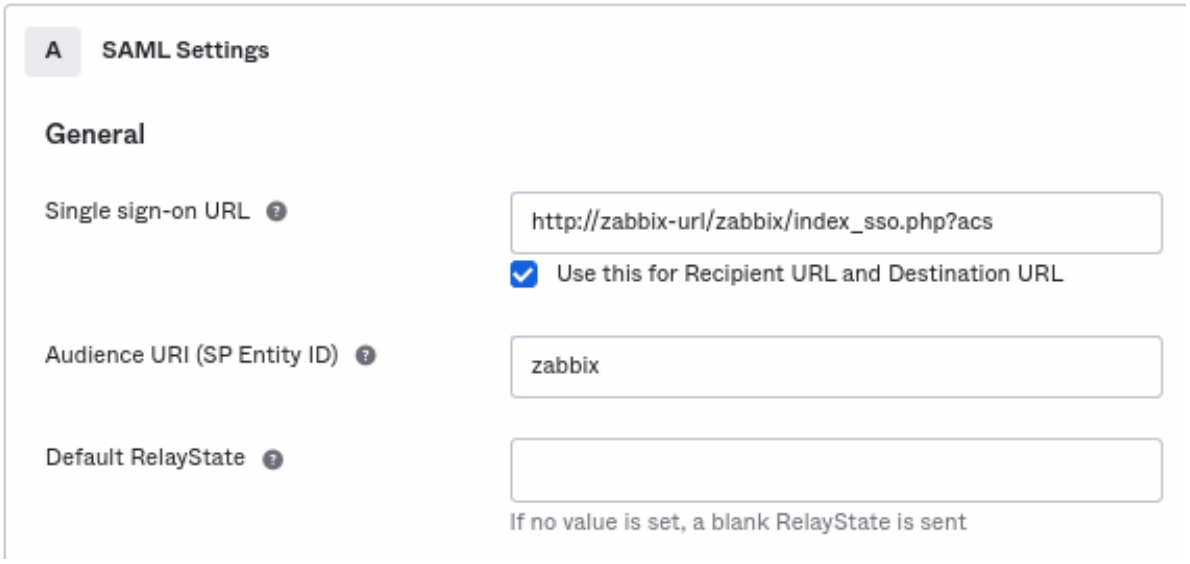
Okta configuration

1. Go to https://developer.okta.com/signup/ and register/sign into your account.
2. In the Okta web interface navigate to *Applications* > *Applications*.
3. Click *Create App Integration*.

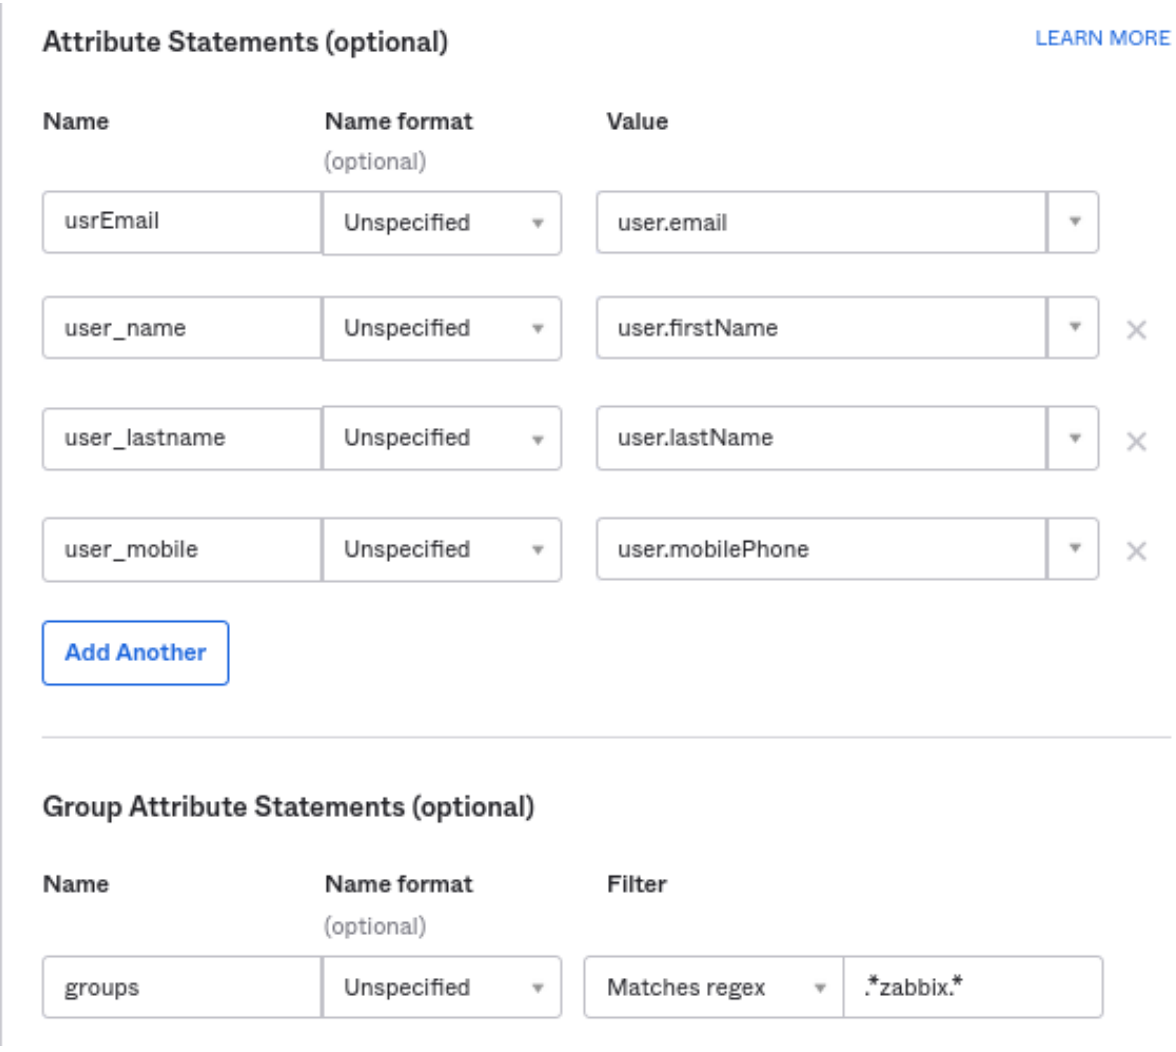


Select "SAML 2.0" as the sign-in method and click *Next*.

4. In general settings, fill in the app name and click *Next*.
5. In SAML configuration, enter the values provided below, then click *Next*.



- In **General** add:
  - *Single sign-on URL*: `http://<your-zabbix-url>/zabbix/index_sso.php?acs<br>` Note the use of "http", and not "https", so that the `acs` parameter is not cut out in the request. The *Use this for Recipient URL and Destination URL* checkbox should also be marked.
  - *Audience URI (SP Entity ID)*: `zabbix<br>` Note that this value will be used within the SAML assertion as a unique service provider identifier (if not matching, the operation will be rejected). It is possible to specify a URL or any string of data in this field.
  - *Default RelayState*:
 Leave this field blank; if a custom redirect is required, it can be added in Zabbix in the *Users* > *Users* settings.
  - Fill in other fields according to your preferences.
- In **Attribute Statements/Group Attribute Statements** add:



These attribute statements are inserted into the SAML assertions shared with Zabbix.

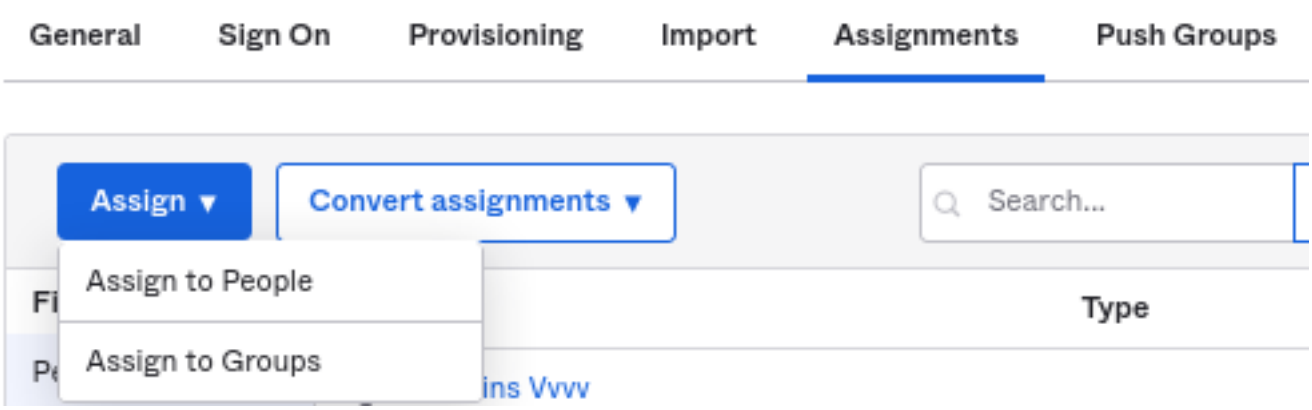
The attribute names used here are arbitrary examples. You may use different attribute names, however, it is required that they match the respective field value in Zabbix SAML settings.

If you want to configure SAML sign-in into Zabbix *without* JIT user provisioning, then only the email attribute is required.

Note:
If planning to use an encrypted connection, generate the private and public encryption certificates, then upload the public certificate to Okta. The certificate upload form appears when *Assertion Encryption* is set to "Encrypted" (click *Show Advanced Settings* to find this parameter).

6. In the next tab, select "I'm a software vendor. I'd like to integrate my app with Okta" and click *Finish*.

7. Navigate to the "Assignments" tab of the newly-created application and click the *Assign* button, then select "Assign to People" from the drop-down.



8. In a popup that appears, assign the app to people that will use SAML 2.0 to authenticate with Zabbix, then click *Save and go back*.

9. Download the IdP certificate. Set 644 permissions to it by running:

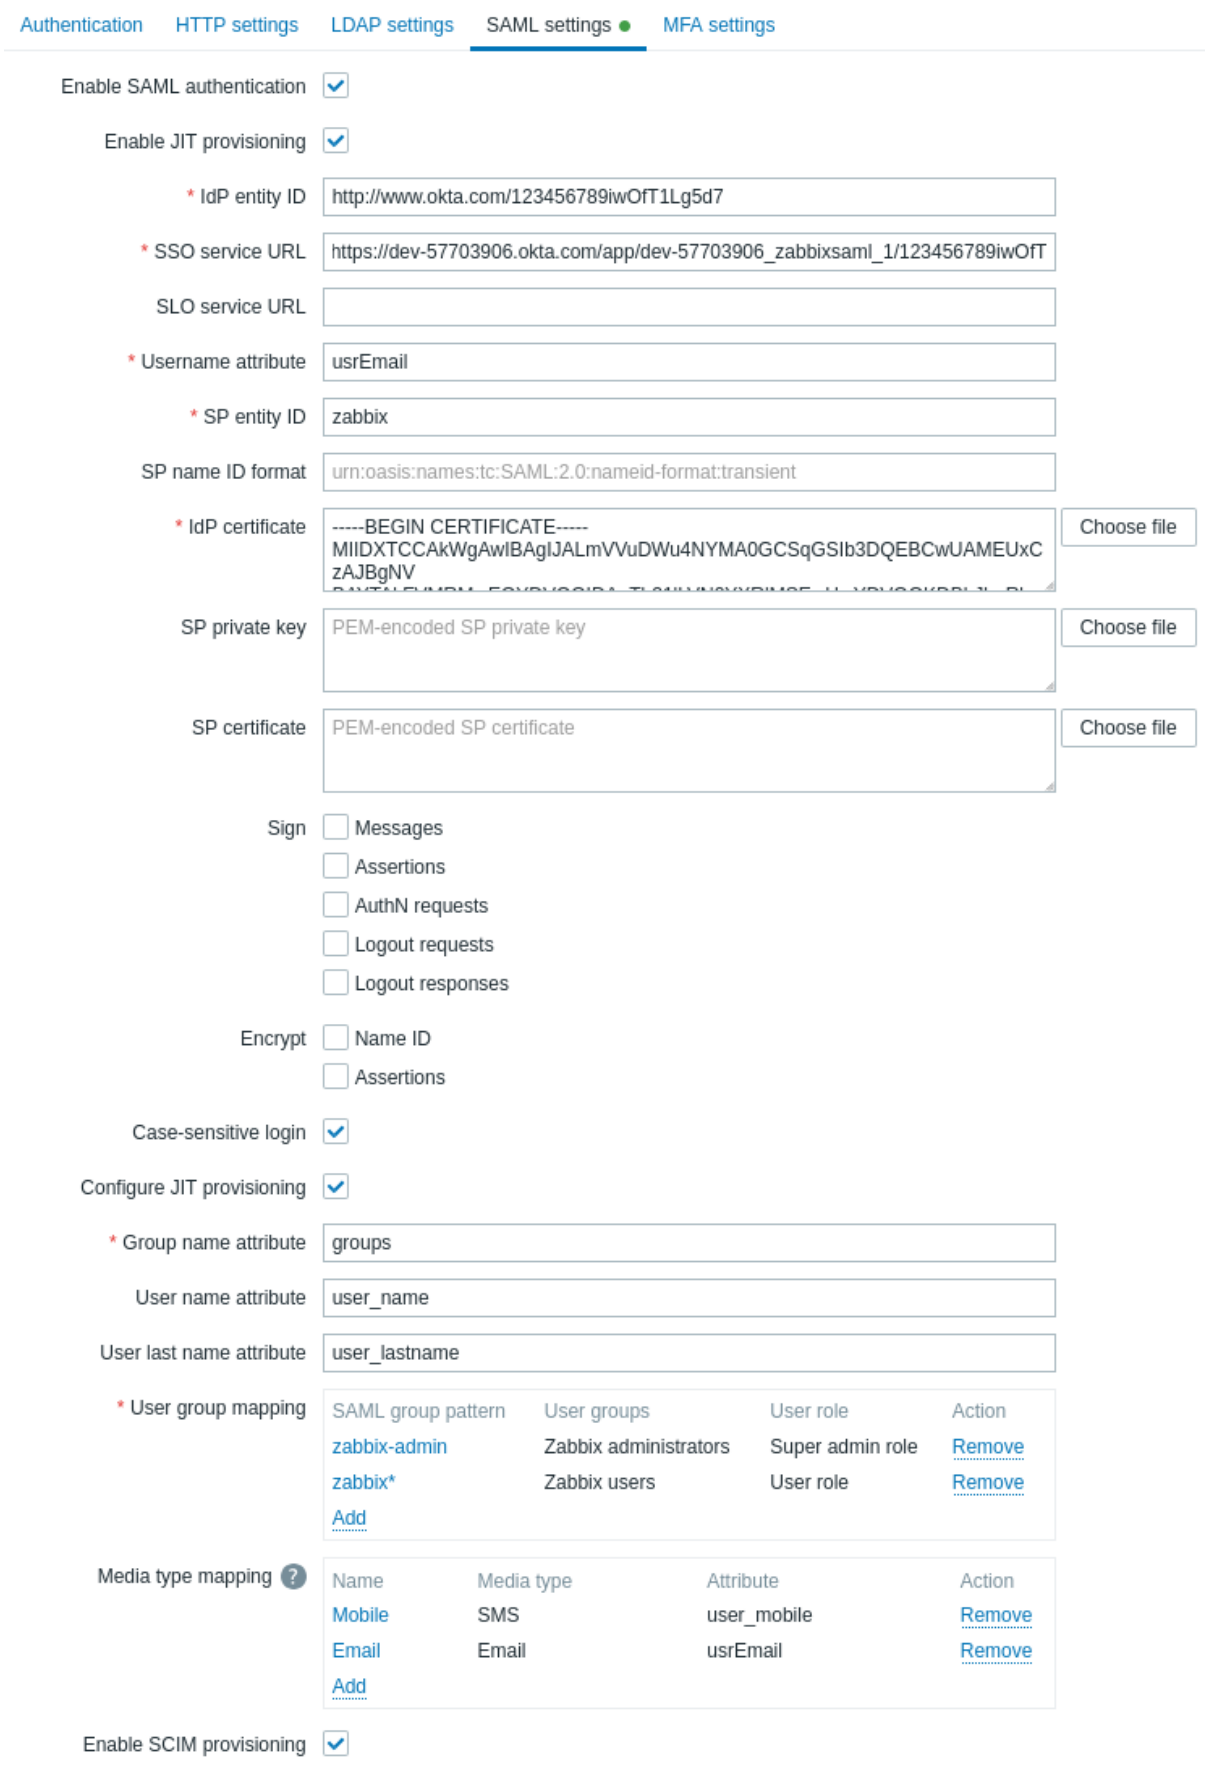
```
chmod 644 idp.crt
```

10. Navigate to the "Sign On" tab and click the *View Setup Instructions* button.

Setup **instructions** will be opened in a new tab; keep this tab open while configuring Zabbix.

Zabbix configuration

1. In Zabbix, go to the SAML settings and fill the configuration options based on setup instructions from Okta:



Zabbix field	Setup field in Okta	Sample value
IdP entity ID	Identity Provider Issuer	
SSO service URL	Identity Provider Single Sign-On URL	
Username attribute	Attribute name	`usrEmail`
SP entity ID	Audience URI	`zabbix`
Group name attribute	Attribute name	`groups`

Zabbix field	Setup field in Okta	Sample value
User name attribute	Attribute name	`user_name`
User last name attribute	Attribute name	`user_lastname`

It is also required to configure user group and media mapping.

2. Add the Base64 certificate provided in the Okta SAML setup instructions.

If `$SSO['CERT_STORAGE'] = 'database'` is set in the frontend configuration file (`zabbix.conf.php`), you can paste the certificate text or upload the certificate file in the frontend during SAML configuration—no files are required on the filesystem.

If `$SSO['CERT_STORAGE'] = 'file'` is set in the configuration file, the certificate must be made available on the filesystem (by default in `ui/conf/certs` or the path configured in the configuration file) and the frontend will not store certificates in the database. Note that if `$SSO['CERT_STORAGE']` is not set or is commented out, file storage is assumed and certificates are read from `ui/conf/certs`.

3. If *Assertion Encryption* has been set to "Encrypted" in Okta, the "Assertions" checkbox of the *Encrypt* parameter should be marked in Zabbix as well.

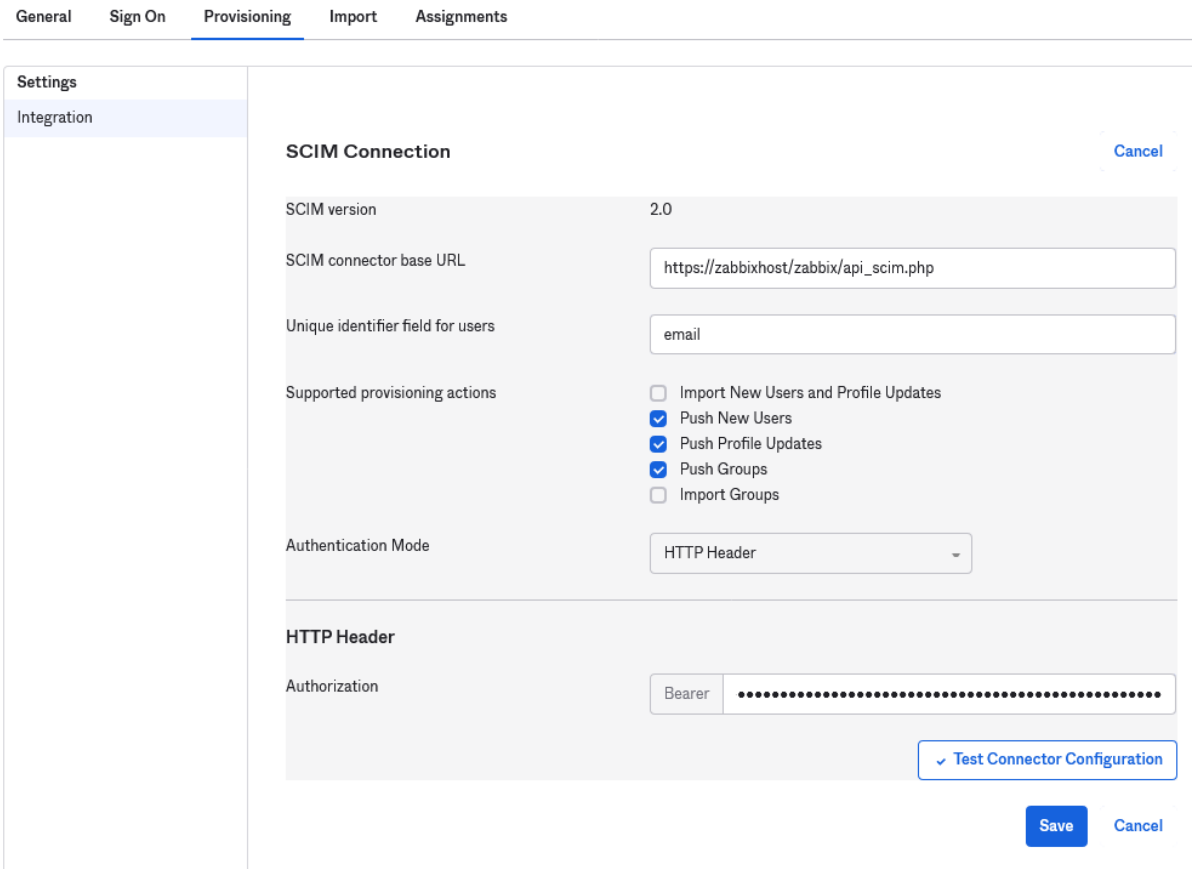
4. Click the *Update* button to save these settings.

SCIM provisioning

1. To turn on SCIM provisioning, go to *General* > *App Settings* of the application in Okta.

Mark the *Enable SCIM provisioning* checkbox. As a result, a new *Provisioning* tab appears.

2. Go to the "Provisioning" tab to set up a SCIM connection: - In *SCIM connector base URL*, specify the path to the Zabbix frontend and append `api_scim.php` to it, i.e.:
 `https://<your-zabbix-url>/zabbix/api_scim.php` - *Unique identifier field for users*: `email` - *Authentication mode*: `HTTP header` - In *Authorization*, enter a valid API token with Super admin rights.



Attention:
If you are experiencing authentication issues, see Authorization header forwarding.

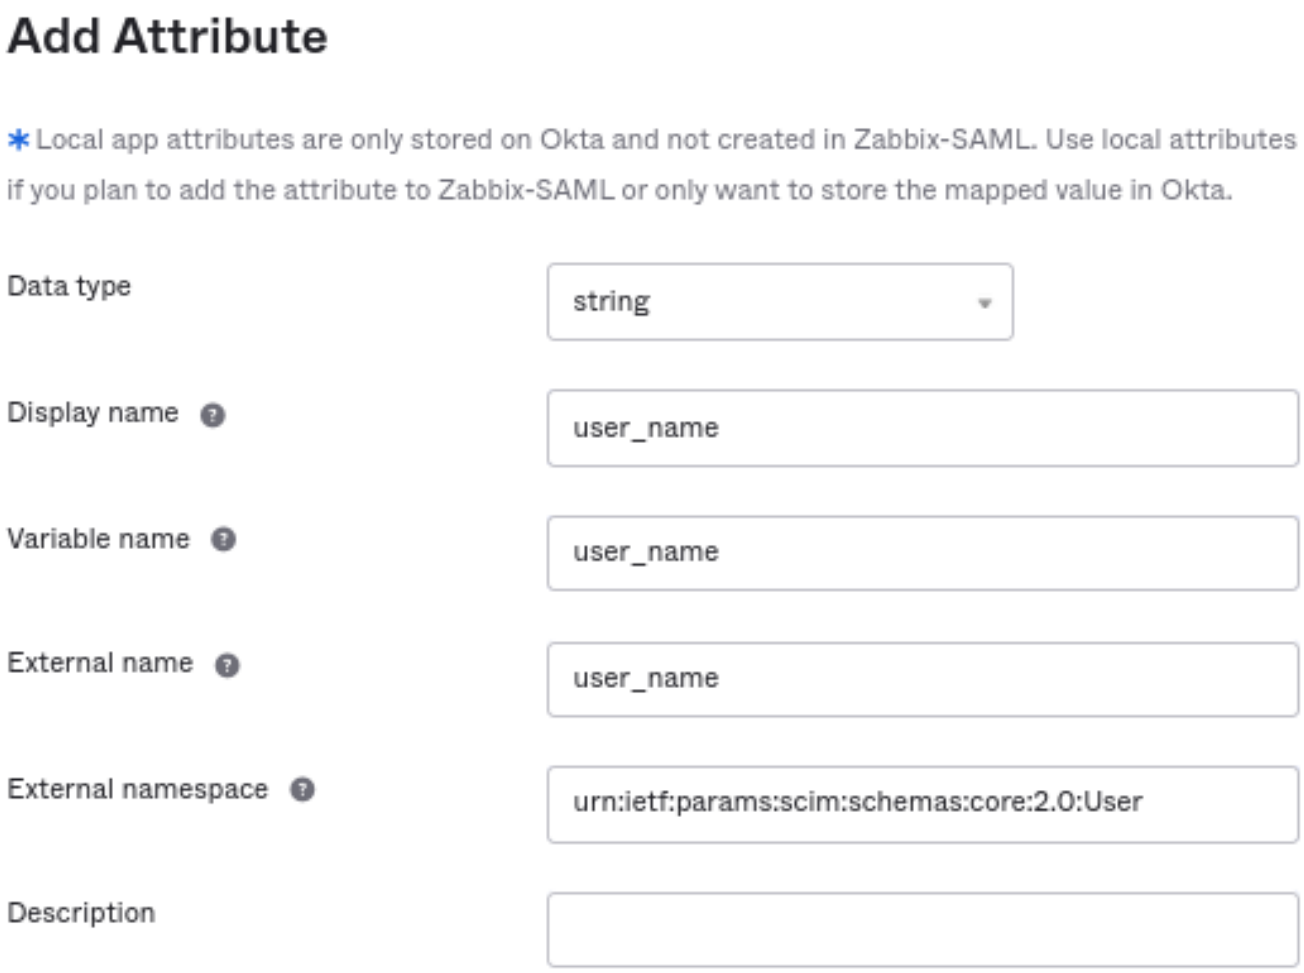
3. Click *Test Connector Configuration* to test the connection. If all is correct a success message will be displayed.

4. In *Provisioning* > *To App*, make sure to mark the following checkboxes: - Create Users - Update User Attributes - Deactivate Users

This will make sure that these request types will be sent to Zabbix.

5. Make sure that all attributes defined in SAML are defined in SCIM. You can access the profile editor for your app in *Provisioning* > *To App*, by clicking *Go to Profile Editor*.

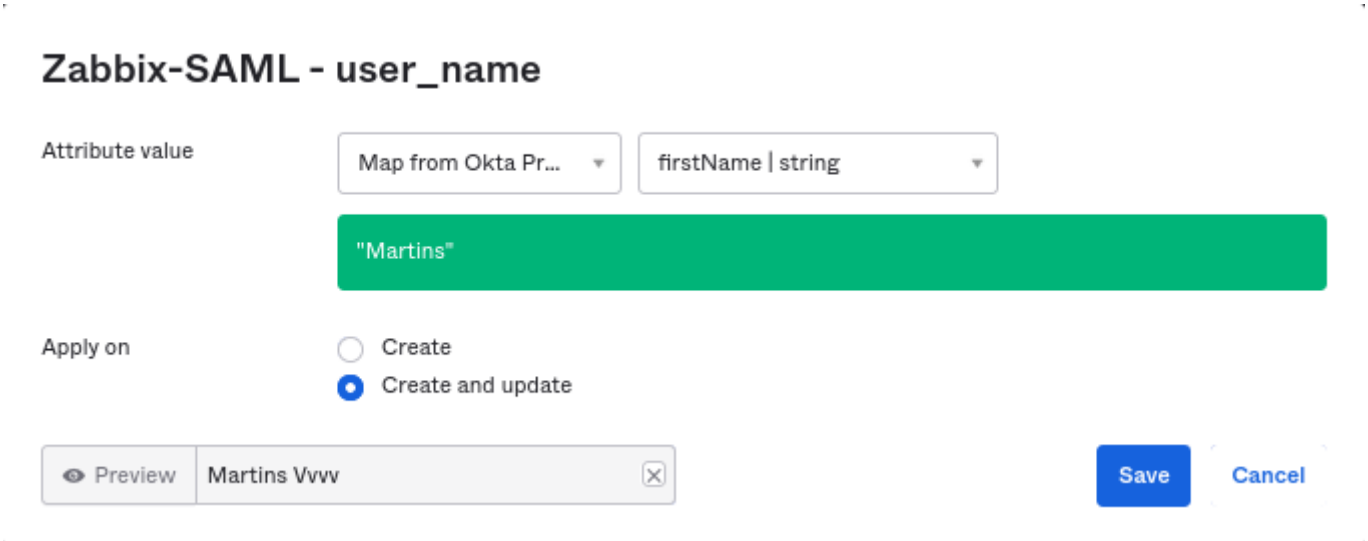
Click *Add Attribute*. Fill the values for *Display name*, *Variable name*, *External name* with the SAML attribute name, for example, `user_name`.



External namespace should be the same as user schema: `urn:ietf:params:scim:schemas:core:2.0:User`

6. Go to *Provisioning* > *To App* > *Attribute Mappings* of your application. Click *Show Unmapped Attributes* at the bottom. Newly added attributes appear.

7. Map each added attribute.



8. Add users in the ”Assignments” tab. The users previously need to be added in *Directory* > *People*. All these assignments will

be sent as requests to Zabbix.

9. Add groups in the "Push Groups" tab. The user group mapping pattern in Zabbix SAML settings must match a group specified here. If there is no match, the user cannot be created in Zabbix.

Information about group members is sent every time when some change is made.

SAML setup with OneLogin

Overview

This section provides guidelines for configuring single sign-on and user provisioning into Zabbix from OneLogin using SAML 2.0 authentication.

OneLogin configuration

Creating application

1. Log into your account at OneLogin. For testing purposes, you may create a free developer account in OneLogin.

2. In the OneLogin web interface navigate to *Applications* > *Applications*.

3. Click "Add App" and search for the appropriate app. The guidelines in this page are based on the *SCIM Provisioner with SAML (SCIM v2 Enterprise, full SAML)* app example.

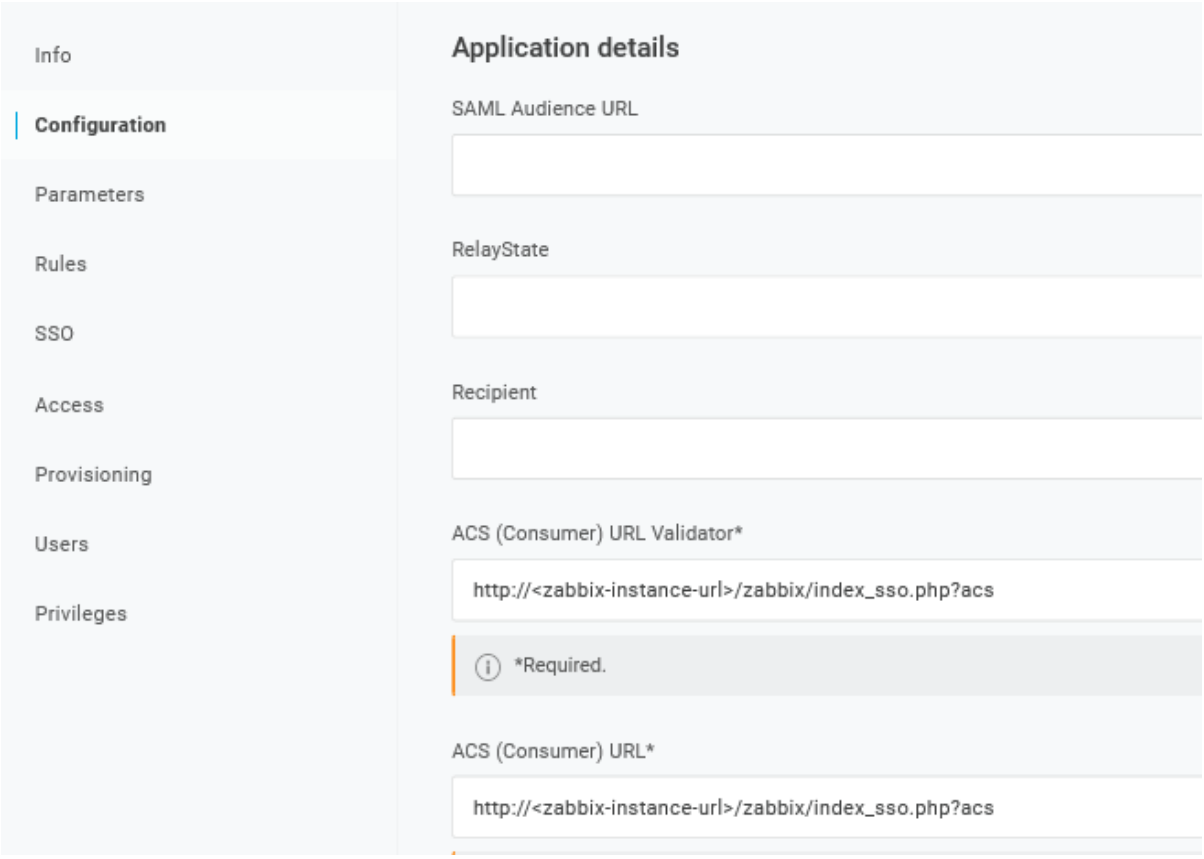
4. To begin with, you may want to customize the display name of your app. You may also want to add the icon and app details. After that, click *Save*.

Setting up SSO/SCIM provisioning

1. In *Configuration* > *Application details*, set the Zabbix single sign-on endpoint `http://<zabbix-instance-url>/zabbix/index_sso.p` as the value of these fields:

- *ACS (Consumer) URL Validator*
- *ACS (Consumer) URL*

Note the use of "http", and not "https", so that the `acs` parameter is not cut out in the request.



It is also possible to use "https". To use this with Zabbix, add the following line to the frontend configuration file (`zabbix.conf.php`):

```
$SSO['SETTINGS'] = ['use_proxy_headers' => true];
```

Leave other options with their default values.

2. In *Configuration* > *API connection*, set the following values:

- *SCIM Base URL*: `https://<zabbix-instance-url>/zabbix/api_scim.php`
- *SCIM JSON Template*: should contain all custom attributes that you would like to pass to Zabbix via SCIM such as `user_name`, `user_lastname`, `user_email`, and `user_mobile`:

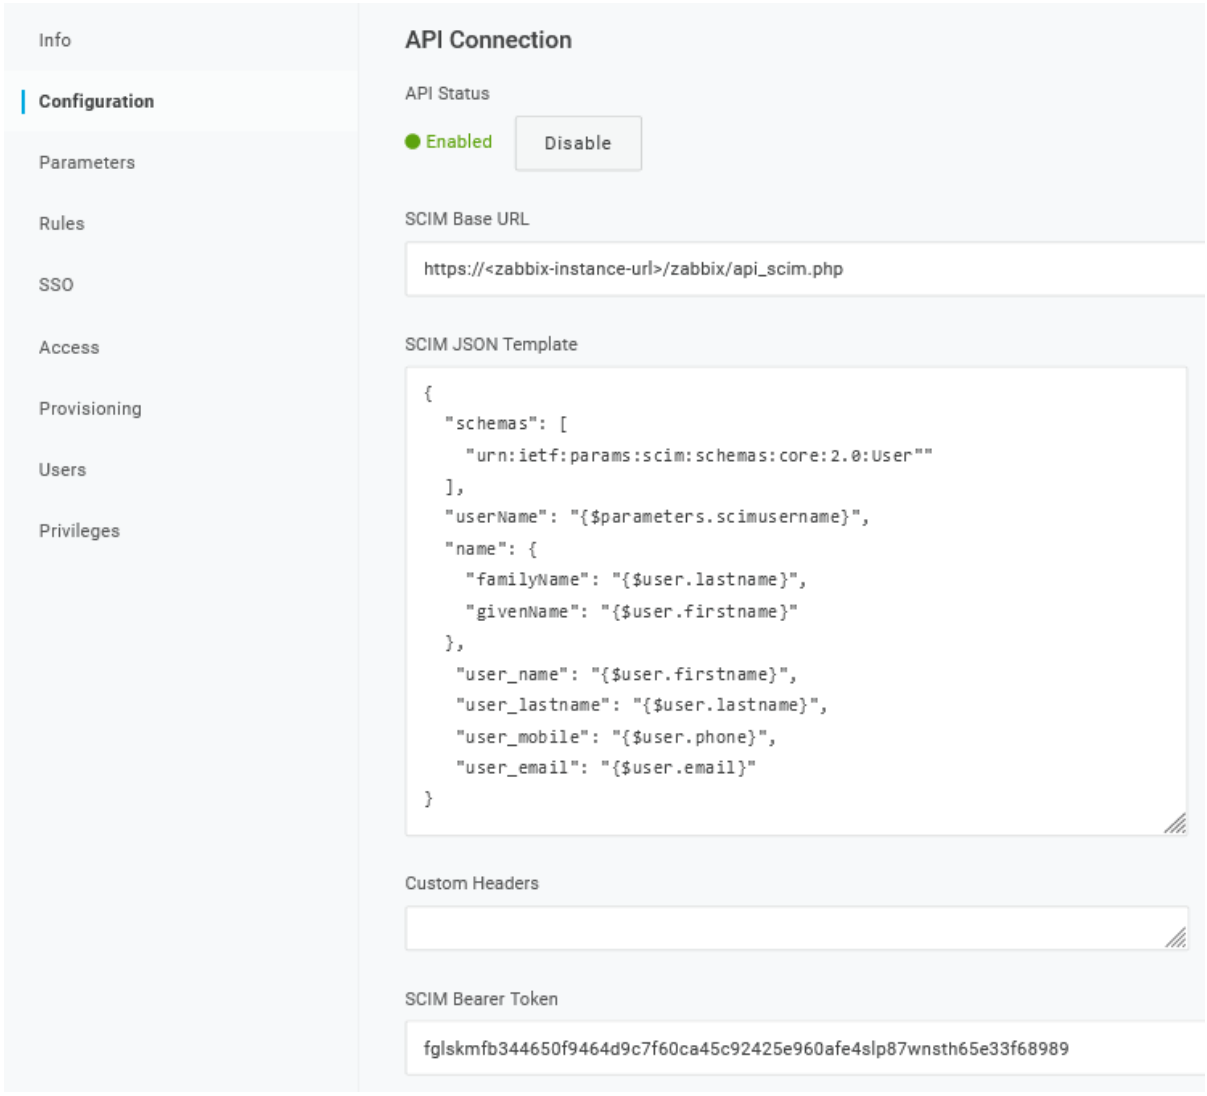
```
{
  "schemas": [
    "urn:ietf:params:scim:schemas:core:2.0:User"
  ],
  "userName": "{$parameters.scimusername}",
  "name": {
    "familyName": "{$user.lastname}",
    "givenName": "{$user.firstname}"
  },
   "user_name": "{$user.firstname}",
   "user_lastname": "{$user.lastname}",
   "user_mobile": "{$user.phone}",
   "user_email": "{$user.email}"
}
```

The attribute names are arbitrary. Different attribute names may be used, however, it is required that they match the respective field value in Zabbix SAML settings.

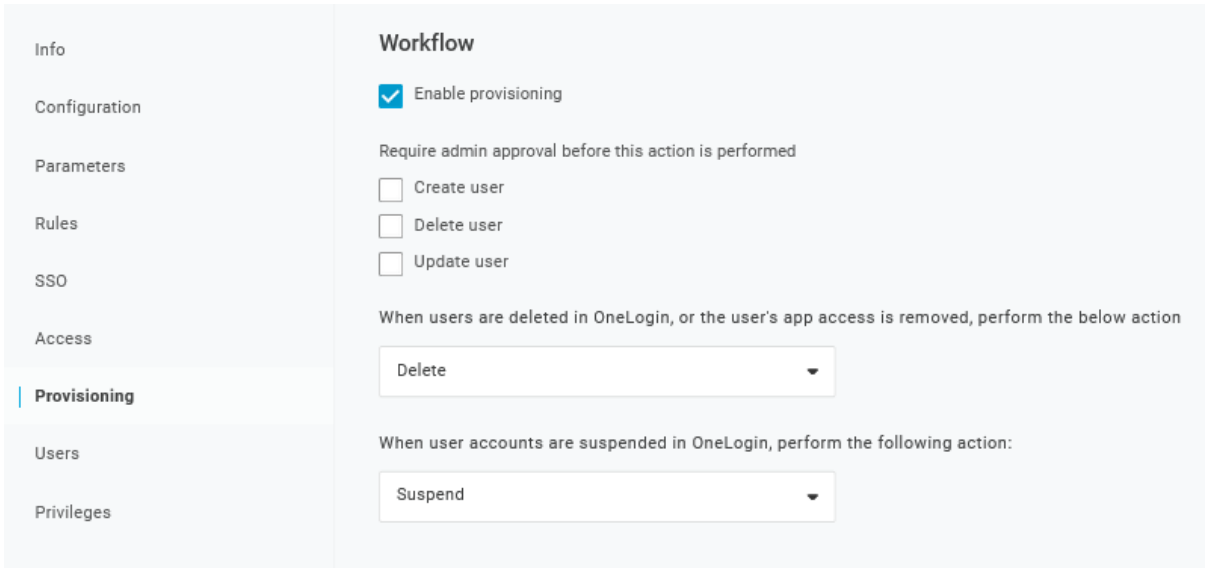
Note that for user provisioning to work, OneLogin needs to receive in response a 'name' attribute with 'givenName' and 'familyName', even if it was not required by the service provider. Thus it is necessary to specify this in the schema in the application configuration part.

- *SCIM Bearer Token*: enter a Zabbix API token with Super admin permissions.

Click *Enable* to activate the connection.

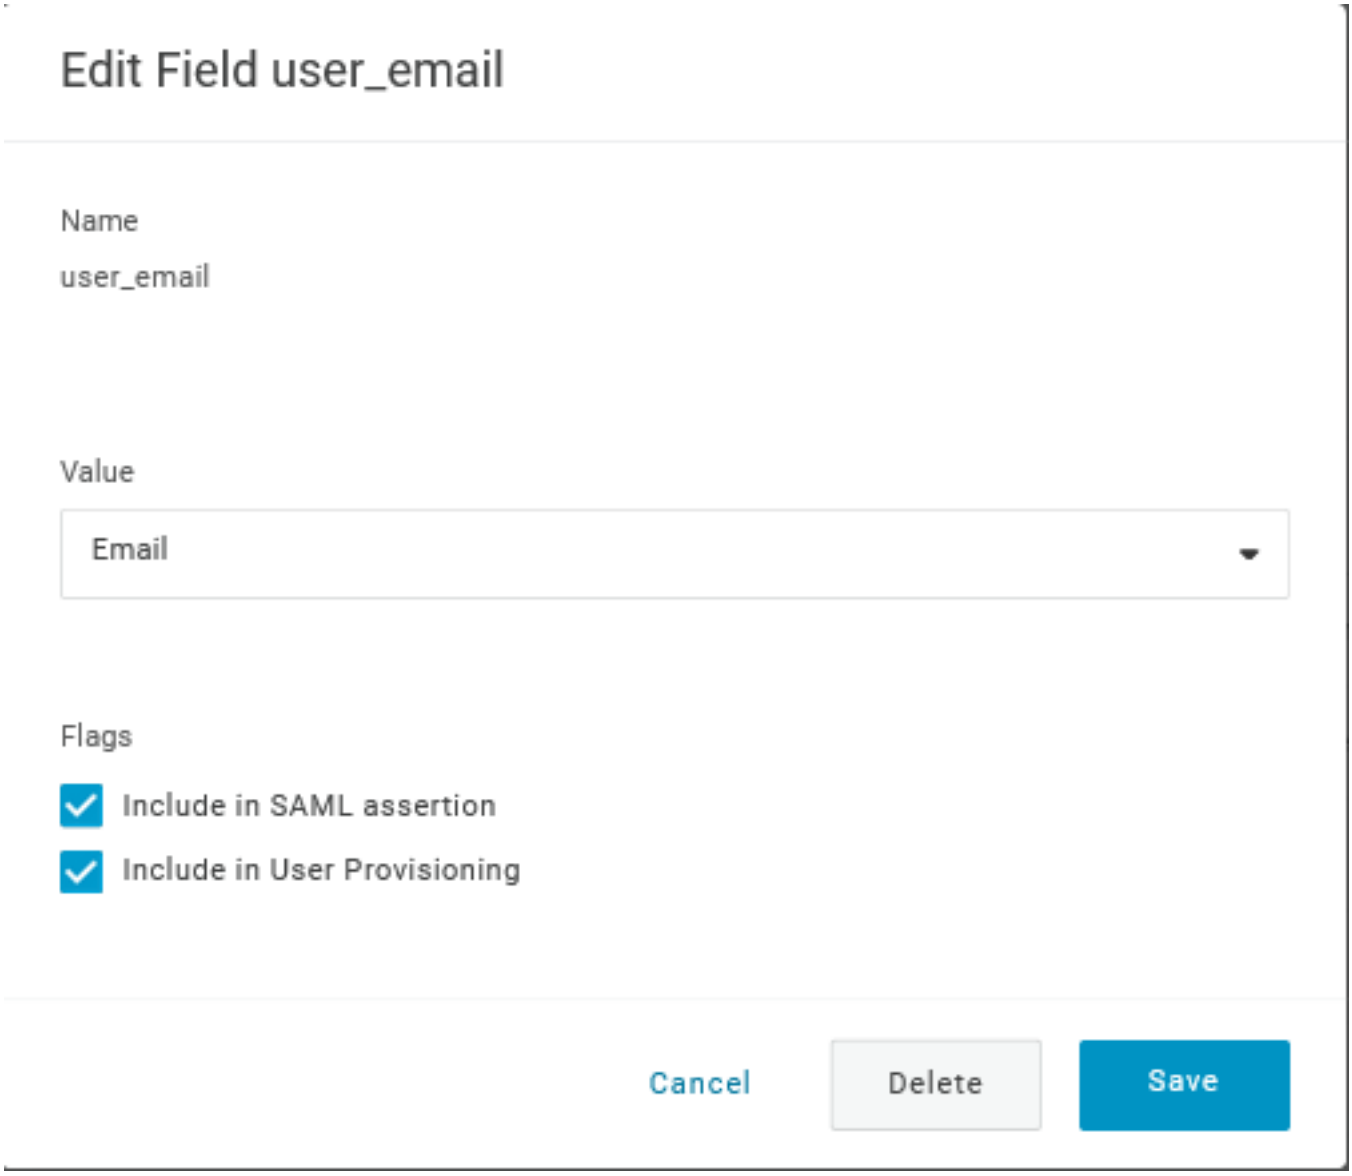


3. In the *Provisioning* page, enable the Provisioning option:



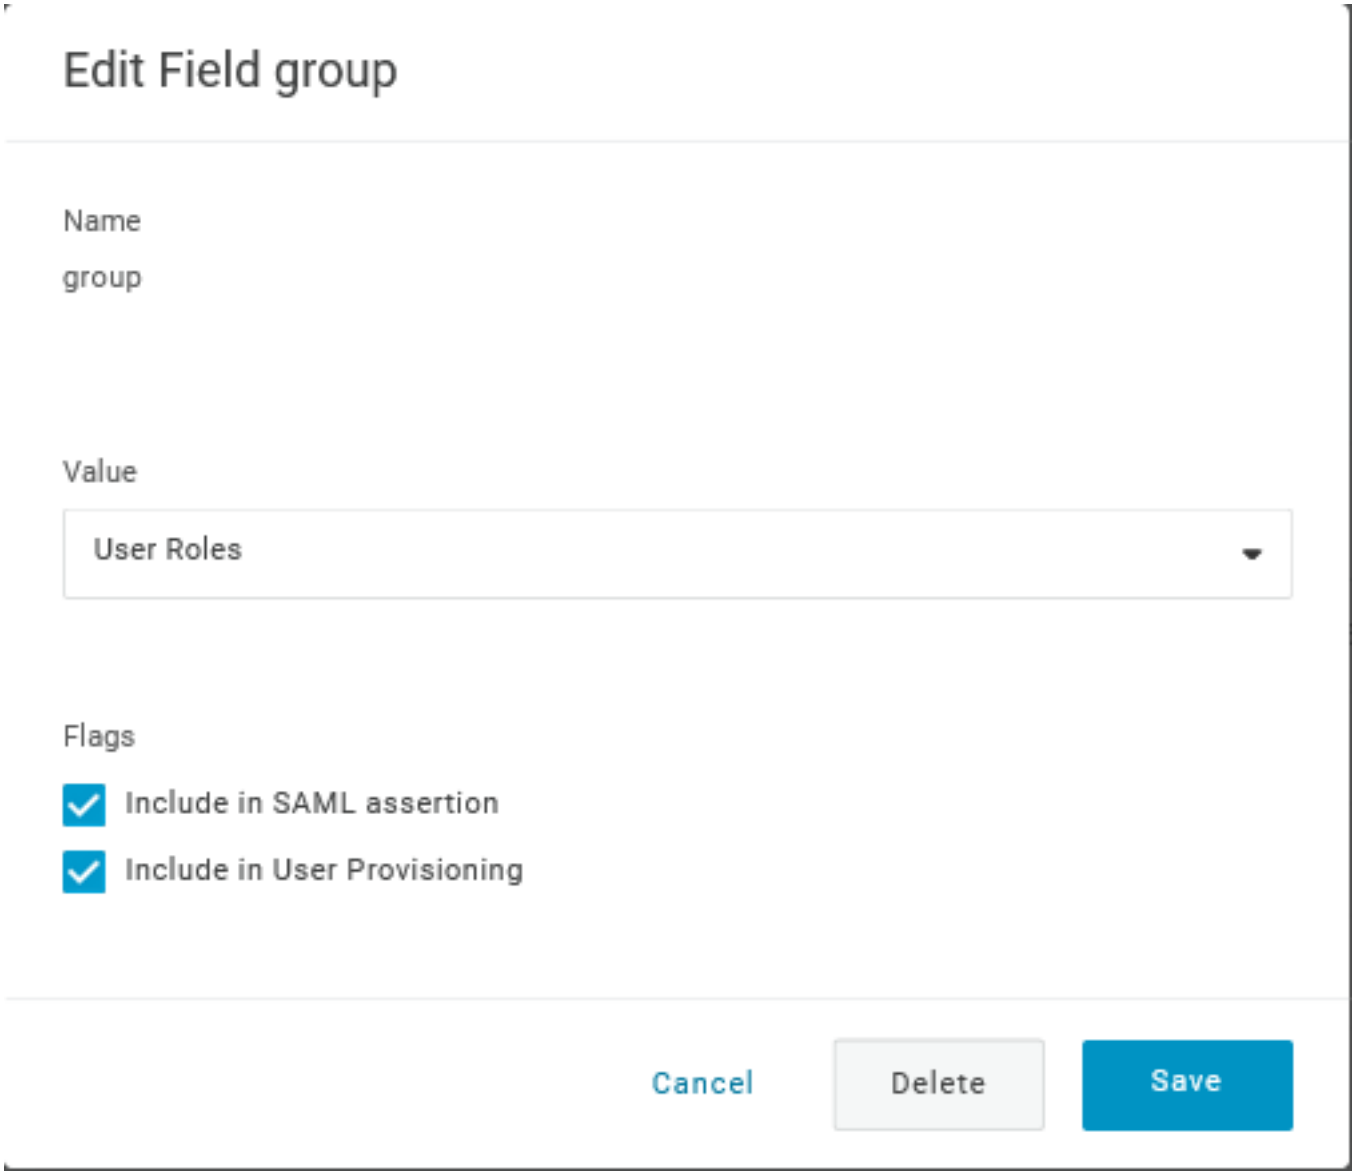
4. The *Parameters* page contains a list of default parameters:

- Make sure that the 'scimusername' matches the user login value in OneLogin (e.g., email);
- Mark the *Include in User Provisioning* option for the 'Groups' parameter;
- Click "+" to create the custom parameters that are required for SAML assertions and user provisioning such as `user_name`, `user_lastname`, `user_email`, and `user_mobile`:

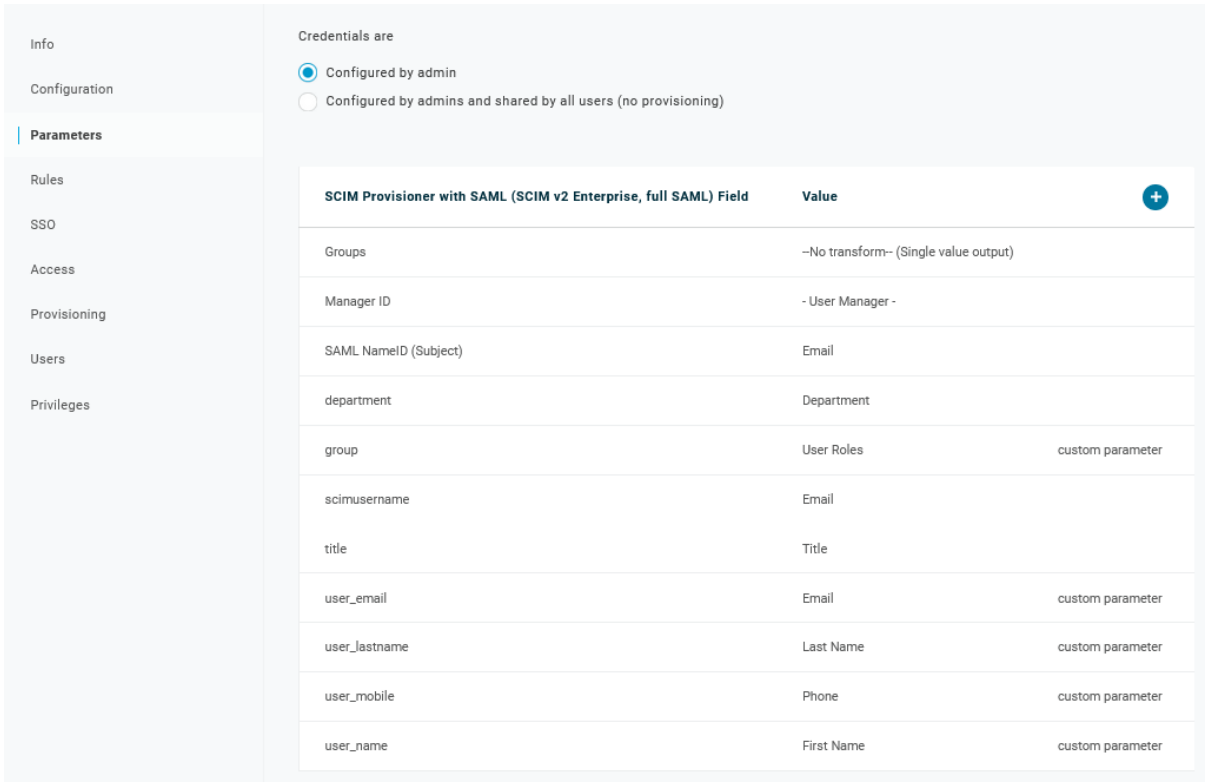


When adding a parameter, make sure to mark both the *Include in SAML assertion* and *Include in User Provisioning* options.

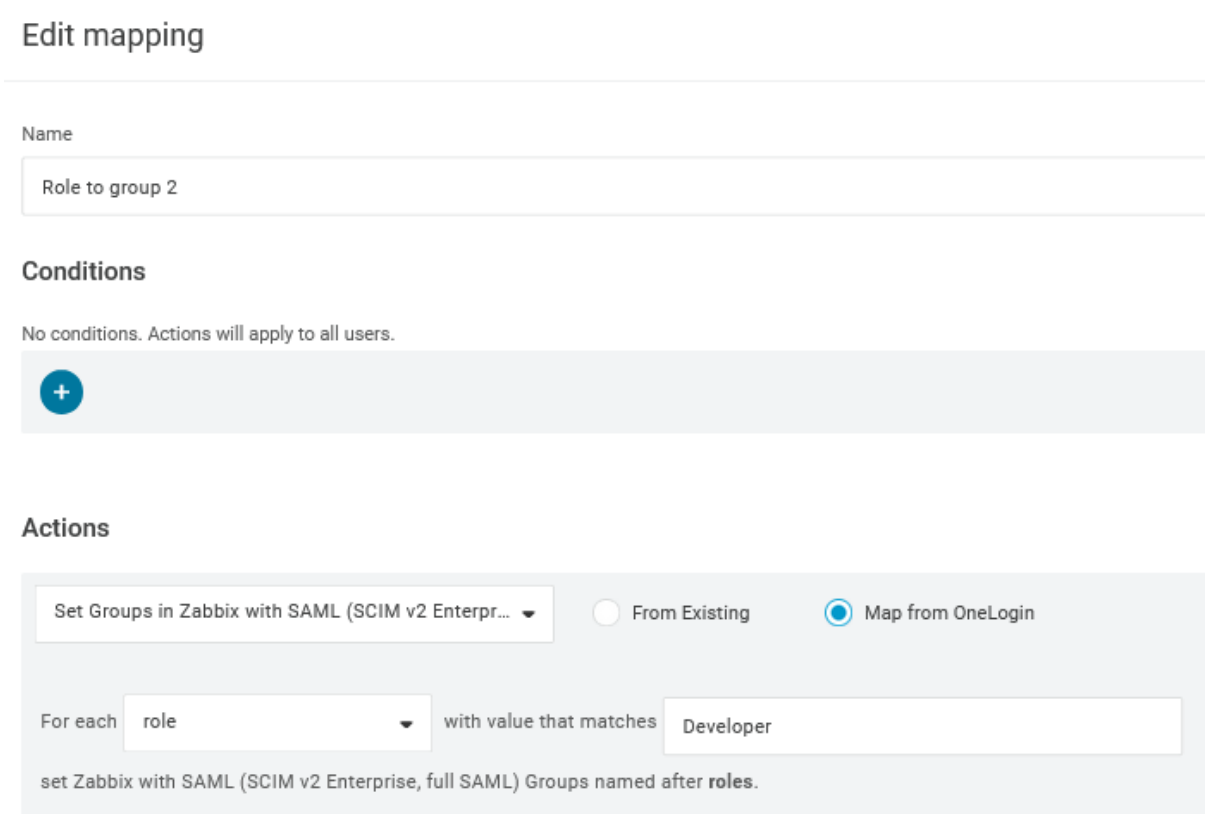
- Add a 'group' parameter that matches user roles in OneLogin. User roles will be passed as a string, separated by a semicolon `;`. The OneLogin user roles will be the used for creating user groups in Zabbix:



Verify the list of parameters:



5. In the *Rules* page, create user role mappings to the default Groups parameter.



You may use a regular expression to pass specific roles as groups. The role names should not contain ; as OneLogin uses it as a separator when sending an attribute with several roles.

6. Download the IdP certificate. Set 644 permissions to it by running:

```
chmod 644 idp.crt
```

Zabbix configuration

1. In Zabbix, go to the SAML settings and fill the configuration options based on the OneLogin configuration:



Zabbix field	Setup field in OneLogin	Sample value
IdP entity ID	Issuer URL (see SSO tab of your application in OneLogin)	
SSO service URL	SAML 2.0 Endpoint (HTTP) (see SSO tab of your application in OneLogin)	

Zabbix field	Setup field in OneLogin	Sample value
SLO service URL	SLO Endpoint (HTTP) (see SSO tab of your application in OneLogin)	
Username attribute	Custom parameter	`user_email`
Group name attribute	Custom parameter	`group`
User name attribute	Custom parameter	`user_name`
User last name attribute	Custom parameter	`user_lastname`

It is also required to configure user group mapping. Media mapping is optional. Click *Update* to save these settings.

2. Add the Base64 certificate provided by OneLogin.

If `$SSO['CERT_STORAGE'] = 'database'` is set in the frontend configuration file (`zabbix.conf.php`), you can paste the certificate text or upload the certificate file in the frontend during SAML configuration—no files are required on the filesystem.

If `$SSO['CERT_STORAGE'] = 'file'` is set in the configuration file, the certificate must be made available on the filesystem (by default in `ui/conf/certs` or the path configured in the configuration file) and the frontend will not store certificates in the database. Note that if `$SSO['CERT_STORAGE']` is not set or is commented out, file storage is assumed and certificates are read from `ui/conf/certs`.

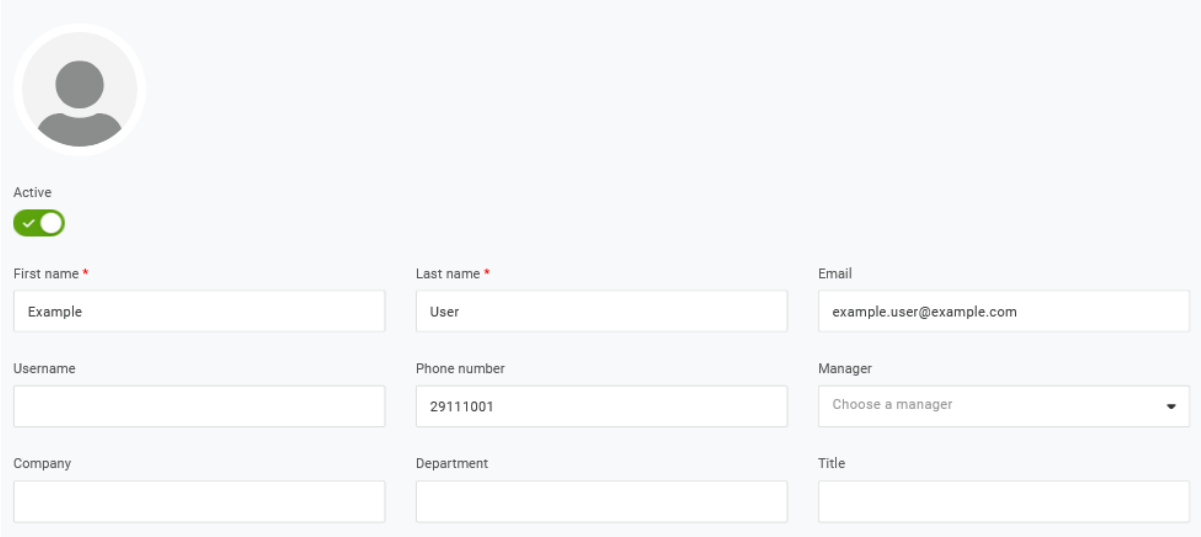
You can access the certificate download in OneLogin in *Applications* > *SSO* > click *View details* under the current certificate.

3. Click the *Update* button to save these settings.

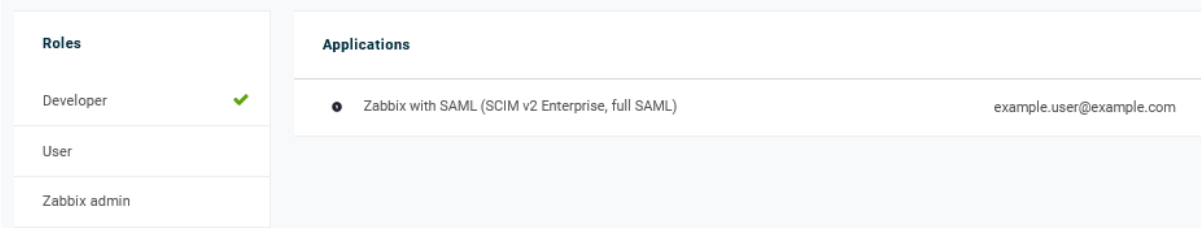
SCIM user provisioning

With user provisioning enabled, it is now possible to add/update users and their roles in OneLogin and have them immediately provisioned to Zabbix.

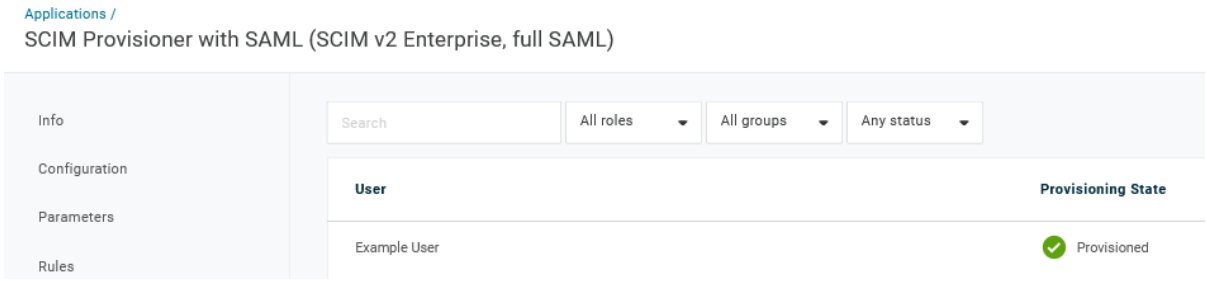
For example, you may create a new user:



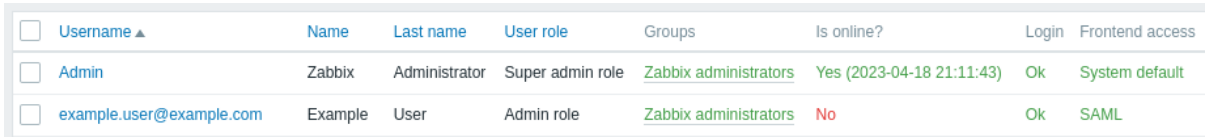
Add it to a user role and the application that will provision the user:



When saving the user, it will be provisioned to Zabbix. In *Application* > *Users* you can check the provisioning status of current application users:



If successfully provisioned, the user can be seen in the Zabbix user list.



Username ▲	Name	Last name	User role	Groups	Is online?	Login	Frontend access
Admin	Zabbix	Administrator	Super admin role	Zabbix administrators	Yes (2023-04-18 21:11:43)	Ok	System default
example.user@example.com	Example	User	Admin role	Zabbix administrators	No	Ok	SAML

Setting up scheduled reports

Overview

This section provides instructions on installing Zabbix web service and configuring Zabbix to enable generation of scheduled reports.

Installation

A new Zabbix web service process and Google Chrome browser should be installed to enable generation of scheduled reports. The web service may be installed on the same machine where the Zabbix server is installed or on a different machine. Google Chrome browser should be installed on the same machine, where the web service is installed.

The official zabbix-web-service package is available in the Zabbix repository. Google Chrome browser is not included into these packages and has to be installed separately.

To compile Zabbix web service from sources, see Installing Zabbix web service.

After the installation, run `zabbix_web_service` on the machine, where the web service is installed:

```
zabbix_web_service
```

Configuration

To ensure proper communication between all elements involved make sure the server configuration and frontend configuration parameters are properly configured.

Zabbix server

The following parameters in Zabbix server configuration file need to be updated: `WebServiceURL` and `StartReportWriters`.

WebServiceURL

This parameter is required to enable communication with the web service. The URL should be in the format `http[s]://host:port/report`.

- By default, the web service listens on port 10053. A different port can be specified in the web service configuration file.
- Specifying the `/report` path is mandatory (the path is hardcoded and cannot be changed).

Example:

```
WebServiceURL=http://localhost:10053/report
```

Attention:
It is strongly recommended to set up encryption between Zabbix server and Zabbix web service using certificates. By default, data transmitted between Zabbix server and Zabbix web service is not encrypted, which can lead to unauthorized access.

StartReportWriters

This parameter determines how many report writer processes should be started. If it is not set or equals 0, report generation is disabled. Based on the number and frequency of reports required, it is possible to enable from 1 to 100 report writer processes.

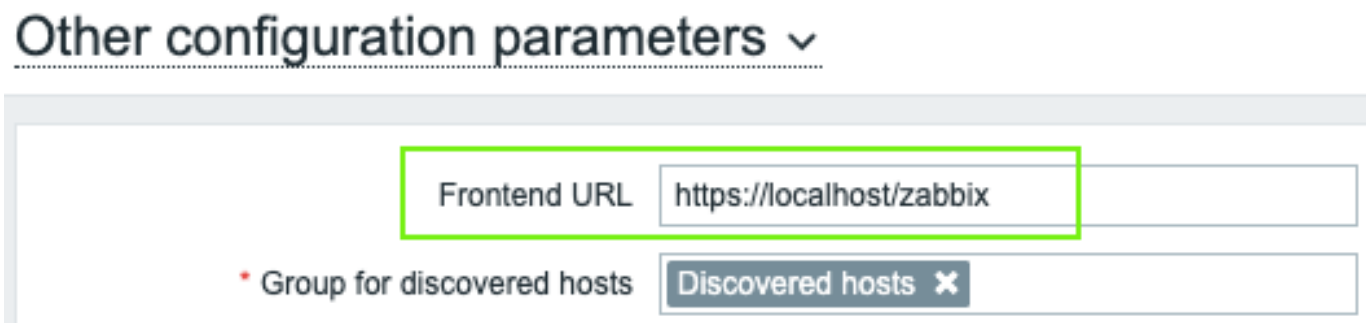
Example:

```
StartReportWriters=3
```

Zabbix frontend

A *Frontend URL* parameter should be set to enable communication between Zabbix frontend and Zabbix web service:

- Proceed to the *Administration* > *General* > *Other* frontend menu section
- Specify the full URL of the Zabbix web interface in the *Frontend URL* parameter



Note:
Once the setup procedure is completed, you may want to configure and send a test report to make sure everything works correctly.

Items

This section contains reference information and technical notes about items, including parameter details for specific item keys, platform-specific details, and notes on special item configurations.

Please use the sidebar to find information relevant to your use case.

vm.memory.size parameters

Overview

This section provides some parameter details for the vm.memory.size[<mode>] agent item.

Parameters

The following parameters are available for this item:

- **active** - memory currently in use or very recently used, and so it is in RAM
- **anon** - memory not associated with a file (cannot be re-read from it)
- **available** - available memory, calculated differently depending on the platform (see the table below)
- **buffers** - cache for things like file system metadata
- **cached** - cache for various things
- **exec** - executable code, typically from a (program) file
- **file** - cache for contents of recently accessed files
- **free** - memory that is readily available to any entity requesting memory
- **inactive** - memory that is marked as not used
- **pavailable** - 'available' memory as percentage of 'total' (calculated as `available/total*100`)
- **pinned** - same as 'wired'
- **pused** - 'used' memory as percentage of 'total' (calculated as `used/total*100`)
- **shared** - memory that may be simultaneously accessed by multiple processes
- **slab** - total amount of memory used by the kernel to cache data structures for its own use
- **total** - total physical memory available
- **used** - used memory, calculated differently depending on the platform (see the table below)
- **wired** - memory that is marked to always stay in RAM. It is never moved to disk.

Warning:
Some of these parameters are platform-specific and might not be available on your platform. See Zabbix agent items for details.

Platform-specific calculation of **available** and **used**:

Platform	**”available”**	**”used”**
AIX	free + cached	real memory in use
FreeBSD	inactive + cached + free	active + wired + cached
HP UX	free	total - free
Linux<3.14	free + buffers + cached	total - free
Linux 3.14+ (also backported to 3.10 on RHEL 7)	/proc/meminfo, see ”MemAvailable” in Linux kernel documentation for details. Note that free + buffers + cached is no longer equal to ’available’ due to not all the page cache can be freed and low watermark being used in calculation.	total - free
NetBSD	inactive + execpages + file + free	total - free
OpenBSD	inactive + free + cached	active + wired
OSX	inactive + free	active + wired
Solaris	free	total - free
Win32	free	total - free

Attention:
The sum of *vm.memory.size[used]* and *vm.memory.size[available]* does not necessarily equal total. For instance, on FreeBSD:
 * Active, inactive, wired, cached memories are considered used, because they store some useful information.
 * At the same time inactive, cached, free memories are considered available, because these kinds of memories can be given instantly to processes that request more memory.
So inactive memory is both used and available simultaneously. Because of this, the *vm.memory.size[used]* item is designed for informational purposes only, while *vm.memory.size[available]* is designed to be used in triggers.

See also

1. Additional details about memory calculation in different OS

Minimum permission level for Windows agent items

Overview

When monitoring systems using an agent, a good practice is to obtain metrics from the host on which the agent is installed. To use the principle of least privilege, it is necessary to determine what metrics are obtained from the agent.

The table in this document allows you to select the minimum rights for guaranteed correct operation of Zabbix agent.

If a different user is selected for the agent to work, rather than ’LocalSystem’, then for the operation of agent as a Windows service, the new user must have the rights ”Log on as a service” from ”Local Policy→User Rights Assignment” and the right to create, write and delete the Zabbix agent log file. An Active Directory user must be added to the *Performance Monitor Users* group.

Note:
When working with the rights of an agent based on the ”minimum technically acceptable” group, prior provision of rights to objects for monitoring is required.

Common agent items supported on Windows

Item key	User group	
	Recommended	Minimum technically acceptable (functionality is limited)
agent.hostname	Guests	Guests
agent.ping	Guests	Guests
agent.variant	Guests	Guests
agent.version	Guests	Guests
log	Administrators	Guests
log.count	Administrators	Guests
logrt	Administrators	Guests
logrt.count	Administrators	Guests
net.dns	Guests	Guests
net.dns.perf	Guests	Guests

Item key	User group	
net.dns.record	Guests	Guests
net.if.discovery	Guests	Guests
net.if.in	Guests	Guests
net.if.out	Guests	Guests
net.if.total	Guests	Guests
net.tcp.listen	Guests	Guests
net.tcp.port	Guests	Guests
net.tcp.service	Guests	Guests
net.tcp.service.perf	Guests	Guests
net.udp.service	Guests	Guests
net.udp.service.perf	Guests	Guests
proc.num	Administrators	Guests
system.cpu.discovery	Performance Monitor Users	Performance Monitor Users
system.cpu.load	Performance Monitor Users	Performance Monitor Users
system.cpu.num	Guests	Guests
system.cpu.util	Performance Monitor Users	Performance Monitor Users
system.hostname	Guests	Guests
system.localtime	Guests	Guests
system.run	Administrators	Guests
system.sw.arch	Guests	Guests
system.swap.size	Guests	Guests
system.uname	Guests	Guests
system.uptime	Performance Monitor Users	Performance Monitor Users
vfs.dir.count	Administrators	Guests
vfs.dir.get	Administrators	Guests
vfs.dir.size	Administrators	Guests
vfs.file.cksum	Administrators	Guests
vfs.file.contents	Administrators	Guests
vfs.file.exists	Administrators	Guests
vfs.file.md5sum	Administrators	Guests
vfs.file.regexp	Administrators	Guests
vfs.file.regmatch	Administrators	Guests
vfs.file.size	Administrators	Guests
vfs.file.time	Administrators	Guests
vfs.fs.discovery	Administrators	Guests
vfs.fs.size	Administrators	Guests
vfs.fs.get	Administrators	Guests
vm.memory.size	Guests	Guests
web.page.get	Guests	Guests
web.page.perf	Guests	Guests
web.page.regexp	Guests	Guests
zabbix.stats	Guests	Guests

Windows-specific item keys

Item key	User group	
	Recommended	Minimum technically acceptable (functionality is limited)
eventlog	Event Log Readers	Guests
net.if.list	Guests	Guests
perf_counter	Performance Monitor Users	Performance Monitor Users
proc_info	Administrators	Guests
service.discovery	Guests	Guests
service.info	Guests	Guests
services	Guests	Guests
wmi.get	Administrators	Guests
vm.vmemory.size	Guests	Guests

Encoding of returned values

Zabbix server expects every returned text value in the UTF8 encoding. This is related to any type of checks: Zabbix agent, SSH, Telnet, etc.

Different monitored systems/devices and checks can return non-ASCII characters in the value. For such cases, almost all possible **zabbix** keys contain an additional item key parameter - **<encoding>**. This key parameter is optional but it should be specified if the returned value is not in the UTF8 encoding and it contains non-ASCII characters. Otherwise, the result can be unexpected and unpredictable.

A description of behavior with different database backends in such cases follows.

MySQL

If a value contains a non-ASCII character in non UTF8 encoding - this character and the following will be discarded when the database stores this value. No warning messages will be written to the zabbix_server.log.
 Relevant for at least MySQL version 5.1.61.

PostgreSQL

If a value contains a non-ASCII character in non UTF8 encoding - this will lead to a failed SQL query (PGRES_FATAL_ERROR:ERROR invalid byte sequence for encoding) and data will not be stored. An appropriate warning message will be written to the zabbix_server.log.
 Relevant for at least PostgreSQL version 9.1.3.

Large file support

Large file support, often abbreviated to LFS, is the term applied to the ability to work with files larger than 2 GB on 32-bit operating systems. Support for large files affects at least log file monitoring and all vfs.file.* items. Large file support depends on the capabilities of a system at Zabbix compilation time, but is completely disabled on a 32-bit Solaris due to its incompatibility with procfs and swapctl.

Sensor

Overview

Each sensor chip gets its own directory in the sysfs /sys/devices tree. To find all sensor chips, it is easier to follow the device symlinks from /sys/class/hwmon/hwmon*, where * is a real number (0,1,2,...).

The sensor readings are located either in /sys/class/hwmon/hwmon*/ directory for virtual devices, or in /sys/class/hwmon/hwmon*/dev directory for non-virtual devices. A file, named name, located inside hwmon* or hwmon*/device directories contains the name of the chip, which corresponds to the name of the kernel driver used by the sensor chip.

There is only one sensor reading value per file. The common scheme for naming the files that contain sensor readings inside any of the directories mentioned above is: <type><number>_<item>, where

- **type** - for sensor chips is "in" (voltage), "temp" (temperature), "fan" (fan), etc.
- **item** - "input" (measured value), "max" (high threshold), "min" (low threshold), etc.
- **number** - always used for elements that can be present more than once (usually starts from 1, except for voltages which start from 0). If files do not refer to a specific element they have a simple name with no number.

The information regarding sensors available on the host can be acquired using **sensors-detect** and **sensors** tools (lm-sensors package). **Sensors-detect** helps to determine which modules are necessary for available sensors. When modules are loaded the **sensors** program can be used to show the readings of all sensor chips. The labeling of sensor readings, used by this program, can be different from the common naming scheme (<type><number>_<item>):

- if there is a file called <type><number>_label, then the label inside this file will be used instead of <type><number><item> name;
- if there is no <type><number>_label file, then the program searches inside the /etc/sensors.conf (could be also /etc/sensors3.conf, or different) for the name substitution.

This labeling allows user to determine what kind of hardware is used. If there is neither <type><number>_label file nor label inside the configuration file the type of hardware can be determined by the name attribute (hwmon*/device/name). The actual names of sensors, which zabbix_agent accepts, can be obtained by running **sensors** program with -u parameter (sensors -u).

In the **sensors** program the available sensors are separated by the bus type (ISA adapter, PCI adapter, SPI adapter, Virtual device, ACPI interface, HID adapter).

Obtaining sensor identifiers

Sensor labels, as printed by the sensors command, cannot always be used directly because labeling varies between chip vendors. For example, sensors output might contain the following lines:

```
sensors
in0:          +2.24 V  (min =  +0.00 V, max =  +3.32 V)
Vcore:        +1.15 V  (min =  +0.00 V, max =  +2.99 V)
+3.3V:        +3.30 V  (min =  +2.97 V, max =  +3.63 V)
+12V:        +13.00 V  (min =  +0.00 V, max = +15.94 V)
M/B Temp:     +30.0°C  (low  = -127.0°C, high = +127.0°C)
```

Out of these, only one display name corresponds directly to an internal identifier:

```
zabbix_get -s 127.0.0.1 -k sensor[lm85-i2c-0-2e,in0]
2.240000
```

Attempting to use other display labels (like `Vcore` or `+12V`) will not work:

```
zabbix_get -s 127.0.0.1 -k sensor[lm85-i2c-0-2e,Vcore]
ZBX_NOTSUPPORTED
```

To find the internal sensor identifier that Zabbix accepts, run `sensors -u`. In the output, the following may be observed:

```
sensors -u
...
Vcore:
  in1_input: 1.15
  in1_min: 0.00
  in1_max: 2.99
  in1_alarm: 0.00
...
+12V:
  in4_input: 13.00
  in4_min: 0.00
  in4_max: 15.94
  in4_alarm: 0.00
...
```

So the `Vcore` label corresponds to identifier `in1`, and `+12V` corresponds to `in4`. According to the kernel hwmon sysfs specification, these are voltages on chip pins and generally speaking may need scaling.

```
zabbix_get -s 127.0.0.1 -k sensor[lm85-i2c-0-2e,in1]
1.301000
```

Not only voltage (`in`), but also current (`curr`), temperature (`temp`) and fan speed (`fan`) readings can be retrieved by Zabbix.

Compatibility

Legacy distributions

Sensor readings are obtained from `/proc/sys/dev/sensors` directory:

- **device** - device name (if <mode> is used, it is a regular expression)
- **sensor** - sensor name (if <mode> is used, it is a regular expression)
- **mode** - possible values: `avg`, `max`, `min` (if this parameter is omitted, device and sensor are treated verbatim)

Example key: `sensor[w83781d-i2c-0-2d,temp1]`.

Current distributions

Sensor readings are obtained from `/sys/class/hwmon` directory:

- **device** - device name (non-regular expression). The device name could be the actual name of the device (e.g `0000:00:18.3`) or the name acquired using sensors program (e.g., `k8temp-pci-00c3`). It is up to the user to choose which name to use.
- **sensor** - sensor name (non-regular expression).
- **mode** - possible values: `avg`, `max`, `min` (if this parameter is omitted, device and sensor are treated verbatim).

Example keys:

`sensor[k8temp-pci-00c3,temp,max]` or `sensor[0000:00:18.3,temp1]`.

`sensor[smsc47b397-isa-0880,in,avg]` or `sensor[smsc47b397.2176,in1]`.

Notes on memtype parameter in proc.mem items

Overview

The **memtype** parameter is supported on Linux, AIX, FreeBSD, and Solaris platforms.

Three common values of memtype are supported on all of these platforms: pmem, rss and vsize. Additionally, platform-specific memtype values are supported on some platforms.

AIX

See values supported for memtype parameter on AIX in the table.

Supported value	Description	Source in procentry64 structure	Tries to be compatible with
vsize [1]	Virtual memory size	pi_size	
pmem	Percentage of real memory	pi_prm	ps -o pmem
rss	Resident set size	pi_trss + pi_drss	ps -o rssize
size	Size of process (code + data)	pi_dvm	ps gvw SIZE column
dsize	Data size	pi_dsize	
tsize	Text (code) size	pi_tsize	ps gvw TSIZ column
sdsize	Data size from shared library	pi_sdsize	
drss	Data resident set size	pi_drss	
trss	Text resident set size	pi_trss	

Notes for AIX:

1. When choosing parameters for proc.mem[] item key on AIX, try to specify narrow process selection criteria. Otherwise, there is a risk of getting unwanted processes counted into proc.mem[] result.

Example:

```
$ zabbix_agentd -t proc.mem[,,,NonExistingProcess,rss]
proc.mem[,,,NonExistingProcess,rss]           [u|2879488]
```

This example shows that specifying only the command line (regular expression to match) parameter results in Zabbix agent self-accounting, which is probably not what you want.

2. Do not use ps -ef to browse processes — it shows only non-kernel processes. Use ps -Af to see all processes which will be seen by Zabbix agent.
3. Let's go through example of topasrec how Zabbix agent proc.mem[] selects processes.

```
$ ps -Af | grep topasrec
root 10747984        1   0   Mar 16      -  0:00 /usr/bin/topasrec  -L -s 300 -R 1 -r 6 -o /var/perf daily
```

proc.mem[] has arguments:

proc.mem[<name>,<user>,<mode>,<cmdline>,<memtype>]

where: - The first criterion is a process name (argument <name>). In our example Zabbix agent will see it as topasrec. In order to match, you need to either specify topasrec or to leave it empty. - The second criterion is a user name (argument <user>). To match, you need to either specify root or to leave it empty. - The third criterion used in process selection is an argument <cmdline>. Zabbix agent will see its value as /usr/bin/topasrec -L -s 300 -R 1 -r 6 -o /var/perf/daily/ -ypersistent=1 -O type=bin -ystart_time=04:08:54,Mar16,2023. To match, you need to either specify a regular expression which matches this string or to leave it empty.

Arguments <mode> and <memtype> are applied after using the three criteria mentioned above.

FreeBSD

See values supported for memtype parameter on FreeBSD in the table.

Supported value	Description	Source in kinfo_proc structure	Tries to be compatible with
vsize	Virtual memory size	kp_eproc.e_vm.vm_map.size or ki_size	ps -o vsz
pmem	Percentage of real memory	calculated from rss	ps -o pmem
rss	Resident set size	kp_eproc.e_vm.vm_rssize or ki_rssize	ps -o rss
size [1]	Size of process (code + data + stack)	tsize + dsize + ssize	
tsize	Text (code) size	kp_eproc.e_vm.vm_tsize or ki_tsize	ps -o tsiz
dsize	Data size	kp_eproc.e_vm.vm_dsize or ki_dsize	ps -o dsiz
ssize	Stack size	kp_eproc.e_vm.vm_ssize or ki_ssize	ps -o ssiz

Linux

See values supported for memtype parameter on Linux in the table.

Supported value	Description	Source in /proc/<pid>/status file
vsize [1]	Virtual memory size	VmSize
pmem	Percentage of real memory	(VmRSS/total_memory) * 100
rss	Resident set size	VmRSS
data	Size of data segment	VmData
exe	Size of code segment	VmExe
hwm	Peak resident set size	VmHWM
lck	Size of locked memory	VmLck
lib	Size of shared libraries	VmLib
peak	Peak virtual memory size	VmPeak
pin	Size of pinned pages	VmPin
pte	Size of page table entries	VmPTE
size	Size of process code + data + stack segments	VmExe + VmData + VmStk
stk	Size of stack segment	VmStk
swap	Size of swap space used	VmSwap

Notes for Linux:

1. Not all memtype values are supported by older Linux kernels. For example, Linux 2.4 kernels do not support hwm, pin, peak, pte and swap values.
2. We have noticed that self-monitoring of the Zabbix agent active check process with proc.mem[...,...,...,...,data] shows a value that is 4 kB larger than reported by VmData line in the agent's /proc/<pid>/status file. At the time of self-measurement the agent's data segment increases by 4 kB and then returns to the previous size.

Solaris

See values supported for memtype parameter on Solaris in the table.

Supported value	Description	Source in psinfo structure	Tries to be compatible with
vsize [1]	Size of process image	pr_size	ps -o vsz
pmem	Percentage of real memory	pr_pctmem	ps -o pmem
rss	Resident set size It may be underestimated (see rss description in man ps).	pr_rssize	ps -o rss

Footnotes

[1] Default value.

Notes on selecting processes in proc.mem and proc.num items

Processes modifying their commandline

Some programs use modifying their commandline as a method for displaying their current activity. A user can see the activity by running ps and top commands. Examples of such programs include *PostgreSQL*, *Sendmail*, *Zabbix*.

Let's see an example from Linux. Let's assume we want to monitor a number of Zabbix agent processes.

ps command shows processes of interest as

```
$ ps -fu zabbix
UID        PID  PPID  C STIME TTY          TIME CMD
...
zabbix    6318     1  0 12:01 ?        00:00:00 sbin/zabbix_agentd -c /home/zabbix/zabbix_agentd.conf
zabbix    6319  6318  0 12:01 ?        00:00:01 sbin/zabbix_agentd: collector [idle 1 sec]
zabbix    6320  6318  0 12:01 ?        00:00:00 sbin/zabbix_agentd: listener #1 [waiting for connection]
zabbix    6321  6318  0 12:01 ?        00:00:00 sbin/zabbix_agentd: listener #2 [waiting for connection]
zabbix    6322  6318  0 12:01 ?        00:00:00 sbin/zabbix_agentd: listener #3 [waiting for connection]
zabbix    6323  6318  0 12:01 ?        00:00:00 sbin/zabbix_agentd: active checks #1 [idle 1 sec]
...
```

Selecting processes by name and user does the job:

```
$ zabbix_get -s localhost -k 'proc.num[zabbix_agentd,zabbix]'
6
```

Now let's rename zabbix_agentd executable to zabbix_agentd_30 and restart it.

ps now shows

```
$ ps -fu zabbix
UID        PID  PPID  C STIME TTY          TIME CMD
...
zabbix    6715     1  0 12:53 ?        00:00:00 sbin/zabbix_agentd_30 -c /home/zabbix/zabbix_agentd.conf
zabbix    6716  6715  0 12:53 ?        00:00:00 sbin/zabbix_agentd_30: collector [idle 1 sec]
zabbix    6717  6715  0 12:53 ?        00:00:00 sbin/zabbix_agentd_30: listener #1 [waiting for connection
zabbix    6718  6715  0 12:53 ?        00:00:00 sbin/zabbix_agentd_30: listener #2 [waiting for connection
zabbix    6719  6715  0 12:53 ?        00:00:00 sbin/zabbix_agentd_30: listener #3 [waiting for connection
zabbix    6720  6715  0 12:53 ?        00:00:00 sbin/zabbix_agentd_30: active checks #1 [idle 1 sec]
...
```

Now selecting processes by name and user produces an incorrect result:

```
$ zabbix_get -s localhost -k 'proc.num[zabbix_agentd_30,zabbix]'
1
```

Why a simple renaming of executable to a longer name lead to quite different result ?

Zabbix agent starts with checking the process name. /proc/<pid>/status file is opened and the line Name is checked. In our case the Name lines are:

```
$ grep Name /proc/{6715,6716,6717,6718,6719,6720}/status
/proc/6715/status:Name:   zabbix_agentd_3
/proc/6716/status:Name:   zabbix_agentd_3
/proc/6717/status:Name:   zabbix_agentd_3
/proc/6718/status:Name:   zabbix_agentd_3
/proc/6719/status:Name:   zabbix_agentd_3
/proc/6720/status:Name:   zabbix_agentd_3
```

The process name in status file is truncated to 15 characters.

A similar result can be seen with ps command:

```
$ ps -u zabbix
  PID TTY          TIME CMD
...
 6715 ?        00:00:00 zabbix_agentd_3
 6716 ?        00:00:01 zabbix_agentd_3
 6717 ?        00:00:00 zabbix_agentd_3
 6718 ?        00:00:00 zabbix_agentd_3
```

```
 6719 ?        00:00:00 zabbix_agentd_3
 6720 ?        00:00:00 zabbix_agentd_3
 ...
```

Obviously, that is not equal to our `proc.num[]` `name` parameter value `zabbix_agentd_30`. Having failed to match the process name from `status` file the Zabbix agent turns to `/proc/<pid>/cmdline` file.

How the agent sees the "cmdline" file can be illustrated with running a command

```
$ for i in 6715 6716 6717 6718 6719 6720; do cat /proc/$i/cmdline | awk '{gsub(/\x0/,"<NUL>"); print};'; d
sbin/zabbix_agentd_30<NUL>-c<NUL>/home/zabbix/zabbix_agentd.conf<NUL>
sbin/zabbix_agentd_30: collector [idle 1 sec]<NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><
sbin/zabbix_agentd_30: listener #1 [waiting for connection]<NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><N
sbin/zabbix_agentd_30: listener #2 [waiting for connection]<NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><N
sbin/zabbix_agentd_30: listener #3 [waiting for connection]<NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><N
sbin/zabbix_agentd_30: active checks #1 [idle 1 sec]<NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL><NUL
```

`/proc/<pid>/cmdline` files in our case contain invisible, non-printable null bytes, used to terminate strings in *C* language. The null bytes are shown as "<NUL>" in this example.

Zabbix agent checks "cmdline" for the main process and takes a `zabbix_agentd_30`, which matches our `name` parameter value `zabbix_agentd_30`. So, the main process is counted by item `proc.num[zabbix_agentd_30,zabbix]`.

When checking the next process, the agent takes `zabbix_agentd_30: collector [idle 1 sec]` from the `cmdline` file and it does not meet our `name` parameter `zabbix_agentd_30`. So, only the main process which does not modify its commandline, gets counted. Other agent processes modify their command line and are ignored.

This example shows that the `name` parameter cannot be used in `proc.mem[]` and `proc.num[]` for selecting processes in this case.

Note:
For `proc.get[]` item, when Zabbix agent checks "cmdline" for the process name, it will only use part of the name starting from the last slash and until the first space or colon sign. Process name received from cmdline file will only be used if its beginning completely matches the shortened process name in the `status` file. The algorithm is the same for both process name in the filter and in the JSON output.

Using `cmdline` parameter with a proper regular expression produces a correct result:

```
$ zabbix_get -s localhost -k 'proc.num[,zabbix,,zabbix_agentd_30[ :]]'
6
```

Be careful when using `proc.get[]`, `proc.mem[]` and `proc.num[]` items for monitoring programs which modify their command lines.

Before putting `name` and `cmdline` parameters into `proc.get[]`, `proc.mem[]` and `proc.num[]` items, you may want to test the parameters using `proc.num[]` item and `ps` command.

Linux kernel threads

Threads cannot be selected with `cmdline` parameter in `proc.get[]`, `proc.mem[]` and `proc.num[]` items

Let's take as an example one of kernel threads:

```
$ ps -ef| grep kthreadd
root         2     0  0 09:33 ?        00:00:00 [kthreadd]
```

It can be selected with process `name` parameter:

```
$ zabbix_get -s localhost -k 'proc.num[kthreadd,root]'
1
```

But selection by process `cmdline` parameter does not work:

```
$ zabbix_get -s localhost -k 'proc.num[,root,,kthreadd]'
0
```

The reason is that Zabbix agent takes the regular expression specified in `cmdline` parameter and applies it to contents of process `/proc/<pid>/cmdline`. For kernel threads their `/proc/<pid>/cmdline` files are empty. So, `cmdline` parameter never matches.

Counting of threads in `proc.mem[]` and `proc.num[]` items

Linux kernel threads are counted by `proc.num[]` item but do not report memory in `proc.mem[]` item. For example:

```
$ ps -ef | grep kthreadd
root         2     0  0 09:51 ?        00:00:00 [kthreadd]

$ zabbix_get -s localhost -k 'proc.num[kthreadd]'
1

$ zabbix_get -s localhost -k 'proc.mem[kthreadd]'
ZBX_NOTSUPPORTED: Cannot get amount of "VmSize" memory.
```

But what happens if there is a user process with the same name as a kernel thread ? Then it could look like this:

```
$ ps -ef | grep kthreadd
root         2     0  0 09:51 ?        00:00:00 [kthreadd]
zabbix    9611  6133  0 17:58 pts/1    00:00:00 ./kthreadd

$ zabbix_get -s localhost -k 'proc.num[kthreadd]'
2

$ zabbix_get -s localhost -k 'proc.mem[kthreadd]'
4157440
```

`proc.num[]` counted both the kernel thread and the user process. `proc.mem[]` reports memory for the user process only and counts the kernel thread memory as if it was 0. This is different from the case above when ZBX_NOTSUPPORTED was reported.

Be careful when using `proc.mem[]` and `proc.num[]` items if the program name happens to match one of the thread.

Before putting parameters into `proc.mem[]` and `proc.num[]` items, you may want to test the parameters using `proc.num[]` item and `ps` command.

Implementation details of net.tcp.service and net.udp.service checks

Implementation of net.tcp.service and net.udp.service checks is detailed on this page for various services specified in the service parameter.

Item net.tcp.service parameters

ftp

Creates a TCP connection and expects the first 4 characters of the response to be "220 ", then sends "QUIT\r\n". Default port 21 is used if not specified.

http

Creates a TCP connection without expecting and sending anything. Default port 80 is used if not specified.

https

Uses (and only works with) libcurl, does not verify the authenticity of the certificate, does not verify the host name in the SSL certificate, only fetches the response header (HEAD request). Default port 443 is used if not specified.

imap

Creates a TCP connection and expects the first 4 characters of the response to be "* OK", then sends "a1 LOGOUT\r\n". Default port 143 is used if not specified.

ldap

Opens a connection to an LDAP server and performs an LDAP search operation with filter set to (objectClass=*). Expects successful retrieval of the first attribute of the first entry. Default port 389 is used if not specified.

nntp

Creates a TCP connection and expects the first 3 characters of the response to be "200" or "201", then sends "QUIT\r\n". Default port 119 is used if not specified.

pop

Creates a TCP connection and expects the first 3 characters of the response to be "+OK", then sends "QUIT\r\n". Default port 110 is used if not specified.

smtp

Creates a TCP connection and expects the first 3 characters of the response to be ”220”, followed by a space, the line ending or a dash. The lines containing a dash belong to a multiline response and the response will be re-read until a line without the dash is received. Then sends ”QUIT\r\n”. Default port 25 is used if not specified.

ssh

Creates a TCP connection. If the connection has been established, both sides exchange an identification string (SSH-major.minor-XXXX), where major and minor are protocol versions and XXXX is a string. Zabbix checks if the string matching the specification is found and then sends back the string ”SSH-major.minor-zabbix_agent\r\n” or ”0\n” on mismatch. Default port 22 is used if not specified.

tcp

Creates a TCP connection without expecting and sending anything. Unlike the other checks requires the port parameter to be specified.

telnet

Creates a TCP connection and expects a login prompt (’:’ at the end). Default port 23 is used if not specified.

Item net.udp.service parameters

ntp

Sends an SNTP packet over UDP and validates the response according to RFC 4330, section 5. Default port 123 is used if not specified.

proc.get parameters

Overview

The item **proc.get**[<name>,<user>,<cmdline>,<mode>] is supported on Linux, Windows, FreeBSD, OpenBSD, and NetBSD.

List of process parameters returned by the item varies depending on the operating system and ’mode’ argument value.

Linux

The following process parameters are returned on Linux for each mode:

mode=process	mode=thread	mode=summary
pid: PID	pid: PID	name: process name
ppid: parent PID	ppid: parent PID	processes: number of processes
name: process name	name: process name	vsize: virtual memory size
cmdline: command with arguments	user: user (real) the process runs under	pmem: percentage of real memory
user: user (real) the process runs under	group: group (real) the process runs under	rss: resident set size
group: group (real) the process runs under	uid: user ID	data: size of data segment
uid: user ID	gid: ID of the group the process runs under	exe: size of code segment
gid: ID of the group the process runs under	tid: thread ID	lib: size of shared libraries
vsize: virtual memory size	tname: thread name	lck: size of locked memory
pmem: percentage of real memory	cputime_user: total CPU time spent in user mode (value expressed in clock ticks—divide by `getconf CLK_TCK` or `sysconf(_SC_CLK_TCK)` to obtain seconds)	pin: size of pinned pages
rss: resident set size	cputime_system: total CPU time spent in system mode (value expressed in clock ticks—divide by `getconf CLK_TCK` or `sysconf(_SC_CLK_TCK)` to obtain seconds)	pte: size of page table entries
data: size of data segment	state: thread state	size: size of process code + data + stack segments

mode=process	mode=thread	mode=summary
exe: size of code segment	ctx_switches: number of context switches	stk: size of stack segment
hwm: peak resident set size	page_faults: number of page faults	swap: size of swap space used
lck: size of locked memory		cputime_user: total CPU time spent in user mode (value expressed in clock ticks—divide by `getconf CLK_TCK` or `sysconf(_SC_CLK_TCK)` to obtain seconds)
lib: size of shared libraries		cputime_system: total CPU time spent in system mode (value expressed in clock ticks—divide by `getconf CLK_TCK` or `sysconf(_SC_CLK_TCK)` to obtain seconds)
peak: peak virtual memory size		ctx_switches: number of context switches
pin: size of pinned pages		threads: number of threads
pte: size of page table entries		page_faults: number of page faults
size: size of process code + data + stack segments		pss: proportional set size memory
stk: size of stack segment		
swap: size of swap space used		
cputime_user: total CPU time spent in user mode (value expressed in clock ticks—divide by `getconf CLK_TCK` or `sysconf(_SC_CLK_TCK)` to obtain seconds)		
cputime_system: total CPU time spent in system mode (value expressed in clock ticks—divide by `getconf CLK_TCK` or `sysconf(_SC_CLK_TCK)` to obtain seconds)		
state: process state (transparently retrieved from procfs, long form)		
ctx_switches: number of context switches		
threads: number of threads		
page_faults: number of page faults		
pss: proportional set size memory		

BSD-based OS

The following process parameters are returned on FreeBSD, OpenBSD, and NetBSD for each mode:

mode=process	mode=thread	mode=summary
pid: PID	pid: PID	name: process name
ppid: parent PID	ppid: parent PID	processes: number of processes
jid: ID of jail (FreeBSD only)	jid: ID of jail (FreeBSD only)	vsize: virtual memory size
jname: name of jail (FreeBSD only)	jname: name of jail (FreeBSD only)	pmem: percentage of real memory (FreeBSD only)
name: process name	name: process name	rss: resident set size
cmdline: command with arguments	user: user (real) the process runs under	size: size of process (code + data + stack)
user: user (real) the process runs under	group: group (real) the process runs under	tsize: text (code) size
group: group (real) the process runs under	uid: user ID	dsize: data size
uid: user ID	gid: ID of the group the process runs under	ssize: stack size

mode=process	mode=thread	mode=summary
gid: ID of the group the process runs under	tid: thread ID	cputime_user: total CPU time spent in user mode (reported in seconds)
vsize: virtual memory size	tname: thread name	cputime_system: total CPU time spent in system mode (reported in seconds)
pmem: percentage of real memory (FreeBSD only)	cputime_user: total CPU time spent in user mode (reported in seconds)	ctx_switches: number of context switches
rss: resident set size	cputime_system: total CPU time spent in system mode (reported in seconds)	threads: number of threads (not supported for NetBSD)
size: size of process (code + data + stack)	state: thread state	stk: size of stack segment
tsize: text (code) size	ctx_switches: number of context switches	page_faults: number of page faults
dsize: data size	io_read_op: number of times the system had to perform input	fds: number of file descriptors (OpenBSD only)
ssize: stack size	io_write_op: number of times the system had to perform output	swap: size of swap space used
cputime_user: total CPU time spent in user mode (reported in seconds)		io_read_op: number of times the system had to perform input
cputime_system: total CPU time spent in system mode (reported in seconds)		io_write_op: number of times the system had to perform output
state: process state (disk sleep/running/sleeping/tracing stop/zombie/other)		
ctx_switches: number of context switches		
threads: number of threads (not supported for NetBSD)		
page_faults: number of page faults		
fds: number of file descriptors (OpenBSD only)		
swap: size of swap space used		
io_read_op: number of times the system had to perform input		
io_write_op: number of times the system had to perform output		

Windows

The following process parameters are returned on Windows for each mode:

mode=process	mode=thread	mode=summary
pid: PID	pid: PID	name: process name
ppid: parent PID	ppid: parent PID	processes: number of processes
name: process name	name: process name	vmsize: virtual memory size
user: user the process runs under	user: user the process runs under	wkset: size of process working set
sid: user SID	sid: user SID	cputime_user: total CPU time spent in user mode (reported in milliseconds)
vmsize: virtual memory size	tid: thread ID	cputime_system: total CPU time spent in system mode (reported in milliseconds)
wkset: size of process working set		threads: number of threads
cputime_user: total CPU time spent in user mode (reported in milliseconds)		page_faults: number of page faults
cputime_system: total CPU time spent in system mode (reported in milliseconds)		handles: number of handles
threads: number of threads		io_read_b: IO bytes read
page_faults: number of page faults		io_write_b: IO bytes written
handles: number of handles		io_read_op: IO read operations
io_read_b: IO bytes read		io_write_op: IO write operations

mode=process	mode=thread	mode=summary
io_write_b: IO bytes written		io_other_b: IO bytes transferred, other than read and write operations
io_read_op: IO read operations		io_other_op: IO operations, other than read and write operations
io_write_op: IO write operations		
io_other_b: IO bytes transferred, other than read and write operations		
io_other_op: IO operations, other than read and write operations		

Unreachable/unavailable host interface settings

Overview

Several configuration parameters define how Zabbix server should behave when an agent check (Zabbix, SNMP, IPMI, JMX) fails and a host interface becomes unreachable.

Unreachable interface

A host interface is treated as unreachable after a failed check (network error, timeout) by Zabbix, SNMP, IPMI or JMX agents. Since Zabbix 6.2.0, active Zabbix agent checks also affect interface availability. When active checks become unavailable, they contribute to the overall availability status of the agent interface.

From the moment an interface becomes unreachable, **UnreachableDelay** defines how often it is rechecked using one of the items (including LLD rules). These rechecks are performed by unreachable pollers (or IPMI pollers for IPMI checks). By default, the interval between consecutive reachability checks is 15 seconds.

Attention:
Checks performed by asynchronous pollers are not moved to unreachable pollers.

In the Zabbix server log unreachability is indicated by messages like:

```
Zabbix agent item "system.cpu.load[percpu,avg1]" on host "New host" failed: first network error, wait for
Zabbix agent item "system.cpu.load[percpu,avg15]" on host "New host" failed: another network error, wait f
```

The log messages specify the exact item that failed and its type (Zabbix agent).

Note:
The *Timeout* parameter will also affect how early an interface is rechecked during unreachability. If Timeout is set to 20 seconds and UnreachableDelay to 30 seconds, the next check will be in 50 seconds after the first attempt.

The **UnreachablePeriod** parameter defines the total duration of the unreachability period. By default, UnreachablePeriod is 45 seconds. This value should be several times larger than UnreachableDelay to ensure that an interface is rechecked multiple times before being marked as unavailable.

An internal item, zabbix[host,active_agent,available], allows monitoring the availability of active checks in unreachable scenarios.

Switching interface back to available

When the unreachability period is over, the interface is polled again, decreasing priority for item that turned the interface into unreachable state. If the unreachable interface reappears, the monitoring returns to normal automatically:

```
resuming Zabbix agent checks on host "New host": connection restored
```

Note:
Once interface becomes available, the host does not poll all its items immediately for two reasons:
- It might overload the host.
- The interface restore time is not always matching planned item polling schedule time.

So, after the interface becomes available, items are not polled immediately, but they are getting rescheduled to their next polling round.

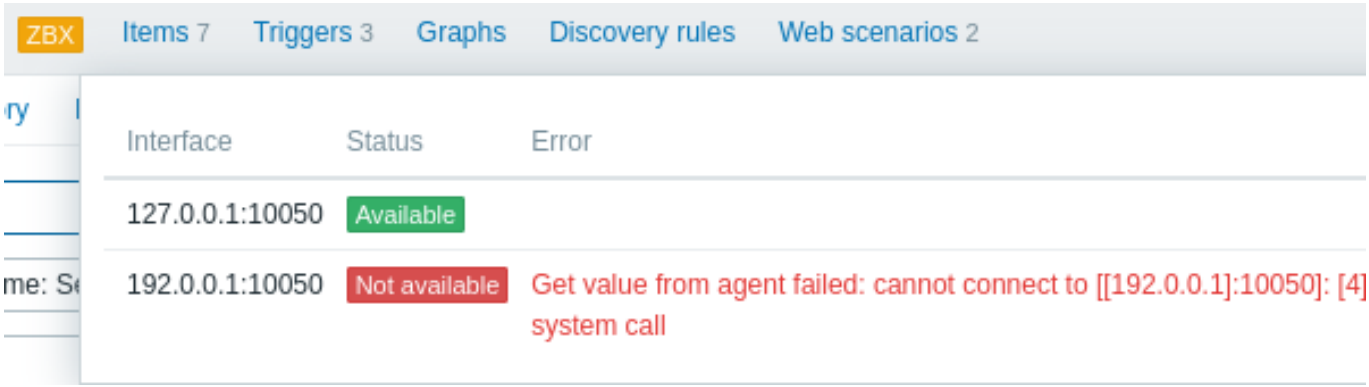
Unavailable interface

After the UnreachablePeriod ends and the interface has not reappeared, the interface is treated as unavailable.

In the server log it is indicated by messages like these:

```
temporarily disabling Zabbix agent checks on host "New host": interface unavailable
```

and in the frontend the host availability icon goes from green/gray to yellow/red (the unavailable interface details can be seen in the hint box that is displayed when a mouse is positioned on the host availability icon):



The **UnavailableDelay** parameter defines how often an interface is checked during interface unavailability.

By default, it is 60 seconds (so in this case "temporarily disabling", from the log message above, will mean disabling checks for one minute).

When the connection to the interface is restored, the monitoring returns to normal automatically, too:

```
enabling Zabbix agent checks on host "New host": interface became available
```

Remote monitoring of Zabbix stats

Overview

It is possible to make some internal metrics of Zabbix server and proxy accessible remotely by another Zabbix instance or a third-party tool. This can be useful so that supporters/service providers can monitor their client Zabbix servers/proxies remotely or, in organizations where Zabbix is not the main monitoring tool, that Zabbix internal metrics can be monitored by a third-party system in an umbrella-monitoring setup.

Zabbix internal stats are exposed to a configurable set of addresses listed in the new 'StatsAllowedIP' server/proxy parameter. Requests will be accepted only from these addresses.

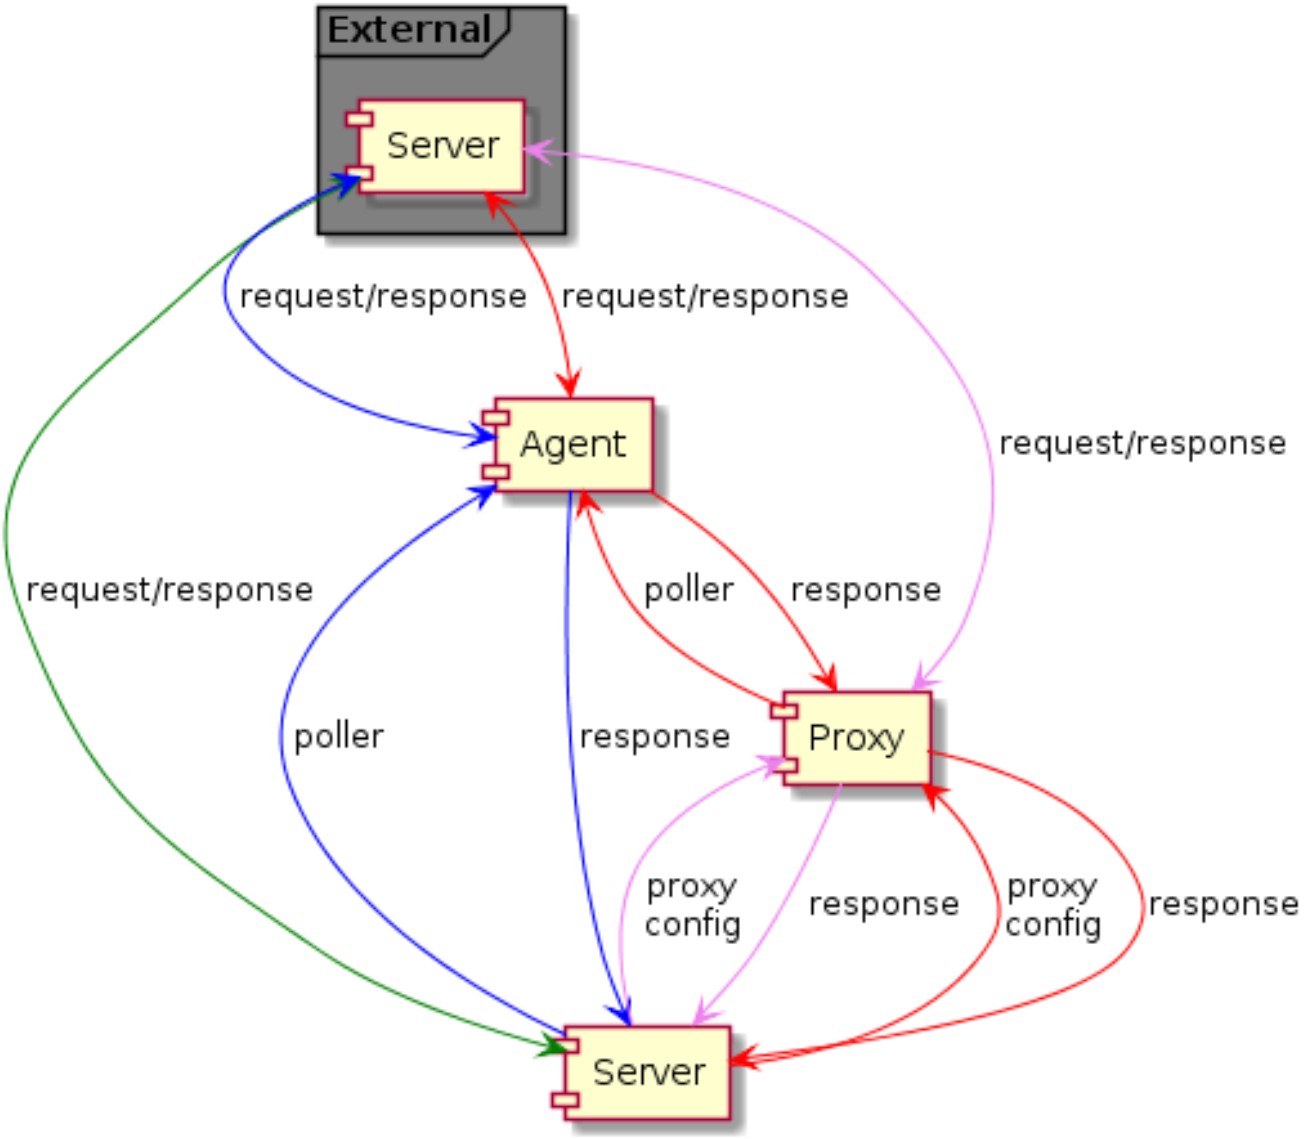
Items

To configure querying of internal stats on another Zabbix instance, you may use two items:

- `zabbix[stats,<ip>,<port>]` internal item - for direct remote queries of Zabbix server/proxy. <ip> and <port> are used to identify the target instance.
- `zabbix.stats[<ip>,<port>]` agent item - for agent-based remote queries of Zabbix server/proxy. <ip> and <port> are used to identify the target instance.

See also: Internal items, Zabbix agent items

The following diagram illustrates the use of either item depending on the context.



- - Server → external Zabbix instance (`zabbix[stats,<ip>,<port>]`)
- - Server → proxy → external Zabbix instance (`zabbix[stats,<ip>,<port>]`)
- - Server → agent → external Zabbix instance (`zabbix.stats[<ip>,<port>]`)
- - Server → proxy → agent → external Zabbix instance (`zabbix.stats[<ip>,<port>]`)

To make sure that the target instance allows querying it by the external instance, list the address of the external instance in the 'StatsAllowedIP' parameter on the target instance.

Exposed metrics

The stats items gather the statistics in bulk and return a JSON, which is the basis for dependent items to get their data from. The following internal metrics are returned by either of the two items:

- `zabbix[boottime]`
- `zabbix[hosts]`
- `zabbix[items]`
- `zabbix[items_unsupported]`
- `zabbix[preprocessing]` (server only)
- `zabbix[preprocessing_queue]` (server only)
- `zabbix[process,<type>,<mode>,<state>]` (only process type based statistics)
- `zabbix[rcache,<cache>,<mode>]`
- `zabbix[requiredperformance]`
- `zabbix[triggers]` (server only)
- `zabbix[uptime]`
- `zabbix[vcache,buffer,<mode>]` (server only)

- zabbix[vcache,cache,<parameter>]
- zabbix[version]
- zabbix[vmware,buffer,<mode>]
- zabbix[wcache,<cache>,<mode>] ('trends' cache type server only)

Templates

Templates are available for remote monitoring of Zabbix server or proxy internal metrics from an external instance:

- Remote Zabbix server health
- Remote Zabbix proxy health

Note that in order to use a template for remote monitoring of multiple external instances, a separate host is required for each external instance monitoring.

Trapper process

Receiving internal metric requests from an external Zabbix instance is handled by the trapper process that validates the request, gathers the metrics, creates the JSON data buffer and sends the prepared JSON back, for example, from server:

```
{
  "response": "success",
  "data": {
    "boottime": N,
    "uptime": N,
    "hosts": N,
    "items": N,
    "items_unsupported": N,
    "preprocessing": {
        "queued": {
          "count": N,
          "size": N
        },
        "direct": {
          "count": N,
          "size": N
        },
        "queue": N
    },
    "preprocessing_queue": N,
    "process": {
      "alert manager": {
        "busy": {
          "avg": N,
          "max": N,
          "min": N
        },
        "idle": {
          "avg": N,
          "max": N,
          "min": N
        },
        "count": N
      },
...
    },
    "queue": N,
    "rcache": {
      "total": N,
      "free": N,
      "pfree": N,
      "used": N,
      "pused": N
    },
    "requiredperformance": N,
    "triggers": N,
```

```
    "uptime": N,
    "vcache": {
      "buffer": {
        "total": N,
        "free": N,
        "pfree": N,
        "used": N,
        "pused": N
      },
      "cache": {
        "requests": N,
        "hits": N,
        "misses": N,
        "mode": N
      }
    },
    "vmware": {
      "total": N,
      "free": N,
      "pfree": N,
      "used": N,
      "pused": N
    },
    "version": "N",
    "wcache": {
      "values": {
        "all": N,
        "float": N,
        "uint": N,
        "str": N,
        "log": N,
        "text": N,
        "not supported": N
      },
      "history": {
        "pfree": N,
        "free": N,
        "total": N,
        "used": N,
        "pused": N
      },
      "index": {
        "pfree": N,
        "free": N,
        "total": N,
        "used": N,
        "pused": N
      },
      "trend": {
        "pfree": N,
        "free": N,
        "total": N,
        "used": N,
        "pused": N
      }
    }
  }
}
```


Internal queue items

There are also two other items that specifically allow remote querying of internal queue statistics on another Zabbix instance:

- `zabbix[stats,<ip>,<port>,queue,<from>,<to>]` internal item - for direct internal queue queries to remote Zabbix

server/proxy
- `zabbix.stats[<ip>,<port>,queue,<from>,<to>]` agent item - for agent-based internal queue queries to remote Zabbix server/proxy

See also: Internal items, Zabbix agent items

Configuring Kerberos with Zabbix

Overview

Kerberos authentication can be used in web monitoring and HTTP items in Zabbix.

This page describes an example of configuring Kerberos for Zabbix server to perform web monitoring of `www.example.com` with a Kerberos principal for Zabbix process on Debian/Ubuntu.

Configuration

1. Install KDC and client utilities:

```
sudo apt update
sudo apt install krb5-kdc krb5-admin-server krb5-user
```

During package setup answer the prompts, e.g.:

```
Default Kerberos version 5 realm: EXAMPLE.COM
Kerberos servers for your realm: localhost (or your FQDN)
Administrative server for your Kerberos realm: localhost (or your FQDN)
```

2. Map a friendly hostname (optional, for local testing).

Edit /etc/hosts and add an entry for your DC and webserver if you don't have DNS:

```
sudo vi /etc/hosts
```

Example line you might add:

```
192.168.1.100  dc01.example.com dc01
```

3. Configure the Kerberos client and KDC realm:

```
sudo vi /etc/krb5.conf
```

Example settings:

```
[libdefaults]
    default_realm = EXAMPLE.COM
    dns_lookup_realm = false
    dns_lookup_kdc = false
    rdns = false
    ticket_lifetime = 24h
    renew_lifetime = 7d
    forwardable = true

[realms]
    EXAMPLE.COM = {
        kdc = dc01.example.com
        admin_server = dc01.example.com
    }

[domain_realm]
    .example.com = EXAMPLE.COM
    example.com = EXAMPLE.COM
```

If you plan to use `.localdomain` or other non-public names, add explicit domain→realm mappings so hostname→realm mapping works. Mismatches here cause `Server not found in Kerberos database` errors.

4. Initialize the Kerberos database (one-time, KDC host). Set a secure master password when prompted:

```
sudo krb5_newrealm
```

5. Create the `HTTP/host.fqdn@REALM` principal using the exact hostname clients will use; prefer lowercase (e.g., `HTTP/web.example.com@`
A case/name mismatch causes `Server not found in Kerberos database`.

```
sudo kadmin.local
```

Inside `kadmin.local`:

```
addprinc kerb-admin@EXAMPLE.COM      # administrative principal
addprinc -randkey HTTP/dc01.example.com@EXAMPLE.COM
ktadd -k /etc/apache2/http.keytab HTTP/dc01.example.com@EXAMPLE.COM
quit
```

Move the keytab to the web host (or keep local if same machine) and set permissions usable by Apache:

```
chown www-data:www-data /etc/apache2/http.keytab
chmod 600 /etc/apache2/http.keytab
#### verify
sudo -u www-data -k /etc/apache2/http.keytab
```

6. Install and enable Apache GSSAPI module:

```
sudo apt install libapache2-mod-auth-gssapi
sudo a2enmod auth_gssapi
sudo a2enmod headers
sudo systemctl restart apache2
```

> **Attention:**
> Not all mod_auth_gssapi versions support every `Gssapi*` directive. If Apache fails with `Invalid command 'GssapiCredStore'` remove the unsupported directive or upgrade the module.

7. Configure a VirtualHost (adjust `DocumentRoot` / path to your Zabbix UI):

```
sudo vi /etc/apache2/sites-available/zabbix.conf
```

Inside `zabbix.conf`:

```
<VirtualHost *:80>
    ServerName dc01.example.com
    DocumentRoot /usr/share/zabbix/ui
    <Directory /usr/share/zabbix/ui>
        Options FollowSymLinks
        AllowOverride None
        Require all granted
        AuthType GSSAPI
        AuthName "Kerberos Login"
        GssapiCredStore keytab:/etc/apache2/http.keytab
        GssapiLocalName On
        Require valid-user
    </Directory>
    RequestHeader set X-Remote-User %{REMOTE_USER}s env=REMOTE_USER
    RequestHeader unset Authorization
</VirtualHost>
```

Restart Apache:

```
sudo systemctl restart apache2
```

8. Enable/start KDC services and verify listening ports (KDC host):

```
sudo systemctl enable --now krb5-kdc krb5-admin-server
ss -tnlp | grep :80    # or: sudo netstat -tnlp | grep :80
```

9. Obtain a TGT for testing (run as the user that will use the ticket).

Expect to see `krbtgt/EXAMPLE.COM@EXAMPLE.COM` in the ticket list. Run `kinit` as the same OS user that needs the ticket (e.g., `zabbix` for web checks or `www-data`/Apache for interactive browser SSO tests). Tickets issued to a different OS user won't be visible unless `KRB5CCNAME` and permissions are adjusted.

```
kinit kerb-admin@EXAMPLE.COM
klist
```

10. Test SPNEGO exchange with curl (from a client with a valid TGT). A `200 OK` (or redirect to app) indicates SPNEGO succeeded:

```
curl -v --negotiate -u : http://dc01.example.com/
```

11. Optionally, if Zabbix UI should accept HTTP-authenticated logins, enable HTTP authentication in the Zabbix frontend.

In the frontend configuration file (`zabbix.conf.php`) set:

```
$ZBX_FEATURE_FLAGS['http_auth_enabled'] = true;
```

Then configure HTTP authentication in the web interface.

12. Browser configuration (Firefox is used as an example): set `network.negotiate-auth.trusted-uris` to the host(s) performing Negotiate (`dc01.example.com`) so the browser will send Kerberos tokens automatically.

Inside `about:config`:

```
network.negotiate-auth.trusted-uris = dc01.example.com
```

Now visiting `http://dc01.example.com` should log you straight into Zabbix without the form.

13. Keep keys/tickets fresh. Default Kerberos ticket lifetime is approximately 10 hours. Add a cron/systemd timer to avoid expirations:

```
####for the web service
kinit -kt /etc/apache2/http.keytab HTTP/dc01.example.com@EXAMPLE.COM
####for the monitoring user
kinit -kt /var/lib/zabbix/kerb.keytab kerb-admin@EXAMPLE.COM
```

14. Housekeeping checks:

- `klist -k /etc/apache2/http.keytab` - verify service principal present in keytab.
- `sudo tail -f /var/log/apache2/error.log` - watch for GSSAPI errors (`gss_acquire_cred[_from]() failed to get server creds` means keytab/permissions or missing principal).
- `curl --negotiate` returning 401/403 often means wrong principal, no ticket, host header mismatch, or filesystem permission issue; check logs and `/etc/krb5.conf` domain mappings.

Security and file-permission notes

Keytab files must be readable only by the account that needs them. Example permissions: `0400` owned by `zabbix:zabbix` for a zabbix user keytab, or `0440` and `root:www-data` for an Apache keytab.

Avoid storing long-lived plaintext passwords on the host. Use keytabs or domain-joined machine principals where possible.

When running tests or scripts that set `KRB5CCNAME` or copy keytabs, double-check ownership and permissions after the operation—a webserver rejecting credentials is commonly a file-permission problem.

modbus.get parameters

Overview

The table below presents details of the `modbus.get` item parameters.

Parameters

Parameter	Description	Defaults	Example
endpoint	Protocol and address of the endpoint, defined as `protocol://connection_string` Possible protocol values: *rtu*, *ascii* (Agent 2 only), *tcp* Connection string format: with *tcp* - `address:port` with serial line: *rtu*, *ascii* - `port_name:speed:params` where 'speed' - 1200, 9600 etc 'params' - data bits (5,6,7 or 8), parity (n,e or o for none/even/odd), stop bits (1 or 2)	protocol: none *rtu/ascii* protocol: port_name: none speed: 115200 params: 8n1 *tcp* protocol: address: none port: 502	tcp://192.168.6.1:511 tcp://192.168.6.2 tcp://[::1]:511 tcp://::1 tcp://localhost:511 tcp://localhost rtu://COM1:9600:8n ascii://COM2:1200:7o2 rtu://ttyS0:9600 ascii://ttyS1

Parameter	Description	Defaults	Example
slave id	Modbus address of the device it is intended for (1 to 247), see MODBUS Messaging Implementation Guide (page 23) tcp device (not GW) will ignore the field	serial: 1 tcp: 255 (0xFF)	2
function	Empty or value of a supported function: 1 - Read Coil, 2 - Read Discrete Input, 3 - Read Holding Registers, 4 - Read Input Registers	empty	3
address	Address of the first registry, coil or input. If 'function' is empty, then 'address' should be in range for: Coil - 00001 - 09999 Discrete input - 10001 - 19999 Input register - 30001 - 39999 Holding register - 40001 - 49999 If 'function' is not empty, the 'address' field will be from 0 till 65535 and used without modification (PDU)	empty function: 00001 non-empty function: 0	9999
count	Count of sequenced 'type' which will be read from device, where: for Coil or Discrete input the 'type' = 1 bit for other cases: (count*sizeof(type))/2 = real count of registers for reading If 'offset' is not 0, the value will be added to 'real count' Acceptable range for 'real count' is 1:65535	1	2
type	Data type: for Read Coil and Read Discrete Input - *bit* for Read Holding Registers and Read Input Registers: *int8* - 8bit *uint8* - 8bit (unsigned) *int16* - 16bit *uint16* - 16bit (unsigned) *int32* - 32bit *uint32* - 32bit (unsigned) *float* - 32bit *uint64* - 64bit (unsigned) *double* - 64bit	bit uint16	uint64
endianness	Endianness type: *be* - Big Endian *le* - Little Endian *mbe* - Mid-Big Endian *mle* - Mid-Little Endian Limitations: for 1 bit - be for 8 bits - be,le for 16 bits - be,le	be	le

Parameter	Description	Defaults	Example
offset	Number of registers, starting from 'address', the result of which will be discarded. The size of each register is 16bit (needed to support equipment that does not support random read access).	0	4

Creating custom performance counter names for VMware

Overview

The VMware performance counter path has the `group/counter[rollup]` format where:

- `group` - the performance counter group, for example, *cpu*
- `counter` - the performance counter name, for example, *usagemhz*
- `rollup` - the performance counter rollup type, for example, *average*

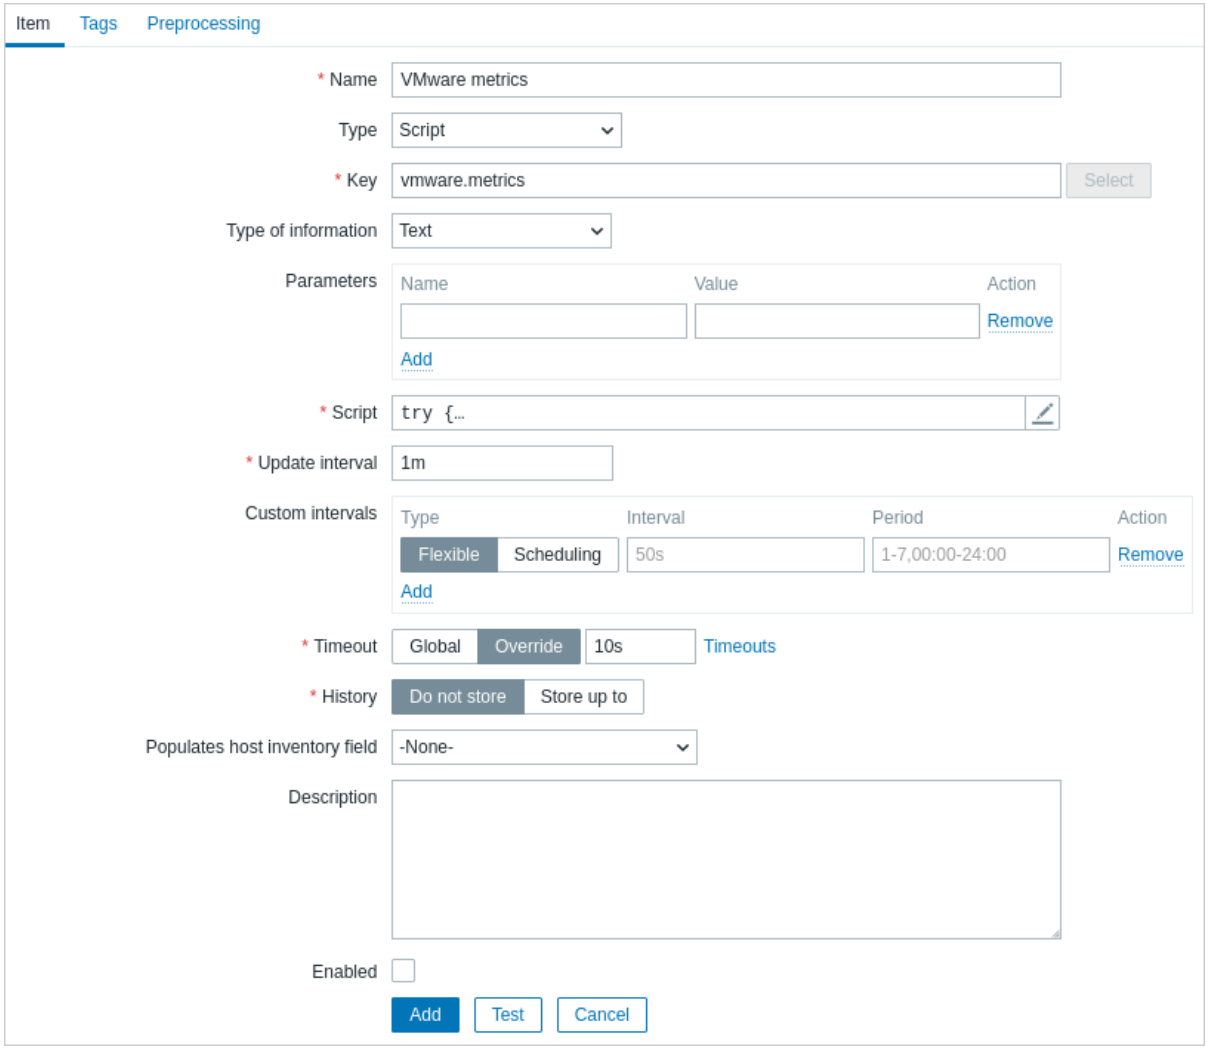
So the above example would give the following counter path: `cpu/usagemhz[average]`

The performance counter group descriptions, counter names and rollup types can be found in VMware documentation.

It is possible to obtain internal names and create custom performance counter names by using script item in Zabbix.

Configuration

1. Create disabled Script item on the main VMware host (where the **eventlog[]** item is present) with the following parameters:



- *Name*: VMware metrics
- *Type*: Script
- *Key*: vmware.metrics
- *Type of information*: Text

- *Script*: copy and paste the script provided below
- *Timeout*: 10
- *History*: Do not store
- *Enabled*: unmarked

Script

```
try {
    Zabbix.log(4, 'vmware metrics script');

    var result, resp,
    req = new HttpRequest();
    req.addHeader('Content-Type: application/xml');
    req.addHeader('SOAPAction: "urn:vim25/6.0"');

    login = '<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/" xmlns:urn="urn:vi
    <soapenv:Header/>\
    <soapenv:Body>\
        <urn:Login>\
            <urn:_this type="SessionManager">SessionManager</urn:_this>\
            <urn:userName>{$VMWARE.USERNAME}</urn:userName>\
            <urn:password>{$VMWARE.PASSWORD}</urn:password>\
        </urn:Login>\
    </soapenv:Body>\
</soapenv:Envelope>'
    resp = req.post("{$VMWARE.URL}", login);
    if (req.getStatus() != 200) {
        throw 'Response code: '+req.getStatus();
    }

    query = '<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/" xmlns:urn="urn:vi
<soapenv:Header/>\
    <soapenv:Body>\
        <urn:RetrieveProperties>\
            <urn:_this type="PropertyCollector">propertyCollector</urn:_this>\
            <urn:specSet>\
                <urn:propSet>\
                   <urn:type>PerformanceManager</urn:type>\
                   <urn:pathSet>perfCounter</urn:pathSet>\
                </urn:propSet>\
                <urn:objectSet>\
                   <urn:obj type="PerformanceManager">PerfMgr</urn:obj>\
                </urn:objectSet>\
            </urn:specSet>\
        </urn:RetrieveProperties>\
    </soapenv:Body>\
</soapenv:Envelope>'
    resp = req.post("{$VMWARE.URL}", query);
    if (req.getStatus() != 200) {
        throw 'Response code: '+req.getStatus();
    }
    Zabbix.log(4, 'vmware metrics=' + resp);
    result = resp;

    logout = '<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/" xmlns:urn="urn:v
    <soapenv:Header/>\
    <soapenv:Body>\
        <urn:Logout>\
            <urn:_this type="SessionManager">SessionManager</urn:_this>\
        </urn:Logout>\
    </soapenv:Body>\
</soapenv:Envelope>'

    resp = req.post("{$VMWARE.URL}",logout);
```

```
    if (req.getStatus() != 200) {
        throw 'Response code: '+req.getStatus();
    }

} catch (error) {
    Zabbix.log(4, 'vmware call failed : '+error);
    result = {};
}

return result;
```

Once the item is configured, click the *Test* button, then *Get value*.



Copy received XML to any XML formatter and find the desired metric.

An example of XML for one metric:

```
<PerfCounterInfo xsi:type="PerfCounterInfo">
    <key>6</key>
    <nameInfo>
        <label>Usage in MHz</label>
        <summary>CPU usage in megahertz during the interval</summary>
        <key>usagemhz</key>
    </nameInfo>
    <groupInfo>
        <label>CPU</label>
        <summary>CPU</summary>
        <key>cpu</key>
    </groupInfo>
    <unitInfo>
        <label>MHz</label>
        <summary>Megahertz</summary>
        <key>megaHertz</key>
    </unitInfo>
    <rollupType>average</rollupType>
    <statsType>rate</statsType>
    <level>1</level>
    <perDeviceLevel>3</perDeviceLevel>
</PerfCounterInfo>
```

Use XPath to extract the counter path from received XML. For the example above, the XPath will be:

field	xPath	value
group	//groupInfo[../key=6]/key	cpu
counter	//nameInfo[../key=6]/key	usagemhz
rollup	//rollupType[../key=6]	average

Resulting performance counter path in this case is: `cpu/usagemhz[average]`

Return values for system.sw.packages.get

Overview

This section provides return value details for the `system.sw.packages.get` Zabbix agent item.

Details

The output of this item is an array of objects each containing the following properties:

- **name** - package name
- **manager** - package manager that reported this data (`rpm`, `dpkg`, `pacman`, `pkgtool`, or `portage`)
- **version** - package version
- **size** - uncompressed package size in bytes (if not available, set to 0)
- **arch** - package architecture
- **buildtime** - an object with 2 entries:
  - **timestamp** - UNIX timestamp when the package was built (if not available, set to 0)
  - **value** - human readable date and time when the package was built (if not available, set to empty string)
- **installtime** - an object with 2 entries:
  - **timestamp** - UNIX timestamp when the package was installed (if not available, set to 0)
  - **value** - human readable date and time when the package was installed (if not available, set to empty string)

Example:

```
[
    {
        "name": "util-linux-core",
        "manager": "rpm",
        "version": "2.37.4-3.el9",
        "size": 1296335,
        "arch": "x86_64",
        "buildtime": {
            "timestamp" : 1653552239,
            "value" : "Sep 20 01:39:40 2021 UTC"
        },
        "installtime": {
            "timestamp" : 1660780885,
            "value" : "Aug 18 00:01:25 2022 UTC"
        }
    },
    {
        "name": "xfonts-base",
        "manager": "dpkg",
        "version": "1:1.0.5",
        "size": 7337984,
        "arch": "all",
        "buildtime": {
            "timestamp": 0,
            "value": ""
        },
        "installtime": {
            "timestamp": 0,
            "value": ""
        }
    }
]
```

Return values for net.dns.get

Overview

This section provides return value details for the `net.dns.get` Zabbix agent 2 item.

Details

The output of this item is an object containing DNS record information based on the parameters provided in the item key.

For example, the `net.dns.get[,example.com]` item may return the following JSON of a refused query:

```
{
    "flags": [
        "RA"
    ],
```

```
    "query_time": "0.019030",
    "question_section": [
        {
            "qclass": "IN",
            "qname": "example.com.",
            "qtype": "SOA"
        }
    ],
    "response_code": "REFUSED",
    "zbx_error_code": 0
}
```


By specifying the IP address of the DNS server, the `net.dns.get[192.0.2.1,example.com]` item may return the following JSON:


```
{
    "answer_section": [
        {
            "class": "IN",
            "name": "example.com.",
            "rdata": {
                "expire": 1209600,
                "mbox": "noc.dns.example.org.",
                "minttl": 3600,
                "ns": "ns.example.org.",
                "refresh": 7200,
                "retry": 3600,
                "serial": 2022091378
            },
            "rdlength": 44,
            "ttl": 1205,
            "type": "SOA"
        }
    ],
    "flags": [
        "RA"
    ],
    "query_time": "0.029556",
    "question_section": [
        {
            "qclass": "IN",
            "qname": "example.com.",
            "qtype": "SOA"
        }
    ],
    "response_code": "NOERROR",
    "zbx_error_code": 0
}
```


If there is a connection problem, the `net.dns.get[192.0.2.1,example.com]` item may an error:


```
{
    "zbx_error_code": -1,
    "zbx_error_msg": "Communication error: read udp 192.0.2.1:12345->192.0.2.1:53: i/o timeout"
}
```


The following types of error codes are possible:

Scenario	"zbx_error_code"	"zbx_error_msg"
No errors and the DNS response was received and parsed.	0	
DNS is down.	-1	"Communication error"
Error occurs during JSON parsing	-2	"Received unexpected response"

With additional parameters, the `net.dns.get[192.0.2.1,example.com,ANY,5,5,tcp,"cdflag,rdflag,dnssec,nsid,edns0,aaf`

item may return the following JSON:


```
{
    "additional_section": [
        {
            "extended_rcode": 32768,
            "name": ".",
            "rdata": {
                "options": [
                {
                    "code": 0,
                    "nsid": "67 70 64 6e 73 2d 6c 70 70"
                }
                ]
        },
            "rdlength": 13,
            "type": "OPT",
            "udp_payload": 512
        }
    ],
    "answer_section": [
        {
            "class": "IN",
            "name": "example.com.",
            "rdata": {
                "a": "192.0.2.1"
            },
            "rdlength": 4,
            "ttl": 19308,
            "type": "A"
        },
        {
            "class": "IN",
            "name": "example.com.",
            "rdata": {
                "algorithm": 13,
                "expiration": 1704715951,
                "inception": 1702910624,
                "key_tag": 21021,
                "labels": 2,
                "orig_ttl": 86400,
                "signature": "HVBOBcJJQy0S08J3f8kviPj8UkEUj7wmyiMyQqPSWgQIY9SCEJ5plq6KuxJmtAek1txZWXDo+6tp
                "signer_name": "example.com.",
                "type_covered": "A"
            },
            "rdlength": 95,
            "ttl": 19308,
            "type": "RRSIG"
        }
    ],
    "flags": [
        "RD",
        "RA",
        "AD",
        "CD"
    ],
    "query_time": "0.058221",
    "question_section": [
        {
            "qclass": "IN",
            "qname": "example.com.",
            "qtype": "ANY"
        }
```

```
    ],
    "response_code": "NOERROR",
    "zbx_error_code": 0
}
```

See also

For more information about DNS records, see:

- Domain Names - Implementation and Specification
- Domain Name System (DNS) Parameters

Return values for net.if.get

Overview

This section provides return value details for the `net.if.get` item. The item is supported on Linux by Zabbix agent and Zabbix agent 2.

Details

This item returns detailed information on network interfaces whose names match the given regular expression. The output is structured as a JSON object containing `config` and `values` arrays.

The output contains the following properties:

- **name** - interface name, read from `/proc/net/dev`;
- **ifalias** - interface alias, read from `/sys/class/net/<name>/ifalias`;
- **mac** - interface MAC address, read from `/sys/class/net/<name>/address`;
- **type** - interface type (`physical`, `virtual`, `loopback`);
- **speed** - the link speed in Mbits/sec, Ethernet-specific, read from `/sys/class/net/<name>/speed`;
- **duplex** - interface latest or current duplex value (`half`, `full`), Ethernet-specific, read from `/sys/class/net/<name>/duplex`;
- **administrative_state** - configured state (`down`, `up`), read from `/sys/class/net/<name>/flags`;
- **operational_state** - interface RFC2863 operational state (`unknown`, `notpresent`, `down`, `lowerlayerdown`, `testing`, `dormant`, `up`), read from `sys/class/net/<name>/operstate`;
- **carrier** - physical link presence (0 - physical link is down, 1 - physical link is up), read from `/sys/class/net/<name>/carrier`;
- **carrier_changes** - total number of times link state changed from up to down and vice versa, read from `/sys/class/net/<name>/car`
- **carrier_up_count** - number of times link has become up, read from `/sys/class/net/<name>/carrier_up_count`;
- **carrier_down_count** - number of times link has become down, read from `/sys/class/net/<name>/carrier_down_count`;
- **in** - incoming traffic, read from `/proc/net/dev`. Contains interface counters:
  - **bytes** - number of good bytes;
  - **packets** - number of good packets;
  - **errors** - total number of bad packets;
  - **dropped** - number of packets received but not processed + count of packets missed by the host;
  - **overruns** - number of overflow events;
  - **frame** - number of packets with framing errors;
  - **compressed** - number of correctly received compressed packets;
  - **multicast** - multicast packets received.
- **out** - outgoing traffic, read from `/proc/net/dev`. Contains interface counters:
  - **bytes** - number of good bytes;
  - **packets** - number of good packets;
  - **errors** - total number of bad packets;
  - **dropped** - number of packets received but not processed + count of packets missed by the host;
  - **overruns** - number of overflow events;
  - **collisions** - number of collisions during packet transmissions;
  - **carrier** - number of frame transmission errors due to loss of carrier during transmission;
  - **compressed** - number of transmitted compressed packets.

Unavailable metrics are excluded from the JSON output. In the following example, no alias is configured for the "lo" interface, so the `ifalias` property is excluded.

Example:

```
{
    "config": [
        {
            "name": "lo",
```

```
            "mac": "00:00:00:00:00:00",
            "type": "loopback",
            "administrative_state": "up",
            "operational_state": "unknown",
            "carrier": 1
        },
        {
            "name": "enp3s0",
            "ifalias": "Wired Ethernet connection",
            "mac": "8c:16:45:e3:1d:58",
            "type": "physical",
            "duplex": "unknown",
            "administrative_state": "up",
            "operational_state": "down",
            "carrier": 0
        },
        {
            "name": "wlp5s0",
            "ifalias": "Wi-Fi network adapter",
            "mac": "34:e1:2d:72:67:c6",
            "type": "physical",
            "administrative_state": "up",
            "operational_state": "up",
            "carrier": 1
        }
    ],
    "values": [
        {
            "name": "lo",
            "mac": "00:00:00:00:00:00",
            "carrier": 1,
            "carrier_changes": 0,
            "carrier_up_count": 0,
            "carrier_down_count": 0,
            "in": {
                "bytes": 1040047,
                "packets": 10784,
                "errors": 0,
                "dropped": 0,
                "overruns": 0,
                "frame": 0,
                "compressed": 0,
                "multicast": 0
            },
            "out": {
                "bytes": 1040047,
                "packets": 10784,
                "errors": 0,
                "dropped": 0,
                "overruns": 0,
                "collisions": 0,
                "carrier": 0,
                "compressed": 0
            }
        },
        {
            "name": "enp3s0",
            "ifalias": "Wired Ethernet connection",
            "mac": "8c:16:45:e3:1d:58",
            "carrier": 0,
            "carrier_changes": 1,
            "carrier_up_count": 0,
```

```
            "carrier_down_count": 1,
            "in": {
                "bytes": 0,
                "packets": 0,
                "errors": 0,
                "dropped": 0,
                "overruns": 0,
                "frame": 0,
                "compressed": 0,
                "multicast": 0
            },
            "out": {
                "bytes": 0,
                "packets": 0,
                "errors": 0,
                "dropped": 0,
                "overruns": 0,
                "collisions": 0,
                "carrier": 0,
                "compressed": 0
            }
        },
        {
            "name": "wlp5s0",
            "ifalias": "Wi-Fi network adapter",
            "mac": "34:e1:2d:72:67:c6",
            "carrier": 1,
            "carrier_changes": 6,
            "carrier_up_count": 3,
            "carrier_down_count": 3,
            "in": {
                "bytes": 165537298,
                "packets": 216182,
                "errors": 0,
                "dropped": 291,
                "overruns": 0,
                "frame": 0,
                "compressed": 0,
                "multicast": 0
            },
            "out": {
                "bytes": 252574347,
                "packets": 246607,
                "errors": 0,
                "dropped": 0,
                "overruns": 0,
                "collisions": 0,
                "carrier": 0,
                "compressed": 0
            }
        }
    ]
}
```


See also

For more information on the sysfs ABI for network devices, see:

- Linux kernel documentation

Return values for vfs.dev.get

Overview

This section provides return value details for the `vfs.dev.get` item. The item is supported on Linux by Zabbix agent and Zabbix agent 2, and on AIX by Zabbix agent.

Details

This item returns detailed information on block devices whose names match the given regular expression.

The output is structured as a JSON object containing `config` and, depending on the <mode> parameter, `values` arrays.

The `config` array contains the properties that are available for the selected mode. The `values` array is returned only for modes that provide block device statistics.

Unavailable metrics are excluded from the JSON output.

The following sections describe the properties returned for each mode.

disks mode

`vfs.dev.get[disks,.*]` returns block device configuration data.

Properties:

- **name** - block device name;
- **devid** - block device identifier, for example, a WWN-based or kernel-generated device ID;
- **type** - block device type (`disk`, `partition`, `rom`);
- **path** - path to the block device node;
- **model** - device model;
- **serial** - device serial number;
- **wwn** - device World Wide Name (WWN);
- **size_bytes** - device size in bytes;
- **logical_block_size** - logical block size in bytes;
- **physical_block_size** - physical block size in bytes.

Example for Linux:


```
{
  "config": [
    {
      "name": "sda",
      "devid": "ata-ST2000LM015-2E8174_ZDZ3YHZY",
      "type": "disk",
      "path": "/dev/sda",
      "model": " ST2000LM015-2E81",
      "wwn": "naa.5000c500b2791423",
      "size_bytes": 2000398934016,
      "logical_block_size": 512,
      "physical_block_size": 4096
    },
    {
      "name": "sdb",
      "devid": "usb-Kingston_DataTraveler_3.0_E0D55E6C382C15A158991CBD-0:0",
      "type": "disk",
      "path": "/dev/sdb",
      "model": "DataTraveler 3.0",
      "size_bytes": 61991813632,
      "logical_block_size": 512,
      "physical_block_size": 512
    },
    {
      "name": "nvme0n1",
      "devid": "nvme-SAMSUNG_MZVLW256HEHP-000L7_S35ENX1K753633",
      "type": "disk",
      "path": "/dev/nvme0n1",
      "model": "SAMSUNG MZVLW256HEHP-000L7",
      "serial": "S35ENX1K753633",
      "wwn": "eui.002538b781b7dc55",
      "size_bytes": 256060514304,
      "logical_block_size": 512,
      "physical_block_size": 512
```

```
    }
  ]
}
```

Example for AIX:

```
{
  "config": [
    {
      "name": "cd0",
      "type": "disk",
      "path": "/dev/cd0",
      "size_bytes": 0,
      "logical_block_size": 512
    },
    {
      "name": "hdisk1",
      "type": "disk",
      "devid": "48333321360050768OC818022F000000000015EA404214503IBMfcp05VDASD03AIXvscsi",
      "path": "/dev/hdisk1",
      "size_bytes": 51136000000,
      "logical_block_size": 512
    },
    {
      "name": "hdisk0",
      "type": "disk",
      "devid": "48333321360050768OC818022F000000000015EA304214503IBMfcp05VDASD03AIXvscsi",
      "path": "/dev/hdisk0",
      "size_bytes": 0,
      "logical_block_size": 512
    }
  ]
}
```

devices mode

`vfs.dev.get[devices,*]` returns block device configuration data, including partitions.

Properties:

- **name** - block device name.
- **devid** - block device identifier.
- **type** - block device type (`disk`, `partition`, `rom`).
- **partitions** - partitions on the device. Contains partition names and their sizes in bytes.

Example for Linux:

```
{
  "config": [
    {
      "name": "sda",
      "devid": "ata-ST2000LM015-2E8174_ZDZ3YHZY",
      "type": "disk",
      "partitions": {
        "sda1": 16759808,
        "sda2": 2000381018112
      }
    },
    {
      "name": "sdb",
      "devid": "usb-Kingston_DataTraveler_3.0_E0D55E6C382C15A158991CBD-0:0",
      "type": "disk",
      "partitions": {
        "sdb1": 27631878144,
        "sdb2": 34359768576,
        "sdb3": 32256
```

```
      }
    },
    {
      "name": "nvme0n1",
      "devid": "nvme-SAMSUNG_MZVLW256HEHP-000L7_S35ENX1K753633",
      "type": "disk",
      "partitions": {
        "nvme0n1p1": 209715200,
        "nvme0n1p2": 16777216,
        "nvme0n1p3": 255060869120,
        "nvme0n1p4": 770703360
      }
    }
  ]
}
```

disk_stats mode

`vfs.dev.get[disk_stats,.*]` returns block device configuration data and disk statistics.

Properties:

- **name** - block device name.
- **devid** - block device identifier.
- **type** - block device type (`disk`, `partition`, `rom`).
- **size_bytes** - device size in bytes.
- **logical_block_size** - logical block size in bytes.
- **stats** - block device I/O statistics. Contains:
  - **reads_completed** - number of completed read operations;
  - **writes_completed** - number of completed write operations;
  - **bytes_read** - number of bytes read;
  - **bytes_written** - number of bytes written;
  - **io_time_ms** - total time spent doing I/O, in milliseconds.

Example for Linux

```
{
  "config": [
    {
      "name": "sda",
      "devid": "ata-ST2000LM015-2E8174_ZDZ3YHZY",
      "type": "disk",
      "size_bytes": 61991813632
    },
    {
      "name": "sdb",
      "devid": "usb-Kingston_DataTraveler_3.0_E0D55E6C382C15A158991CBD-0:0",
      "type": "disk",
      "size_bytes": 61991813632
    },
    {
      "name": "nvme0n1",
      "devid": "nvme-SAMSUNG_MZVLW256HEHP-000L7_S35ENX1K753633",
      "type": "disk",
      "size_bytes": 256060514304
    }
  ],
  "values": [
    {
      "name": "sda",
      "stats": {
        "reads_completed": 871,
        "writes_completed": 0,
        "bytes_read": 13045760,
        "bytes_written": 0,
```

```
        "io_time_ms": 1256
      }
    },
    {
      "name": "sdb",
      "stats": {
        "reads_completed": 51017,
        "writes_completed": 10682,
        "bytes_read": 10514627,
        "bytes_written": 1530472,
        "io_time_ms": 284988
      }
    },
    {
      "name": "nvme0n1",
      "stats": {
        "reads_completed": 1259,
        "writes_completed": 7,
        "bytes_read": 49698,
        "bytes_written": 6,
        "io_time_ms": 243
      }
    }
  ]
}
```


Example for AIX:


```
{
  "config": [
    {
      "name": "cd0",
      "type": "disk",
      "path": "/dev/cd0",
      "size_bytes": 0,
      "logical_block_size": 2048
    },
    {
      "name": "hdisk1",
      "type": "disk",
      "devid": "4833332136005076 80C818022F000000000015EA404214503IBMfcp05VDASD03AIXvscsi",
      "path": "/dev/hdisk1",
      "size_bytes": 51136000000,
      "logical_block_size": 512
    },
    {
      "name": "hdisk0",
      "type": "disk",
      "devid": "4833332136005076 80C818022F000000000015EA304214503IBMfcp05VDASD03AIXvscsi",
      "path": "/dev/hdisk0",
      "size_bytes": 0,
      "logical_block_size": 512
    }
  ],
  "values": [
    {
      "name": "cd0",
      "stats": {
        "reads_completed": 21,
        "writes_completed": 0,
        "bytes_read": 75776,
        "bytes_written": 0,
        "io_time_ms": 0
```

```
      }
    },
    {
      "name": "hdisk1",
      "stats": {
        "reads_completed": 35888,
        "writes_completed": 1165147,
        "bytes_read": 593699840,
        "bytes_written": 10009646592,
        "io_time_ms": 69
      }
    },
    {
      "name": "hdisk0",
      "stats": {
        "reads_completed": 3,
        "writes_completed": 0,
        "bytes_read": 12800,
        "bytes_written": 0,
        "io_time_ms": 0
      }
    }
  ]
}
```

device_stats mode

`vfs.dev.get[device_stats,.*]` returns block device configuration data and device statistics.

Properties:

- **name** - block device name.
- **devid** - block device identifier.
- **type** - block device type (`disk`, `partition`, `rom`).
- **size_bytes** - device size in bytes.
- **stats** - block device I/O statistics. Contains:
  - **reads_completed** - number of completed read operations;
  - **writes_completed** - number of completed write operations;
  - **bytes_read** - number of bytes read;
  - **bytes_written** - number of bytes written;
  - **io_time_ms** - total time spent doing I/O, in milliseconds.

Example for Linux:

```
{
  "config": [
    {
      "name": "sda",
      "devid": "ata-ST2000LM015-2E8174_ZDZ3YHZY",
      "type": "disk",
      "size_bytes": 61991813632
    },
    {
      "name": "sda1",
      "type": "partition",
      "devid": "ata-ST2000LM015-2E8174_ZDZ3YHZY-part1",
      "size_bytes": 1234567890
    },
    {
      "name": "sda2",
      "type": "partition",
      "devid": "ata-ST2000LM015-2E8174_ZDZ3YHZY-part2",
      "size_bytes": 1234567890
    },
    {
```

```
      "name": "sdb",
      "devid": "usb-Kingston_DataTraveler_3.0_E0D55E6C382C15A158991CBD-0:0",
      "type": "disk",
      "size_bytes": 61991813632
    },
    {
      "name": "sdb1",
      "type": "partition",
      "devid": "usb-Kingston_DataTraveler_3.0_E0D55E6C382C15A158991CBD-0:0-part1",
      "size_bytes": 1234567890
    },
    {
      "name": "sdb2",
      "type": "partition",
      "devid": "usb-Kingston_DataTraveler_3.0_E0D55E6C382C15A158991CBD-0:0-part2",
      "size_bytes": 1234567890
    },
    {
      "name": "sdb3",
      "type": "partition",
      "devid": "usb-Kingston_DataTraveler_3.0_E0D55E6C382C15A158991CBD-0:0-part3",
      "size_bytes": 1234567890
    },
    {
      "name": "nvme0n1",
      "devid": "nvme-SAMSUNG_MZVLW256HEHP-000L7_S35ENX1K753633",
      "type": "disk",
      "size_bytes": 256060514304
    },
    {
      "name": "nvme0n1p1",
      "type": "partition",
      "devid": "nvme-SAMSUNG_MZVLW256HEHP-000L7_S35ENX1K753633-part1",
      "size_bytes": 1234567890
    },
    {
      "name": "nvme0n1p2",
      "type": "partition",
      "devid": "nvme-SAMSUNG_MZVLW256HEHP-000L7_S35ENX1K753633-part2",
      "size_bytes": 1234567890
    },
    {
      "name": "nvme0n1p3",
      "type": "partition",
      "devid": "nvme-SAMSUNG_MZVLW256HEHP-000L7_S35ENX1K753633-part3",
      "size_bytes": 1234567890
    },
    {
      "name": "nvme0n1p4",
      "type": "partition",
      "devid": "nvme-SAMSUNG_MZVLW256HEHP-000L7_S35ENX1K753633-part4",
      "size_bytes": 1234567890
    }
  ],
  "values": [
    {
      "name": "sda",
      "stats": {
        "reads_completed": 871,
        "writes_completed": 0,
        "bytes_read": 13045760,
        "bytes_written": 0,
```

```
      "io_time_ms": 1256
    }
  },
  {
    "name": "sda1",
    "stats": {
      "reads_completed": 123,
      "writes_completed": 123,
      "bytes_read": 123,
      "bytes_written": 123,
      "io_time_ms": 123
    }
  },
  {
    "name": "sda2",
    "stats": {
      "reads_completed": 123,
      "writes_completed": 123,
      "bytes_read": 123,
      "bytes_written": 123,
      "io_time_ms": 123
    }
  },
  {
    "name": "sdb",
    "stats": {
      "reads_completed": 51017,
      "writes_completed": 10682,
      "bytes_read": 10514627,
      "bytes_written": 1530472,
      "io_time_ms": 284988
    }
  },
  {
    "name": "sdb1",
    "stats": {
      "reads_completed": 123,
      "writes_completed": 123,
      "bytes_read": 123,
      "bytes_written": 123,
      "io_time_ms": 123
    }
  },
  {
    "name": "sdb2",
    "stats": {
      "reads_completed": 123,
      "writes_completed": 123,
      "bytes_read": 123,
      "bytes_written": 123,
      "io_time_ms": 123
    }
  },
  {
    "name": "sdb3",
    "stats": {
      "reads_completed": 123,
      "writes_completed": 123,
      "bytes_read": 123,
      "bytes_written": 123,
      "io_time_ms": 123
    }
```

```
    },
    {
      "name": "nvme0n1",
      "stats": {
        "reads_completed": 1259,
        "writes_completed": 7,
        "bytes_read": 49698,
        "bytes_written": 6,
        "io_time_ms": 243
      }
    },
    {
      "name": "nvme0n1p1",
      "stats": {
        "reads_completed": 123,
        "writes_completed": 123,
        "bytes_read": 123,
        "bytes_written": 123,
        "io_time_ms": 123
      }
    },
    {
      "name": "nvme0n1p2",
      "stats": {
        "reads_completed": 123,
        "writes_completed": 123,
        "bytes_read": 123,
        "bytes_written": 123,
        "io_time_ms": 123
      }
    },
    {
      "name": "nvme0n1p3",
      "stats": {
        "reads_completed": 123,
        "writes_completed": 123,
        "bytes_read": 123,
        "bytes_written": 123,
        "io_time_ms": 123
      }
    },
    {
      "name": "nvme0n1p4",
      "stats": {
        "reads_completed": 123,
        "writes_completed": 123,
        "bytes_read": 123,
        "bytes_written": 123,
        "io_time_ms": 123
      }
    }
  ]
}
```


Notes on system.cpu.util items on Windows

system.cpu.util item provides the CPU utilization percentage.

When a collector process is started on Zabbix agent for Windows, a buffer for N items for N logical processors (threads) is initialized. Values are periodically updated in the buffer using Windows performance counters. These values are reported when `system.cpu.util[n]` is requested, where 0 <= n < N.

One more item in the buffer is used for `system.cpu.util[total]`.

Zabbix automatically selects different performance counters to support both NUMA systems and older Windows versions with no processor group support.

If logical processor (thread) count is less or equal than 64 and number of processor groups is equal to 1, then:

```
N is logical processor (thread) count.

\Processor(_Total)\% Processor Time
\Processor(0)\% Processor Time
\Processor(1)\% Processor Time
\Processor(2)\% Processor Time
...
\Processor(N-1)\% Processor Time
```

Else:

```
G is processor group count and N is logical processor (thread) count in group.

\Processor Information(_Total)\% Processor Time
\Processor Information(0,0)\% Processor Time
\Processor Information(0,1)\% Processor Time
\Processor Information(0,2)\% Processor Time
...
\Processor Information(0,N-1)\% Processor Time
...
\Processor Information(G-1,0)\% Processor Time
\Processor Information(G-1,1)\% Processor Time
\Processor Information(G-1,2)\% Processor Time
...
\Processor Information(G-1,N-1)\% Processor Time
```

Large JSON value support

To ensure reliable handling of large values with the JSON data type (1MiB or more), review and apply the following configuration changes as needed:

1. If you're using MySQL or MariaDB, increase the maximum allowed packet size in `/etc/mysql/my.cnf` (it may default to 1MB on some systems):

```
[mysqld]
max_allowed_packet=128M

[mysqldump]
max_allowed_packet=1G
```

2. Adjust additional MySQL/MariaDB server system variables and InnoDB system variables as needed. For example:

```
innodb_io_capacity=1900
connect_timeout=600
wait_timeout=57600
interactive_timeout=57600
net_read_timeout=7200
net_write_timeout=7200
net-buffer-length=32704
```

3. Increase `HistoryCacheSize` to 2G in the Zabbix server configuration file.

4. If you're using Zabbix proxy, increase `HistoryCacheSize` to 2G and `ProxyMemoryBufferSize` to 2G in the proxy configuration file.

5. If you're using Zabbix agent, increase communication timeouts between Zabbix server/proxy and agent (e.g., item timeouts or the `Timeout` parameter value in the Zabbix server/proxy and agent configuration files). Otherwise, messages between them may not complete, and you may see errors such as *"message length does not match expected length"*. Also, review other timeouts as needed, such as the script timeout or item test timeout.

Supported functions

See Trigger expression and functions.

Macros

It is possible to use out-of-the-box Supported macros and User macros supported by location.

Macros supported by location

Overview

The macros supported by location table contains reference of all macros (built-in, expression, and user-definable) that are supported in the specific location.

If all macros from a related group are supported, then only the group name is given. Click each macro name or group name to see full details of the macros.

The macro details list contains a detailed list of all **built-in** macros, grouped by application area. Note that to customize macro values (for example, shorten or extract specific substrings), you can use macro functions.

Macros supported by location

Location	Supported macros	Comments
Action		
Time period condition, Default operation step duration, Step duration	User-definable macros (global only)	A single macro has to fill the whole field; multiple macros/mix with text are not supported
Notifications and commands		
Trigger-based	Action macros Date/time macros {ESC.HISTORY} Event macros, Cause/symptom event macros Event update macros: {EVENT.UPDATE.HISTORY}, {EVENT.UPDATE.STATUS} Expression macro: {?EXPRESSION} {FUNCTION.VALUE} Host macros: {HOST.CONN}, {HOST.DESCRIPTION}, {HOST.DNS}, {HOST.HOST}, {HOST.ID}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Target host macros (commands only) Host inventory macros Item macros (except {ITEM.STATE}, {ITEM.STATE.ERROR}) Proxy macros Trigger macros: {TRIGGER.DESCRIPTION}, {TRIGGER.EVENTS.*}, {TRIGGER.EXPRESSION}, {TRIGGER.EXPRESSION.EXPLAIN}, {TRIGGER.EXPRESSION.RECOVERY}, {TRIGGER.ID}, {TRIGGER.HOSTGROUP.NAME}, {TRIGGER.NAME}, {TRIGGER.NAME.ORIG}, {TRIGGER.NSEVERITY}, {TRIGGER.SEVERITY}, {TRIGGER.STATUS}, {TRIGGER.TEMPLATE.NAME}, {TRIGGER.URL}, {TRIGGER.URL.NAME}, {TRIGGER.VALUE} User-definable macros	See also: Using macros in messages

Location	Supported macros	Comments
Problem update	Action macros Date/time macros {ESC.HISTORY} Event macros, Cause/symptom event macros, Recovery event macros Event update macros: all (except {EVENT.UPDATE.NSEVERITY}, {EVENT.UPDATE.SEVERITY}) Expression macro: {?EXPRESSION} Function macros Host macros: {HOST.CONN}, {HOST.DESCRIPTION}, {HOST.DNS}, {HOST.HOST}, {HOST.ID}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Target host macros (commands only) Host inventory macros Item macros (except {ITEM.STATE}, {ITEM.STATE.ERROR}) Proxy macros Trigger macros: {TRIGGER.DESCRIPTION}, {TRIGGER.EVENTS.*}, {TRIGGER.EXPRESSION}, {TRIGGER.EXPRESSION.EXPLAIN}, {TRIGGER.EXPRESSION.RECOVERY}, {TRIGGER.ID}, {TRIGGER.HOSTGROUP.NAME}, {TRIGGER.NAME}, {TRIGGER.NAME.ORIG}, {TRIGGER.NSEVERITY}, {TRIGGER.SEVERITY}, {TRIGGER.STATUS}, {TRIGGER.TEMPLATE.NAME}, {TRIGGER.URL}, {TRIGGER.URL.NAME}, {TRIGGER.VALUE} {USER.FULLNAME} User-definable macros	
Problem recovery	Recovery event macros Event macros {FUNCTION.RECOVERY.VALUE} {TRIGGER.EXPRESSION.RECOVERY.EXPLAIN}	
Discovery	Action macros Date/time macros Discovery macros Event macros: {EVENT.AGE}, {EVENT.DATE}, {EVENT.ID}, {EVENT.OBJECT}, {EVENT.SOURCE}, {EVENT.TIME}, {EVENT.TIMESTAMP} Target host macros (commands only) Proxy macros	
Autoregistration	Action macros Date/time macros Event macros: {EVENT.AGE}, {EVENT.DATE}, {EVENT.ID}, {EVENT.OBJECT}, {EVENT.SOURCE}, {EVENT.TIME}, {EVENT.TIMESTAMP} Host macros: {HOST.HOST}, {HOST.IP}, {HOST.PORT}, {HOST.METADATA} Target host macros (commands only) Proxy macros	
Service-based	Action macros Date/time macros {ESC.HISTORY} Event macros: all except {EVENT.ACK.STATUS}, {EVENT.OPDATA} Service macros User-definable macros	
Service recovery	Event macros (except {EVENT.ACK.STATUS}, {EVENT.OPDATA}), Recovery event macros	

Location	Supported macros	Comments
Service update	Action macros Date/time macros {ESC.HISTORY} Event macros (except {EVENT.ACK.STATUS}, {EVENT.OPDATA}) Event update macros: {EVENT.UPDATE.DATE}, {EVENT.UPDATE.NSEVERITY}, {EVENT.UPDATE.SEVERITY}, {EVENT.UPDATE.STATUS}, {EVENT.UPDATE.TIME}, {EVENT.UPDATE.TIMESTAMP} Service macros User-definable macros	
Internal notifications	Action macros Date/time macros {ESC.HISTORY} Event macros (except {EVENT.NSEVERITY}, {EVENT.SEVERITY}), Recovery event macros Host macros: {HOST.CONN}, {HOST.DESCRIPTION}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Host inventory macros Item macros: {ITEM.DESCRIPTION}, {ITEM.DESCRIPTION.ORIG}, {ITEM.ID}, {ITEM.KEY}, {ITEM.KEY.ORIG}, {ITEM.NAME}, {ITEM.NAME.ORIG}, {ITEM.VALUETYPE} Proxy macros	
Item-based	Item macros: {ITEM.STATE}, {ITEM.STATE.ERROR}	
LLD rule-based	Low-level discovery notification macros	
Trigger-based	Trigger macros: {TRIGGER.DESCRIPTION}, {TRIGGER.EXPRESSION}, {TRIGGER.EXPRESSION.RECOVERY}, {TRIGGER.HOSTGROUP.NAME}, {TRIGGER.ID}, {TRIGGER.NAME}, {TRIGGER.NAME.ORIG}, {TRIGGER.NSEVERITY}, {TRIGGER.SEVERITY}, {TRIGGER.STATE}, {TRIGGER.STATE.ERROR}, {TRIGGER.TEMPLATE.NAME}, {TRIGGER.URL}, {TRIGGER.URL.NAME} User-definable macros	
Alert script parameters	Alert macros User-definable macros (global only)	
Connector		
URL, Username, Password, Bearer token, HTTP proxy, SSL certificate, SSL key file, SSL key password field	User-definable macros (global only)	
Timeout field	User-definable macros (global only)	A single macro has to fill the whole field; multiple macros/mix with text are not supported.

Location	Supported macros	Comments
Graph name	Expression macro: {?EXPRESSION}	Only a single **avg**, **last**, **max**, **min** function with seconds as parameter (time suffixes can be used) is supported in this macro in this location. The {HOST.HOST<1-9>} macro can be used as host within the macro (see also indexed macros). Example usage: {?avg(/{HOST.HOST}/item.key,1h)} Unsupported usage: {?last(/host/item1)/last(/host/item2)} {?last(/host/item1)*10} {?count(/host/item1,5m)}
Host, host prototype		
Interface IP/DNS	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME} User-definable macros [1]	[1] For IP address, a single macro has to fill the whole field; multiple macros/mix with text are not supported
Interface port	User-definable macros	A single macro has to fill the whole field; multiple macros/mix with text are not supported
SNMP community	User-definable macros	
SNMPv2,v3 Max repetition count	User-definable macros, LLD macros (for host prototypes)	
SNMP Retries	User-definable macros, LLD macros (for host prototypes)	
SNMPv3 Context name, Security name, Authentication passphrase, Privacy passphrase		
IPMI username, password		
Item, item prototype, LLD rule		
Name	User-definable macros	Not supported in LLD rule names.
Description	User-definable macros	

Location	Supported macros	Comments
Key parameters	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} User-definable macros	The supported {HOST.*} macros will resolve to the interface that is selected for the item. When used in items without interfaces they will resolve to either the Zabbix agent, SNMP, JMX or IPMI interface of the host in this order of priority or to 'UNKNOWN' if the host does not have any interface. Macros are not supported in the item key of Zabbix trapper items.
Update interval, *Custom intervals*, *Timeout (available for supported item types)*, *Store up to (for history/trends)* fields	User-definable macros	A single macro has to fill the whole field; multiple macros/mix with text are not supported.
Delete lost resources, *Disable lost resources* fields	User-definable macros	These fields are supported for LLD rules only. A single macro has to fill the whole field; multiple macros/mix with text are not supported.
Preprocessing step parameters (including custom scripts), *custom error-handling parameters* (*Set value to* and *Set error to* fields)	User-definable macros	
Filter regular expressions	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} User-definable macros	These fields are supported for LLD rules only.
Override filter regular expressions	User-definable macros	These fields are supported for LLD rules only.
Override operations: *Update interval*, *History storage period*, *Trend storage period* (for item prototypes)	User-definable macros	These fields are supported for LLD rules only. A single macro has to fill the whole field; multiple macros/mix with text are not supported.
Item types		

Location	Supported macros	Comments
Browser item *parameter names and values*	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Host inventory macros Item macros: {ITEM.ID}, {ITEM.KEY}, {ITEM.KEY.ORIG} User-definable macros	
Script item *parameter names and values*		
Calculated/aggregate item formula (*expression constants*, *function parameters*, *item key parameters*, (aggregate item only) *filter conditions* (host group name and tag name))	User-definable macros	
Database monitor item *Username, Password, SQL query* fields	User-definable macros	
HTTP agent item *URL, Query fields, Request body, Headers (names and values), SSL certificate file, SSL key file* fields	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Item macros: {ITEM.ID}, {ITEM.KEY}, {ITEM.KEY.ORIG} User-definable macros	URLs that contain a secret macro will not work, as the macro in them will be resolved as "******"
HTTP agent item *HTTP Proxy* field	Item macros: {ITEM.ID}, {ITEM.KEY}, {ITEM.KEY.ORIG} User-definable macros	
HTTP agent item *Required status codes, HTTP authentication username/password, SSL key password* fields	User-definable macros	
HTTP agent item *Allowed hosts* field	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME} User-definable macros	
JMX agent item *Endpoint* field	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.PORT} User-definable macros	
Trapper item *Allowed hosts* field	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP} User-definable macros	
SNMP agent item *SNMP OID* field	User-definable macros	
SSH and Telnet item *script*	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} User-definable macros	
SSH agent item *Username, Public key file, Private key file, Password* fields	User-definable macros	
Telnet agent item *Username, Password* fields	User-definable macros	
Map		

Location	Supported macros	Comments
URL, URL name	Host macros: {HOST.CONN}, {HOST.DESCRIPTION}, {HOST.DNS}, {HOST.HOST}, {HOST.ID}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} {HOSTGROUP.ID} Host inventory macros Map macros {TRIGGER.ID}	
Element label	Expression macro: {?EXPRESSION} [1] Host macros: {HOST.CONN}, {HOST.DESCRIPTION}, {HOST.DNS}, {HOST.HOST}, {HOST.ID}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} {HOSTGROUP.ID} Host inventory macros Map macros Trigger macros: {TRIGGER.EVENTS.*}, {TRIGGER.ID}, {TRIGGER.PROBLEM.EVENTS.*}, {TRIGGERS.ACK}, {TRIGGERS.PROBLEM.ACK}, {TRIGGERS.PROBLEM.UNACK}, {TRIGGERS.UNACK}	[1] Only a single **avg**, **last**, **max**, **min** function with seconds as parameter (time suffixes can be used) is supported in this macro in this location. The {HOST.HOST<1-9>} macro can be used as host within the macro. Example usage: {?avg(/{HOST.HOST}/item.key,1h)} Unsupported usage: {?last(/host/item1)/last(/host/item2)} {?last(/host/item1)*10} {?count(/host/item1,5m)}
Link label	Expression macro: {?EXPRESSION} [1]	
Shape label	Expression macro: {?EXPRESSION} [1]	
Text field in map shapes	{MAP.NAME}	
Network discovery		
Update interval field	User-definable macros (global only)	A single macro has to fill the whole field; multiple macros/mix with text are not supported.
SNMP community, SNMP OID fields *SNMPv3 Context name, Security name, Authentication passphrase, Privacy passphrase* fields	User-definable macros (global only)	
Proxy		
Address for active agents > *Port* (if the proxy belongs to a group) For passive proxy mode: interface *Address* and *Port* fields *Timeouts for item types* fields in override mode	User-definable macros (global only)	A single macro has to fill the whole field; multiple macros/mix with text are not supported.
Proxy group		
Failover period, Minimum number of proxies fields	User-definable macros (global only)	A single macro has to fill the whole field; multiple macros/mix with text are not supported.

Location	Supported macros	Comments
Scripts		
Commands, confirmation text (Type: script, manual **host action**)	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.ID} [1], {HOST.IP}, {HOST.NAME}, {HOST.PORT} Host inventory macros [2] {MANUALINPUT} Username macros User-definable macros	[1] Confirmation text only [2] Supported for Zabbix server and Zabbix proxy
Commands (Type: script, manual **event action**)	Date/time macros Event macros, Cause/symptom event macros, Recovery event macros (if recovery took place) Event update macros: {EVENT.UPDATE.HISTORY}, {EVENT.UPDATE.STATUS} {FUNCTION.VALUE} Host macros: {HOST.CONN}, {HOST.DESCRIPTION}, {HOST.DNS}, {HOST.HOST}, {HOST.ID}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Host inventory macros Item macros (except {ITEM.STATE}, {ITEM.STATE.ERROR}) {MANUALINPUT} Proxy macros Trigger macros: {TRIGGER.DESCRIPTION}, {TRIGGER.EVENTS.*}, {TRIGGER.EXPRESSION}, {TRIGGER.EXPRESSION.EXPLAIN}, {TRIGGER.EXPRESSION.RECOVERY}, {TRIGGER.EXPRESSION.RECOVERY.EXPLAIN}, {TRIGGER.HOSTGROUP.NAME}, {TRIGGER.ID}, {TRIGGER.NAME}, {TRIGGER.NAME.ORIG}, {TRIGGER.NSEVERITY}, {TRIGGER.SEVERITY}, {TRIGGER.STATUS}, {TRIGGER.TEMPLATE.NAME}, {TRIGGER.URL}, {TRIGGER.URL.NAME}, {TRIGGER.VALUE} Username macros User-definable macros	
Confirmation text (Type: script, manual **event action**)	Event macros: {EVENT.ID}, {EVENT.NAME}, {EVENT.NSEVERITY}, {EVENT.SEVERITY}, {EVENT.STATUS}, {EVENT.VALUE} Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.ID}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Host inventory macros {MANUALINPUT} Username macros User-definable macros	
URL, confirmation text (Type: URL, manual **host/event action**)		
Tag name and value	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.ID}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Host inventory macros Item macros: {ITEM.LASTVALUE*}, {ITEM.LOG.*}, {ITEM.VALUE*} Trigger macros: {TRIGGER.ID} (trigger tag value only) User-definable macros	Tag name and value macros are resolved during the event generation process only.
Trigger		
Name	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Item macros: {ITEM.LASTVALUE*}, {ITEM.LOG.*}, {ITEM.VALUE*} Positional macros/references: $1...$9 User-definable macros	

Location	Supported macros	Comments
Event name	Date/time macros: {TIME}, {TIMESTAMP} Expression macro: {?EXPRESSION} Function macros Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Item macros: {ITEM.LASTVALUE*}, {ITEM.LOG.*}, {ITEM.VALUE*} Trigger macros: {TRIGGER.EXPRESSION.EXPLAIN}	
Operational data	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Item macros: {ITEM.LASTVALUE*}, {ITEM.LOG.*}, {ITEM.VALUE*} User-definable macros	
Expression	{TRIGGER.VALUE} User-definable macros [1]	[1] Only in constants and function parameters; secret macros are not supported
Tag for matching, *Menu entry name*, *Menu entry URL* fields	User-definable macros	URLs that contain a secret macro will not work, as the macro in them will be resolved as "******".
Description	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Item macros: {ITEM.LASTVALUE*}, {ITEM.LOG.*}, {ITEM.VALUE*} User-definable macros	
URL	{EVENT.ID} Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.ID}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Item macros: {ITEM.LASTVALUE*}, {ITEM.LOG.*}, {ITEM.VALUE*} {TRIGGER.ID}	
User parameter commands	Positional macros/references: $1...$9	
Web scenario		
Name, *Agent*, *HTTP proxy* fields	User-definable macros	
Update interval field	User-definable macros	A single macro has to fill the whole field; multiple macros/mix with text are not supported.
Variables, *Headers*, *SSL certificate*, *SSL key file* fields	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} User-definable macros	
User, *Password*, *SSL key password* fields	User-definable macros	
Web scenario step *Name*, *Variables (values only)* fields	User-definable macros	

Location	Supported macros	Comments
Web scenario step *Timeout, Required status codes* fields	User-definable macros	A single macro has to fill the whole field; multiple macros/mix with text are not supported.
Web scenario step *URL, Headers (names and values), Required string* fields	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} User-definable macros	
Web scenario step *Post* field	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.IP}, {HOST.NAME}, {HOST.PORT}	
Webhook media type		
Parameter values	Alert macros All macros supported in trigger-based problem notifications (see above)	A single macro has to fill the whole field; multiple macros/mix with text are not supported.
JavaScript script, *JavaScript* script parameter name	User-definable macros (global only)	
Menu entry name, Menu entry URL	{EVENT.TAGS.<tag name>}	
Widgets		
Gauge widget *Description* field	Host macros: {HOST.CONN}, {HOST.DESCRIPTION}, {HOST.DNS}, {HOST.HOST}, {HOST.ID}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Host inventory macros Item macros (except {ITEM.STATE.ERROR}) User-definable macros (global only)	
Honeycomb widget *Primary/secondary label* field **Item value** widget *Description* field		
Top hosts widget *Text data* column	Host macros: {HOST.CONN}, {HOST.DESCRIPTION}, {HOST.DNS}, {HOST.HOST}, {HOST.ID}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} Host inventory macros User-definable macros (global only)	
URL (Dynamic) widget *URL* field	Host macros: {HOST.CONN}, {HOST.DNS}, {HOST.HOST}, {HOST.ID}, {HOST.IP}, {HOST.NAME}, {HOST.PORT} User-definable macros	
***Custom text* column**		
Monitoring > Problems	Host macros (except {HOST.METADATA}) Host inventory macros Item macros (except {ITEM.STATE.ERROR}) Trigger macros: {TRIGGER.ID} Event macros: {EVENT.ID} User-definable macros	
Monitoring > Hosts	Host macros (except {HOST.METADATA}) Host inventory macros User-definable macros	
Monitoring > Latest data	Host macros (except {HOST.METADATA}) Host inventory macros Item macros (except {ITEM.STATE.ERROR}) User-definable macros	

Location	Supported macros	Comments
Data collection > Templates	Host macros: {HOST.DESCRIPTION}, {HOST.HOST}, {HOST.ID}, {HOST.NAME}	
Data collection > Hosts	Host macros (except {HOST.METADATA}) Host inventory macros User-definable macros	
Other locations		
Working time field	User-definable macros (global only)	A single macro has to fill the whole field; multiple macros/mix with text are not supported.
Timeouts for item types field (global setting)		
User media *When active* field	User-definable macros (global only)	
Email media type *Username, Password* fields	User-definable macros (global only)	Secret macros are recommended.
Scheduled reports *Subject, Message* fields	Date/time macros: {TIME}, {TIMESTAMP}	

Macro details

This list contains details for each built-in macro grouped by application area:

- Action macros
- Alert macros
- Date/time macros
- Discovery macros
- Event macros
  - Event update macros
  - Cause/symptom event macros
  - Recovery event macros
- Escalation macros
- Function macros
- Host macros
  - Target host macros
- Host group macros
- Host inventory macros
- Item macros
- Low-level discovery notification macros
- Map macros
- Proxy macros
- Script macros
- Service macros
- Trigger macros
- Username macros

Action macros

{ACTION.ID}

The numeric ID of the triggered action.

{ACTION.NAME}

The name of the triggered action.

Alert macros

{ALERT.MESSAGE}

The *Default message* value from action configuration.

{ALERT.SENDTO}

The *Send to* value from user media configuration.

{ALERT.SUBJECT}

The *Default subject* value from action configuration.

Date and time macros

{DATE}

The current date in `yyyy.mm.dd` format.

{TIME}

The current time in `hh:mm:ss` format.

{TIMESTAMP}

The current time in UNIX timestamp format.

Discovery macros

{DISCOVERY.DEVICE.IPADDRESS}

The IP address of the discovered device.
 Available always, does not depend on host being added.

{DISCOVERY.DEVICE.DNS}

The DNS name of the discovered device.
 Available always, does not depend on host being added.

{DISCOVERY.DEVICE.STATUS}

The status of the discovered device (UP/DOWN).

{DISCOVERY.DEVICE.UPTIME}

The time since the last change of discovery status for a particular device.
 For example: 1h 29m 01s.
 For devices with status DOWN, this is the period of their downtime.

{DISCOVERY.RULE.NAME}

The name of the discovery rule that discovered the presence or absence of device/service.

{DISCOVERY.SERVICE.NAME}

The name of the discovered service.
 For example: HTTP.

{DISCOVERY.SERVICE.PORT}

The port of the discovered service.
 For example: 80.

{DISCOVERY.SERVICE.STATUS}

The status of the discovered service (UP/DOWN).

{DISCOVERY.SERVICE.UPTIME}

The time since the last change of discovery status for a particular service.
 For example: 1h 29m 01s.
 For services with status DOWN, this is the period of their downtime.

Escalation macros

{ESC.HISTORY}

The log of previously sent notifications, their escalation step and status (*sent*, *in progress* or *failed*).

Event macros

{EVENT.ACK.STATUS}

The acknowledgment status of the event (Yes/No).

{EVENT.AGE}

The age of the event that triggered an action, with precision down to a second.
 Useful in escalated messages. For example: 9m 13s.

{EVENT.DATE}

The date of the event that triggered an action in `yyyy.mm.dd` format.
 For example: 2025.04.14.

{EVENT.DURATION}

The duration of the event (time difference between problem and recovery events), with precision down to a second.
 Useful in problem recovery messages.

{EVENT.ID}

The numeric ID of the event that triggered an action.

{EVENT.NAME}

The name of the event that triggered an action.

{EVENT.NSEVERITY}

The numeric value of the event severity.
 Possible values: *0* - Not classified, *1* - Information, *2* - Warning, *3* - Average, *4* - High, *5* - Disaster.

{EVENT.OBJECT}

The numeric value of the event object.
 Possible values: *0* - Trigger, *1* - Discovered host, *2* - Discovered service, *3* - Autoregistration, *4* - Item, *5* - Low-level discovery rule.

{EVENT.OPDATA}

The operational data of the underlying trigger of the problem.

{EVENT.SEVERITY}

The name of the event severity.

{EVENT.SOURCE}

The numeric value of the event source.
 Possible values: *0* - Trigger, *1* - Discovery, *2* - Autoregistration, *3* - Internal, *4* - Service.

{EVENT.STATUS}

The verbal value of the event that triggered an action.

{EVENT.TAGS}

A comma-separated list of event tags.
Expanded to an empty string if no tags exist.

{EVENT.TAGSJSON}

A JSON array containing event tag objects.
Expanded to an empty array if no tags exist.

{EVENT.TAGS.<tag name>}

The event tag value referenced by the tag name.
 A tag name containing non-alphanumeric characters (including non-English multibyte-UTF characters) must be double quoted. Quotes and backslashes inside a quoted tag name must be escaped with a backslash.

{EVENT.TIME}

The time of the event that triggered an action in `hh:mm:ss` format.
 For example: 12:57:53.

{EVENT.TIMESTAMP}

The UNIX timestamp of the event that triggered an action.

{EVENT.VALUE}

The numeric value of the event that triggered an action.
 Possible values: *1* - problem, *0* - recovering.

Cause/symptom event macros

{EVENT.CAUSE.ACK.STATUS}

The acknowledgment status of the cause event (Yes/No).
 Used in the context of a symptom event.

{EVENT.CAUSE.AGE}

The age of the cause event, with precision down to a second.
 Useful in escalated messages.
 Used in the context of a symptom event.

{EVENT.CAUSE.DATE}

The date of the cause event in `yyyy.mm.dd` format.
 Used in the context of a symptom event.

{EVENT.CAUSE.DURATION}

The duration of the cause event (time difference between problem and recovery events), with precision down to a second.
 Useful in problem recovery messages.
 Used in the context of a symptom event.

{EVENT.CAUSE.ID}

The numeric ID of the cause event.
 Used in the context of a symptom event.

{EVENT.CAUSE.NAME}

The name of the cause event.
 Used in the context of a symptom event.

{EVENT.CAUSE.NSEVERITY}

The numeric value of the cause event severity.
 Possible values: *0* - Not classified, *1* - Information, *2* - Warning, *3* - Average, *4* - High, *5* - Disaster.
 Used in the context of a symptom event.

{EVENT.CAUSE.OBJECT}

The numeric value of the cause event object.
 Possible values: *0* - Trigger, *1* - Discovered host, *2* - Discovered service, *3* - Autoregistration, *4* - Item, *5* - Low-level discovery rule.
 Used in the context of a symptom event.

{EVENT.CAUSE.OPDATA}

The operational data of the underlying trigger of the cause problem.
 Used in the context of a symptom event.

{EVENT.CAUSE.SEVERITY}

The name of the cause event severity.
 Possible values: *Not classified*, *Information*, *Warning*, *Average*, *High*, *Disaster*.
 Used in the context of a symptom event.

{EVENT.CAUSE.SOURCE}

The numeric value of the cause event source.
 Possible values: *0* - Trigger, *1* - Discovery, *2* - Autoregistration, *3* - Internal.
 Used in the context of a symptom event.

{EVENT.CAUSE.STATUS}

The verbal value of the cause event.
 Used in the context of a symptom event.

{EVENT.CAUSE.TAGS}

A comma-separated list of cause event tags.
Expanded to an empty string if no tags exist.
 Used in the context of a symptom event.

{EVENT.CAUSE.TAGSJSON}

A JSON array containing cause event tag objects.
Expanded to an empty array if no tags exist.
 Used in the context of a symptom event.

{EVENT.CAUSE.TAGS.<tag name>}

The cause event tag value referenced by the tag name.
 A tag name containing non-alphanumeric characters (including non-English multibyte-UTF characters) must be double quoted. Quotes and backslashes inside a quoted tag name must be escaped with a backslash.
 Used in the context of a symptom event.

{EVENT.CAUSE.TIME}

The time of the cause event in `hh:mm:ss` format.
 Used in the context of a symptom event.

{EVENT.CAUSE.TIMESTAMP}

The UNIX timestamp of the cause event.
 Used in the context of a symptom event.

{EVENT.CAUSE.UPDATE.HISTORY}

The log of cause problem updates (acknowledgments, etc).
 Used in the context of a symptom event.

{EVENT.CAUSE.VALUE}

The numeric value of the cause event.
 Possible values: *1* - Problem, *0* - Recovering.
 Used in the context of a symptom event.

{EVENT.SYMPTOMS}

The list of symptom events. Includes the following details: host name, event name, severity, age, service tags and values.
 This macro is used in the context of the cause event and returns information about symptom events.

Recovery event macros

{EVENT.RECOVERY.DATE}

The date of the recovery event in `yyyy.mm.dd` format.

{EVENT.RECOVERY.ID}

The numeric ID of the recovery event.

{EVENT.RECOVERY.NAME}

The name of the recovery event.

{EVENT.RECOVERY.STATUS}

The verbal value of the recovery event.

{EVENT.RECOVERY.TAGS}

A comma-separated list of recovery event tags. Expanded to an empty string if no tags exist.

{EVENT.RECOVERY.TAGSJSON}

A JSON array containing recovery event tag objects. Expanded to an empty array if no tags exist.

{EVENT.RECOVERY.TIME}

The time of the recovery event in `hh:mm:ss` format.

{EVENT.RECOVERY.TIMESTAMP}

The UNIX timestamp of the recovery event.

{EVENT.RECOVERY.VALUE}

The numeric value of the recovery event.

Event update macros

{EVENT.UPDATE.ACTION}

Human-readable name of the action(s) performed during problem update.
 Resolves to the following values: *acknowledged*, *unacknowledged*, *commented*, *changed severity from (original severity) to (updated severity)* and *closed* (depending on how many actions are performed in one update).

{EVENT.UPDATE.ACTIONJSON}

A JSON array containing details of the action(s) performed during problem update.
 Possible JSON property values:
- true (for `acknowledge`, `unacknowledge`, `close`, `unsuppress`, `cause`, and `symptom` properties);
- <message string> (for `message` property);
- timestamp (for `suppress_until` and `timestamp` properties) or 0 (for `suppress_until` property if suppressed indefinitely);
- 0, 1, 2, 3, 4, 5 (for `old` and `new` severity properties).

 For example: `{"acknowledge":true,"message":"Monthly maintenance.","severity":{"old":2,"new":1},"suppress_until":1730851`

{EVENT.UPDATE.DATE}

The date of the event update (acknowledgment, etc) in `yyyy.mm.dd` format.

{EVENT.UPDATE.HISTORY}

The log of problem updates (acknowledgments, etc).

{EVENT.UPDATE.MESSAGE}

The problem update message.

{EVENT.UPDATE.NSEVERITY}

The numeric value of the new event severity set during problem update operation.

{EVENT.UPDATE.SEVERITY}

The name of the new event severity set during problem update operation.

{EVENT.UPDATE.STATUS}

The Numeric value of the problem update status.
 Possible values: *0* - Webhook was called because of problem/recovery event, *1* - Update operation.

{EVENT.UPDATE.TIME}

The time of the event update (acknowledgment, etc) in `hh:mm:ss` format.

{EVENT.UPDATE.TIMESTAMP}

The UNIX timestamp of the event update (acknowledgment, etc).

Function macros

{FUNCTION.VALUE}

The value of the Nth item-based function in the trigger expression at the time of the event.
 Only functions with `/host/key` as the first parameter are counted.

 This macro may be used with a numeric index, e.g., {FUNCTION.VALUE<1-9>} to point to the first, second, third, etc. function in a trigger expression. See indexed macros.

{FUNCTION.RECOVERY.VALUE}

The value of the Nth item-based function in the recovery expression at the time of the event.*
 Only functions with `/host/keys` as the first parameter are counted.

 This macro may be used with a numeric index, e.g., {FUNCTION.RECOVERY.VALUE<1-9>} to point to the first, second, third, etc. function in a trigger expression. See indexed macros.

Host macros

{HOST.CONN}

The host IP address or DNS name, depending on host settings.
 In global scripts, interface IP/DNS fields and web scenarios the macro will resolve to the main agent interface. If an agent interface is not defined, the main SNMP interface will be used. If an SNMP interface is also not defined, the main JMX interface will be used. If a JMX interface is not defined either, the main IPMI interface will be used. If the host does not have any interface, the macro resolves to 'UNKNOWN'.

 This macro may be used with a numeric index as {HOST.CONN<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{HOST.DESCRIPTION}

The host description.

 This macro may be used with a numeric index as {HOST.DESCRIPTION<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{HOST.DNS}

The host DNS name.
 In global scripts, interface IP/DNS fields and web scenarios the macro will resolve to the main agent interface. If an agent interface is not defined, the main SNMP interface will be used. If an SNMP interface is also not defined, the main JMX interface will be used. If a JMX interface is not defined either, the main IPMI interface will be used. If the host does not have any interface, the macro resolves to 'UNKNOWN'.

 This macro may be used with a numeric index as {HOST.DNS<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{HOST.HOST}

The technical host name.
 Macro functions are not supported for this macro if it is used as a placeholder in the first parameter of a history function, for example, `last(/{HOST.HOST}/{ITEM.KEY})`.

 This macro may be used with a numeric index as {HOST.HOST<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{HOST.ID}

The host ID.

 This macro may be used with a numeric index as {HOST.ID<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{HOST.IP}

The host IP address.
 In global scripts, interface IP/DNS fields and web scenarios the macro will resolve to the main agent interface. If an agent interface is not defined, the main SNMP interface will be used. If an SNMP interface is also not defined, the main JMX interface will be used. If a JMX interface is not defined either, the main IPMI interface will be used. If the host does not have any interface, the macro resolves to 'UNKNOWN'.

 This macro may be used with a numeric index as {HOST.IP<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{HOST.METADATA}

The host metadata.

 Used only for active agent autoregistration.

{HOST.NAME}

The visible host name.

 This macro may be used with a numeric index as {HOST.NAME<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{HOST.PORT}

The host (agent) port.
 In global scripts, interface IP/DNS fields and web scenarios the macro will resolve to the main agent interface. If an agent interface is not defined, the main SNMP interface will be used. If an SNMP interface is also not defined, the main JMX interface will be used. If a JMX interface is not defined either, the main IPMI interface will be used. If the host does not have

any interface, the macro resolves to 'UNKNOWN'.

 This macro may be used with a numeric index as {HOST.PORT<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

Target host macros

{HOST.TARGET.CONN}

The IP address or DNS name of the target host.

{HOST.TARGET.DNS}

The DNS name of the target host.

{HOST.TARGET.HOST}

The technical name of the target host.

{HOST.TARGET.IP}

The IP address of the target host.

{HOST.TARGET.NAME}

The visible name of the target host.

Host group macros

{HOSTGROUP.ID}

The host group ID.

Host inventory macros

{INVENTORY.ALIAS}

The *Alias* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.ALIAS<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.ASSET.TAG}

The *Asset tag* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.ASSET.TAG<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.CHASSIS}

The *Chassis* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.CHASSIS<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.CONTACT}

The *Contact* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.CONTACT<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.CONTRACT.NUMBER}

The *Contract number* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.CONTRACT.NUMBER<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.DEPLOYMENT.STATUS}

The *Deployment status* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.DEPLOYMENT.STATUS<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.HARDWARE}

The *Hardware* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.HARDWARE<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.HARDWARE.FULL}

The *Hardware (Full details)* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.HARDWARE.FULL<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.HOST.NETMASK}

The *Host subnet mask* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.HOST.NETMASK<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.HOST.NETWORKS}

The *Host networks* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.HOST.NETWORKS<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.HOST.ROUTER}

The *Host router* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.HOST.ROUTER<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.HW.ARCH}

The *Hardware architecture* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.HW.ARCH<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.HW.DATE.DECOMM}

The *Date hardware decommissioned* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.HW.DATE.DECOMM<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.HW.DATE.EXPIRY}

The *Date hardware maintenance expires* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.HW.DATE.EXPIRY<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.HW.DATE.INSTALL}

The *Date hardware installed* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.HW.DATE.INSTALL<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.HW.DATE.PURCHASE}

The *Date hardware purchased* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.HW.DATE.PURCHASE<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.INSTALLER.NAME}

The *Installer name* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.INSTALLER.NAME<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.LOCATION}

The *Location* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.LOCATION<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.LOCATION.LAT}

The *Location latitude* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.LOCATION.LAT<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.LOCATION.LON}

The *Location longitude* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.LOCATION.LON<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.MACADDRESS.A}

The *MAC address A* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.MACADDRESS.A<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.MACADDRESS.B}

The *MAC address B* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.MACADDRESS.B<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.MODEL}

The *Model* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.MODEL<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.NAME}

The *Name* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.NAME<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.NOTES}

The *Notes* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.NOTES<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.OOB.IP}

The *OOB IP address* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.OOB.IP<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.OOB.NETMASK}

The *OOB subnet mask* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.OOB.NETMASK<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.OOB.ROUTER}

The *OOB router* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.OOB.ROUTER<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.OS}

The *OS* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.OS<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.OS.FULL}

The *OS (Full details)* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.OS.FULL<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.OS.SHORT}

The *OS (Short)* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.OS.SHORT<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.PRIMARY.CELL}

The *Primary POC cell* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.PRIMARY.CELL<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.PRIMARY.EMAIL}

The *Primary POC email* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.PRIMARY.EMAIL<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.PRIMARY.NAME}

The *Primary POC name* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.PRIMARY.NAME<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.PRIMARY.NOTES}

The *Primary POC notes* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.PRIMARY.NOTES<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.PRIMARY.PHONE.A}

The *Primary POC phone A* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.PRIMARY.PHONE.A<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.PRIMARY.PHONE.B}

The *Primary POC phone B* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.PRIMARY.PHONE.B<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.PRIMARY.SCREEN}

The *Primary POC screen name* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.PRIMARY.SCREEN<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.SECONDARY.CELL}

The *Secondary POC cell* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.SECONDARY.CELL<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.SECONDARY.EMAIL}

The *Secondary POC email* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.SECONDARY.EMAIL<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.SECONDARY.NAME}

The *Secondary POC name* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.SECONDARY.NAME<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.SECONDARY.NOTES}

The *Secondary POC notes* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.SECONDARY.NOTES<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.SECONDARY.PHONE.A}

The *Secondary POC phone A* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.SECONDARY.PHONE.A<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.SECONDARY.PHONE.B}

The *Secondary POC phone B* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.SECONDARY.PHONE.B<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.POC.SECONDARY.SCREEN}

The *Secondary POC screen name* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.POC.SECONDARY.SCREEN<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SERIALNO.A}

The *Serial number A* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SERIALNO.A<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SERIALNO.B}

The *Serial number B* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SERIALNO.B<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SITE.ADDRESS.A}

The *Site address A* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SITE.ADDRESS.A<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SITE.ADDRESS.B}

The *Site address B* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SITE.ADDRESS.B<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SITE.ADDRESS.C}

The *Site address C* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SITE.ADDRESS.C<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SITE.CITY}

The *Site city* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SITE.CITY<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SITE.COUNTRY}

The *Site country* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SITE.COUNTRY<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SITE.NOTES}

The *Site notes* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SITE.NOTES<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SITE.RACK}

The *Site rack location* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SITE.RACK<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SITE.STATE}

The *Site state/province* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SITE.STATE<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SITE.ZIP}

The *Site ZIP/postal* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SITE.ZIP<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SOFTWARE}

The *Software* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SOFTWARE<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SOFTWARE.APP.A}

The *Software application A* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SOFTWARE.APP.A<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SOFTWARE.APP.B}

The *Software application B* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SOFTWARE.APP.B<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SOFTWARE.APP.C}

The *Software application C* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SOFTWARE.APP.C<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SOFTWARE.APP.D}

The *Software application D* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SOFTWARE.APP.D<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SOFTWARE.APP.E}

The *Software application E* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SOFTWARE.APP.E<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.SOFTWARE.FULL}

The *Software (Full details)* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.SOFTWARE.FULL<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.TAG}

The *Tag* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.TAG<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.TYPE}

The *Type* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.TYPE<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.TYPE.FULL}

The *Type (Full details)* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.TYPE.FULL<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.URL.A}

The *URL A* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.URL.A<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.URL.B}

The *URL B* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.URL.B<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.URL.C}

The *URL C* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.URL.C<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{INVENTORY.VENDOR}

The *Vendor* field in host inventory.

 This macro may be used with a numeric index, e.g., {INVENTORY.VENDOR<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

Item macros

{ITEM.DESCRIPTION}

The description of the Nth item in the trigger expression that caused a notification.

 This macro may be used with a numeric index as {ITEM.DESCRIPTION<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{ITEM.DESCRIPTION.ORIG}

The original description (with macros unresolved) of the Nth item in the trigger expression that caused a notification.

 This macro may be used with a numeric index as {ITEM.DESCRIPTION.ORIG<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{ITEM.ID}

The numeric ID of the Nth item in the trigger expression that caused a notification.

 This macro may be used with a numeric index as {ITEM.ID<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{ITEM.KEY}

The key of the Nth item in the trigger expression that caused a notification.
 Macro functions are not supported for this macro if it is used as a placeholder in the first parameter of a history function, for example, `last(/{HOST.HOST}/{ITEM.KEY})`.

 This macro may be used with a numeric index as {ITEM.KEY<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{ITEM.KEY.ORIG}

The original key (with macros unresolved) of the Nth item in the trigger expression that caused a notification.

 This macro may be used with a numeric index as {ITEM.KEY.ORIG<1-9>} to point to the first, second, third, etc. host in a trigger expression. See indexed macros.

{ITEM.LASTVALUE}

The latest value of the Nth item in the trigger expression that caused a notification.
 It will resolve to *UNKNOWN* in the frontend if the latest history value has been collected more than the *Max history display period* time ago (set in the *Administration > General* menu section).
 When used in the problem name, the macro will not resolve to the latest item value when viewing problem events; instead, it will keep the item value from the time when the problem happened.
When used in notifications, in some cases the macro might not resolve to the latest item value at the moment the trigger fired. For example, if an item quickly receives two values, ”A” and ”B”, and the trigger fires for ”A”, notifications might show ”B” as the latest value due to a slight processing delay - the latest item value changed between the time the trigger fired and when the notification was created. To avoid this, you can use the {ITEM.VALUE} macro, which resolves to the value at the moment the trigger fires, ensuring the correct value is used in the notification.
 It is alias to `last(/{HOST.HOST}/{ITEM.KEY})`.
 The resolved value for text/log items is truncated to 20 characters by the frontend in the following locations:
- Operational data;
- Trigger description;
- Trigger URLs;
- Trigger URL labels;
- Description of the item value widget.
To resolve to a full value, you may use macro functions, as no values are truncated by the server. For example: `{{ITEM.LASTVALUE}.regsub("(.*)", \1)}`.

 This macro may be used with a numeric index, e.g., {ITEM.LASTVALUE<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.LASTVALUE.AGE}

The time that elapsed between the latest item value collection and macro evaluation.
 Displayed in a human-readable format (e.g., 1m 45s).

 This macro may be used with a numeric index, e.g., {ITEM.LASTVALUE.AGE<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.LASTVALUE.DATE}

The date when the latest item value was collected.
 Displayed in a `YYYYMMDD` format.
 Server time zone is used in trigger names (in *Monitoring* > *Problems* list), event names, and tag names and values. In all other cases user time zone is used.

 This macro may be used with a numeric index, e.g., {ITEM.LASTVALUE.DATE<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.LASTVALUE.TIME}

The time when the latest item value was collected.
 Displayed in an `HHMMSS` format.
 Server time zone is used in trigger names (in *Monitoring* > *Problems* list), event names, and tag names and values. In all other cases user time zone is used.

 This macro may be used with a numeric index, e.g., {ITEM.LASTVALUE.TIME<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.LASTVALUE.TIMESTAMP}

The UNIX timestamp when the latest item value was collected.

 This macro may be used with a numeric index, e.g., {ITEM.LASTVALUE.TIMESTAMP<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.LOG.AGE}

The age of the log event.
With precision down to a second.

 This macro may be used with a numeric index, e.g., {ITEM.LOG.AGE<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.LOG.DATE}

The date when the log entry was written to the log.

 This macro may be used with a numeric index, e.g., {ITEM.LOG.DATE<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.LOG.EVENTID}

The ID of the event in the event log.
 For Windows event log monitoring only.

 This macro may be used with a numeric index, e.g., {ITEM.LOG.EVENTID<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.LOG.NSEVERITY}

The numeric severity of the event in the event log.
 For Windows event log monitoring only.

 This macro may be used with a numeric index, e.g., {ITEM.LOG.NSEVERITY<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.LOG.SEVERITY}

The verbal severity of the event in the event log.
 For Windows event log monitoring only.

 This macro may be used with a numeric index, e.g., {ITEM.LOG.SEVERITY<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.LOG.SOURCE}

The source of the event in the event log.
 For Windows event log monitoring only.

 This macro may be used with a numeric index, e.g., {ITEM.LOG.SOURCE<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.LOG.TIME}

The time when the log entry was written to the log.

 This macro may be used with a numeric index, e.g., {ITEM.LOG.TIME<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.LOG.TIMESTAMP}

The UNIX timestamp when the log entry was written to the log.

 This macro may be used with a numeric index, e.g., {ITEM.LOG.TIMESTAMP<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.NAME}

The name of the item with all macros resolved.

 This macro may be used with a numeric index, e.g., {ITEM.NAME<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.NAME.ORIG}

The original name (with macros unresolved) of the item.

 This macro may be used with a numeric index, e.g., {ITEM.NAME.ORIG<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.STATE}

The latest state of the Nth item in the trigger expression that caused a notification.
 Possible values: *Not supported*, *Normal*.

 This macro may be used with a numeric index, e.g., {ITEM.STATE<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.STATE.ERROR}

The error message with details why an item became unsupported.
 If an item goes into the unsupported state and then immediately gets supported again the error field can be empty.

{ITEM.VALUE}

Resolved to either:

- The historical (at-the-time-of-event) value of the Nth item in the trigger expression, if used in the context of change of trigger state, for example, when displaying events or sending notifications;
- The latest value of the Nth item in the trigger expression, if used without the context of trigger state change, for example, when displaying a list of triggers in a pop-up selection window. In this case works the same as {ITEM.LASTVALUE}.

In the first case it will resolve to *UNKNOWN* if the history value has already been deleted or has never been stored.
In the second case, and in the frontend only, it will resolve to *UNKNOWN* if the latest history value has been collected more than the *Max history display period* time ago (set in the *Administration > General* menu section).
The resolved value for text/log items is truncated to 20 characters by the frontend in the following locations:
- Operational data;
- Trigger description;
- Trigger URLs;
- Trigger URL labels;
- Description of the item value widget.
To resolve to a full value, you may use macro functions, as no values are truncated by the server. For example: `{{ITEM.VALUE}.regsub("(.*)", \1)}`.

This macro may be used with a numeric index, e.g., {ITEM.VALUE<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.VALUE.AGE}

The time that elapsed between the item value collection and macro evaluation.
 Displayed in a human-readable format (e.g., 1m 45s).

 This macro may be used with a numeric index, e.g., {ITEM.VALUE.AGE<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.VALUE.DATE}

The date when the item value was collected.
 Displayed in a `YYYYMMDD` format.
 The server time zone is used in trigger names (in *Monitoring* > *Problems* list), event names, and tag names and values. In all other cases user time zone is used.

This macro may be used with a numeric index, e.g., {ITEM.VALUE.DATE<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.VALUE.TIME}

The time when the item value was collected.
 Displayed in an `HHMMSS` format.
 The server time zone is used in trigger names (in *Monitoring* > *Problems* list), event names, and tag names and values. In all other cases user time zone is used.

This macro may be used with a numeric index, e.g., {ITEM.VALUE.TIME<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.VALUE.TIMESTAMP}

The UNIX timestamp when the item value was collected.

 This macro may be used with a numeric index, e.g., {ITEM.VALUE.TIMESTAMP<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{ITEM.VALUETYPE}

The value type of the Nth item in the trigger expression that caused a notification.
 Possible values: *0* - Numeric float, *1* - Character, *2* - Log, *3* - Numeric unsigned, *4* - Text.

 This macro may be used with a numeric index, e.g., {ITEM.VALUETYPE<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

Low-level discovery notification macros

{LLDRULE.DESCRIPTION}

The description of the low-level discovery rule that caused a notification.

{LLDRULE.DESCRIPTION.ORIG}

The original description (with macros unresolved) of the low-level discovery rule that caused a notification.

{LLDRULE.ID}

The numeric ID of the low-level discovery rule that caused a notification.

{LLDRULE.KEY}

The key of the low-level discovery rule that caused a notification.

{LLDRULE.KEY.ORIG}

The original key (with macros unresolved) of the low-level discovery rule that caused a notification.

{LLDRULE.NAME}

The name of the low-level discovery rule that caused a notification.

{LLDRULE.NAME.ORIG}

The original name (with macros unresolved) of the low-level discovery rule that caused a notification.

{LLDRULE.STATE}

The latest state of the low-level discovery rule.
 Possible values: *Not supported*, *Normal*.

 This macro may be used with a numeric index, e.g., {LLDRULE.STATE<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{LLDRULE.STATE.ERROR}

The error message with details why the low-level discovery rule became unsupported.
 If a low-level discovery rule goes into the unsupported state and then immediately gets supported again the error field can be empty.

Map macros

{MAP.ID}

The network map ID.

{MAP.NAME}

The network map name.

Proxy macros

{PROXY.DESCRIPTION}

The proxy description.
 Resolves to either:

- Proxy of the Nth item in the trigger expression (in trigger-based notifications). You may use indexed macros here;
- Proxy, which executed discovery (in discovery notifications). Use {PROXY.DESCRIPTION} here, without indexing;
- Proxy to which an active agent registered (in autoregistration notifications). Use {PROXY.DESCRIPTION} here, without indexing.

This macro may be used with a numeric index, e.g., {PROXY.DESCRIPTION<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

{PROXY.NAME}

The proxy name.
 Resolves to either:

- Proxy of the Nth item in the trigger expression (in trigger-based notifications). You may use indexed macros here;
- Proxy, which executed discovery (in discovery notifications). Use {PROXY.NAME} here, without indexing;
- Proxy to which an active agent registered (in autoregistration notifications). Use {PROXY.NAME} here, without indexing.

This macro may be used with a numeric index, e.g., {PROXY.NAME<1-9>} to point to the first, second, third, etc. item in a trigger expression. See indexed macros.

Script macros

{MANUALINPUT}

The manual input value specified by user at script execution time.

Service macros

{SERVICE.DESCRIPTION}

The service description with macros resolved.

{SERVICE.ID}

The numeric ID of the service that triggered an action.

{SERVICE.NAME}

The service name with macros resolved.

{SERVICE.ROOTCAUSE}

The list of trigger problem events that caused a service to fail, sorted by severity and host name.
 Includes the following details: host name, event name, severity, age, service tags and values.

{SERVICE.TAGS}

A comma-separated list of service event tags.
 Service event tags can be defined in the service configuration tags section. Expanded to an empty string if no tags exist.

{SERVICE.TAGSJSON}

A JSON array containing service event tag objects.
 Service event tags can be defined in the service configuration tags section. Expanded to an empty array if no tags exist.

{SERVICE.TAGS.<tag name>}

The service event tag value referenced by the tag name.
 The service event tags can be defined in the service configuration tags section.
 A tag name containing non-alphanumeric characters (including non-English multibyte-UTF characters) must be double quoted. Quotes and backslashes inside a quoted tag name must be escaped with a backslash.

Trigger macros

{TRIGGER.DESCRIPTION}

The trigger description.
 All macros supported in a trigger description will be expanded if `{TRIGGER.DESCRIPTION}` is used in notification text.

{TRIGGER.EVENTS.ACK}

The number of acknowledged events for a map element in maps, or for the trigger which generated current event in notifications.

{TRIGGER.EVENTS.PROBLEM.ACK}

The number of acknowledged problem events for all triggers disregarding their state.

{TRIGGER.EVENTS.PROBLEM.UNACK}

The number of unacknowledged problem events for all triggers disregarding their state.

{TRIGGER.EVENTS.UNACK}

The number of unacknowledged events for a map element in maps, or for the trigger which generated current event in notifications.

{TRIGGER.EXPRESSION}

The trigger expression.

{TRIGGER.EXPRESSION.EXPLAIN}

A partially-evaluated trigger expression.
 Item-based functions are evaluated and replaced by the results at the time of event generation whereas all other functions are displayed as written in the expression. Can be used for debugging trigger expressions.

{TRIGGER.EXPRESSION.RECOVERY}

The trigger recovery expression if *OK event generation* in trigger configuration is set to "Recovery expression"; otherwise, an empty string is returned.

{TRIGGER.EXPRESSION.RECOVERY.EXPLAIN}

A partially-evaluated trigger recovery expression.
 Item-based functions are evaluated and replaced by the results at the time of event generation whereas all other functions are displayed as written in the expression. Can be used for debugging trigger recovery expressions.

{TRIGGER.HOSTGROUP.NAME}

A sorted (by SQL query), comma-space separated list of host groups in which the trigger is defined.

{TRIGGER.ID}

The numeric trigger ID which triggered this action.

{TRIGGER.NAME}

The trigger name with macros resolved.
 Note that `{EVENT.NAME}` can be used in actions to display the triggered event/problem name with macros resolved.

{TRIGGER.NAME.ORIG}

The original trigger name (with macros unresolved).

{TRIGGER.NSEVERITY}

The numerical trigger severity.
 Possible values: *0* - Not classified, *1* - Information, *2* - Warning, *3* - Average, *4* - High, *5* - Disaster.

{TRIGGER.PROBLEM.EVENTS.PROBLEM.ACK}

The number of acknowledged problem events for triggers in problem state.

{TRIGGER.PROBLEM.EVENTS.PROBLEM.UNACK}

The number of unacknowledged problem events for triggers in problem state.

{TRIGGER.SEVERITY}

The trigger severity name.
 Can be defined in *Administration > General > Trigger displaying options*.

{TRIGGER.STATE}

The latest state of the trigger expression.
 Possible values: *Unknown*, *Normal*.

{TRIGGER.STATE.ERROR}

The error message with details why a trigger became unsupported.
 If a trigger goes into the unsupported state and then immediately gets supported again the error field can be empty.

{TRIGGER.STATUS}

The trigger value at the time of operation step execution.
 Possible values: *Problem*, *OK*.

{TRIGGER.TEMPLATE.NAME}

A sorted (by SQL query), comma-space separated list of templates in which the trigger is defined, or *UNKNOWN* if the trigger is defined in a host.

{TRIGGER.URL}

The trigger URL.

{TRIGGER.URL.NAME}

The label for the trigger URL.

{TRIGGER.VALUE}

The current numeric trigger value.
 Possible values: *0* - trigger is in OK state, *1* - trigger is in Problem state.

{TRIGGERS.ACK}

The number of acknowledged triggers for a map element, disregarding trigger state.
 A trigger is considered to be acknowledged if all of its problem events are acknowledged.

{TRIGGERS.PROBLEM.ACK}

The number of acknowledged problem triggers for a map element.
 A trigger is considered to be acknowledged if all of its problem events are acknowledged.

{TRIGGERS.PROBLEM.UNACK}

The number of unacknowledged problem triggers for a map element.
 A trigger is considered to be unacknowledged if at least one of its problem events is unacknowledged.

{TRIGGERS.UNACK}

The number of unacknowledged triggers for a map element, disregarding trigger state.
 A trigger is considered to be unacknowledged if at least one of its problem events is unacknowledged.

Username macros

{USER.FULLNAME}

The name, surname and username of the user who added event acknowledgment or started the script.

{USER.NAME}

The name of the user who started the script.

{USER.SURNAME}

The surname of the user who started the script.

{USER.USERNAME}

The username of the user who started the script.

Indexed macros

The indexed macro syntax **{MACRO<1-9>}** can only reference the N-th item, function, or host in a trigger's *Expression* field:

- **{HOST.IP1}, {HOST.IP2}, {HOST.IP3}** resolve to the IP addresses of the first, second, and third hosts in the trigger expression (if present).
- **{ITEM.VALUE1}, {ITEM.VALUE2}, {ITEM.VALUE3}** resolve to the values of the first, second, and third items in the trigger expression at the time of the event (if present).
- **{FUNCTION.VALUE1}, {FUNCTION.VALUE2}, {FUNCTION.VALUE3}** resolve to the values of the first, second, and third item-based functions at the time of the event (if present).

In the context of triggers, indexed macros always refer to the *Expression* field of the trigger configuration, not the *Recovery expression*. For example, in a recovery event, **{ITEM.VALUE2}** will resolve to the value of the second item from the problem expression at the time of recovery.

The **{HOST.HOST<1-9>}** macro is also supported within the `{?func(/host/key,param)}` expression macro in **graph names**. For example, `{?func(/{HOST.HOST2}/key,param)}` in a graph name will resolve to the host of the second item in the graph.

Warning:
Indexed macros will not resolve in any other context beyond the cases mentioned here. For other contexts, use macros without an index (**{HOST.HOST}, {HOST.IP}**, etc.).

User macros supported by location

Overview

This section contains a list of locations, where user-definable macros are supported.

Note:
Only global-level user macros are supported for *Actions*, *Network discovery*, *Proxies* and all locations listed under *Other locations* section of this page. In the mentioned locations, host-level and template-level macros will not be resolved.

Note:
To customize macro values (for example, shorten or extract specific substrings), you can use macro functions.

Actions

In actions, user macros can be used in the following fields:

Location		Multiple macros/mix with text[1]
Trigger-based notifications and commands		yes
Trigger-based internal notifications		yes
Problem update notifications		yes
Service-based notifications and commands		yes
Service update notifications		yes
Time period condition		no
Operations		
	Default operation step duration	no
	Step duration	no

Hosts/host prototypes

In a host and host prototype configuration, user macros can be used in the following fields:

Location		Multiple macros/mix with text[1]
Interface IP/DNS		DNS only
Interface port		no
SNMP		
	SNMP community	yes
	Max repetition count	no
	Retries	no
	Context name	yes
	Security name	yes
	Authentication passphrase	yes
	Privacy passphrase	yes
IPMI		
	Username	yes
	Password	yes
Tags[2]		
	Tag names	yes
	Tag values	yes

Items / item prototypes

In an item or an item prototype configuration, user macros can be used in the following fields:

Location	Multiple macros/mix with text[1]
Item name	yes

Location	Multiple macros/mix with text[1]
Item key parameters	yes[4]
Update interval	no
Custom intervals	no
Timeout (available for supported item types)	no
Store up to (for history and trends)	no
Description	yes
Calculated/aggregate item	
Formula (expression constants and function parameters; item key parameters; (*aggregate item only*) filter conditions (host group name and tag name))	yes
Database monitor	
Username	yes
Password	yes
SQL query	yes
HTTP agent	
URL[3]	yes
Query fields	yes
Request body	yes
Headers (names and values)	yes
Required status codes	yes
HTTP proxy	yes
HTTP authentication username	yes
HTTP authentication password	yes
SSl certificate file	yes
SSl key file	yes
SSl key password	yes
Allowed hosts	yes
JMX agent	
JMX endpoint	yes
Script item	

Location	Multiple macros/mix with text[1]
Parameter names and values	yes
Browser item	
Parameter names and values	yes
Telemetry query item	
Time shift	no
Lookback limit	no
Granularity	no
SNMP agent	
SNMP OID	yes
SSH agent	
Username	yes
Public key file	yes
Private key file	yes
Password	yes
Script	yes
TELNET agent	
Username	yes
Password	yes
Script	yes
Zabbix trapper	
Allowed hosts	yes
Tags[2]	
Tag names	yes
Tag values	yes
Preprocessing steps	
Parameters (including custom scripts)	yes
Custom error-handling parameters (*Set value to* and *Set error to* fields)	yes

Low-level discovery

In a low-level discovery rule, user macros can be used in the following fields:

Location		Multiple macros/mix with text[1]
Key parameters		yes
Update interval		no
Custom interval		no
Timeout (available for supported item types)		no
Delete lost resources		no
Disable lost resources		no
Description		yes
SNMP agent		
	SNMP OID	yes
SSH agent		
	Username	yes
	Public key file	yes
	Private key file	yes
	Password	yes
	Script	yes
TELNET agent		
	Username	yes
	Password	yes

Location		Multiple macros/mix with text[1]
	Script	yes
Zabbix trapper		
	Allowed hosts	yes
Database monitor		
	Username	yes
	Password	yes
	SQL query	yes
JMX agent		
	JMX endpoint	yes
HTTP agent		
	URL[3]	yes
	Query fields	yes
	Request body	yes
	Headers (names and values)	yes
	Required status codes	yes
	HTTP authentication username	yes
	HTTP authentication password	yes
Filters		
	Regular expression	yes
Overrides		
	Filters: regular expression	yes
	Operations: update interval (for item prototypes)	no
	Operations: history storage period (for item prototypes)	no
	Operations: trend storage period (for item prototypes)	no

Network discovery

In a network discovery rule, user macros can be used in the following fields:

Location		Multiple macros/mix with text[1]
Update interval		no
SNMP v1, v2		
	SNMP community	yes
	SNMP OID	yes
SNMP v3		
	Context name	yes
	Security name	yes
	Authentication passphrase	yes
	Privacy passphrase	yes
	SNMP OID	yes

Proxies

In a proxy configuration, user macros can be used in the following fields:

Location	Multiple macros/mix with text[1]
Address for active agents > *Port* (if the proxy belongs to a group)	no
Interface address and port (for passive proxy)	no
Timeouts for item types	no

Proxy groups

In a proxy group configuration, user macros can be used in the following fields:

Location	Multiple macros/mix with text[1]
Failover period	no
Minimum number of proxies	no

Templates

In a template configuration, user macros can be used in the following fields:

Location		Multiple macros/mix with text[1]
Tags[2]		
	Tag names	yes
	Tag values	yes

Triggers

In a trigger configuration, user macros can be used in the following fields:

Location	Multiple macros/mix with text[1]
Name	yes
Operational data	yes
Expression (only in constants and function parameters; secret macros are not supported)	yes
Tag for matching	yes
Menu entry name	yes
Menu entry URL[3]	yes
Description	yes
Tags[2]	
Tag names	yes
Tag values	yes

Web scenario

In a web scenario configuration, user macros can be used in the following fields:

Location	Multiple macros/mix with text[1]
Name	yes
Update interval	no

Location		Multiple macros/mix with text[1]
Agent		yes
HTTP proxy		yes
Variables (values only)		yes
Headers (names and values)		yes
Steps		
	Name	yes
	URL[3]	yes
	Variables (values only)	yes
	Headers (names and values)	yes
	Timeout	no
	Required string	yes
	Required status codes	no
Authentication		
	User	yes
	Password	yes
	SSL certificate	yes
	SSL key file	yes
	SSL key password	yes
Tags[2]		
	Tag names	yes
	Tag values	yes

Other locations

In addition to the locations listed here, user macros can be used in the following fields:

Location	Multiple macros/mix with text[1]
Global scripts (URL, script, SSH, Telnet, IPMI), including confirmation text	yes
Webhooks	
JavaScript script	no
JavaScript script parameter name	no
JavaScript script parameter value	yes
Dashboards	
Column of data type *Text* in *Top hosts* dashboard widget	yes
Description parameter in *Item value* and *Gauge* dashboard widget	yes
Primary/Secondary label *Text* parameter in *Honeycomb* dashboard widget	yes
URL[3] parameter in *URL* dashboard widget	yes
Monitoring > *Problems*, *Hosts*, *Latest data*	

Location	Multiple macros/mix with text[1]
Custom text column	yes
Data collection > *Templates*, *Hosts*	
Custom text column	yes
Data collection > *Event processing*	
Duration field in CEP rules	yes
Capacity field in CEP rules	yes
Alerts > *Media types* > *Message templates*	
Subject	yes
Message	yes
Alerts > *Media types* > *Script*	
Script parameters	yes
Alerts > *Media types* > *Media type*	
Username and *Password* fields for the *Email* media type (when *Authentication* is set to "Username and password"; secret macros recommended)	yes
Users > *Users* > *Media*	
When active	yes

Location	Multiple macros/mix with text[1]
Administration > General > GUI	
Working time	no
Administration > General > Timeouts	
Timeouts for item types	no
Administration > General > Connectors	
URL	yes
Username	yes
Password	yes
Bearer token	yes
Timeout	no
HTTP proxy	yes
SSL certificate file	yes
SSL key file	yes
SSL key password	yes

For a complete list of all macros supported in Zabbix, see supported macros.

Footnotes

[1] If multiple macros in a field or macros mixed with text are not supported for the location, a single macro has to fill the whole field.

[2] Macros used in tag names and values are resolved only during event generation process.

[3] URLs that contain a secret macro will not work, as the macro in them will be resolved as "******".

[4] Macros are not supported in the item key of Zabbix trapper items.

Version compatibility

Supported agents

To be compatible with Zabbix 8.2, Zabbix agent must not be older than version 1.6 and must not be newer than 8.2.

You may need to review the configuration of older agents as some parameters have changed, for example, parameters related to logging for versions before 3.0.

To take full advantage of the latest functionality, metrics, improved performance and reduced memory usage, use the latest supported agent.

Notes for Windows XP

- On 32-bit Windows XP, do not use Zabbix agents newer than 6.0.x;
- On Windows XP/Server 2003, do not use agent templates that are newer than Zabbix 4.0.x. The newer templates use English performance counters, which are only supported since Windows Vista/Server 2008.

Supported agents 2

Older Zabbix agents 2 from version 4.4 onwards are compatible with Zabbix 8.2; Zabbix agent 2 must not be newer than 8.2.

Note that when using Zabbix agent 2 versions 4.4 and 5.0, the default interval of 10 minutes is used for refreshing unsupported items.

To take full advantage of the latest functionality, metrics, improved performance and reduced memory usage, use the latest supported agent 2.

Supported Zabbix proxies

For full compatibility with Zabbix 8.2, proxies must match the server's major version. Only Zabbix 8.2.x proxies are fully compatible with a Zabbix 8.2.x server.

Outdated proxies are partially supported: they can still collect data and execute scripts but cannot receive configuration updates, such as new items.

In relation to Zabbix server, proxies can be:

- *Current* (proxy and server have the same major version);
- *Outdated* (proxy version is older than server version, but is partially supported);
- *Unsupported* (proxy version is older than server previous LTS release version *or* proxy version is newer than server major version).

Examples:

Server version	*Current* proxy version	*Outdated* proxy version	*Unsupported* proxy version
6.4	6.4	6.0, 6.2	Older than 6.0; newer than 6.4
7.0	7.0	6.0, 6.2, 6.4	Older than 6.0; newer than 7.0
7.2	7.2	7.0	Older than 7.0; newer than 7.2
7.4	7.4	7.0	Older than 7.0; newer than 7.4
8.0	8.0	7.0, 7.2, 7.4	Older than 7.0; newer than 8.0
8.2	8.2	8.0	Older than 8.0; newer than 8.2

Functionality supported by proxies:

Proxy version	Data update	Configuration update	Tasks
Current	Yes	Yes	Yes
Outdated	Yes	No	Remote commands (e.g., shell scripts); Immediate item value checks (i.e., *Execute now*); Note: Preprocessing tests with a real value are not supported.
Unsupported	No	No	No

Warnings about using incompatible Zabbix daemon versions are logged.

Supported XML files

XML files not older than version 1.8 are supported for import in Zabbix 8.2.

Attention:
In the XML export format, trigger dependencies are stored by name only. If there are several triggers with the same name (for example, having different severities and expressions) that have a dependency defined between them, it is not possible to import them. Such dependencies must be manually removed from the XML file and re-added after import.

Time period syntax

Overview

To set a time period, the following format has to be used:

`d-d,hh:mm-hh:mm`

where the symbols stand for the following:

Symbol	Description
d	Day of the week: 1 - Monday, 2 - Tuesday ,... , 7 - Sunday
hh	Hours: 00-24
mm	Minutes: 00-59

You can specify more than one time period using a semicolon (;) separator:

`d-d,hh:mm-hh:mm;d-d,hh:mm-hh:mm...`

Leaving the time period empty equals 1-7,00:00-24:00, which is the default value.

Attention:
The upper limit of a time period is not included. Thus, if you specify 09:00-18:00 the last second included in the time period is 17:59:59.

Examples

Working hours. Monday - Friday from 9:00 till 18:00:

`1-5,09:00-18:00`

Working hours plus weekend. Monday - Friday from 9:00 till 18:00 and Saturday, Sunday from 10:00 till 16:00:

`1-5,09:00-18:00;6-7,10:00-16:00`

Unit symbols

Overview

Working with large values such as 86400, 104857600, or 1000000 can be challenging and can lead to errors. To simplify configuration and improve readability, Zabbix supports unit symbols (suffixes) that function as value multipliers.

Trigger expressions without suffixes:

```
last(/host/system.uptime)<86400
avg(/host/system.cpu.load,600s)<10
last(/host/vm.memory.size[available])<20971520
```

Trigger expressions with suffixes:

```
last(/host/system.uptime)<1d
avg(/host/system.cpu.load,10m)<10
last(/host/vm.memory.size[available])<20M
```

Suffixes can also simplify the configuration of other entities, such as items, widgets, etc., and help display item values in a human-readable format.

Note:
To see if an entity configuration field supports suffixes, always see the relevant page for the entity being configured.

Time suffixes

Zabbix frontend supports the following time suffixes in entity configuration:
 s - seconds *(when used, works the same as the raw value)*;
 m - minutes;
 h - hours;
 d - days;
 w - weeks;
 M - months *(trend functions only)*;
 y - years *(trend functions only)*.

Time suffixes are supported only with integer numbers. For example, 1h is supported, but 1,5h or 1.5h is not supported; use 90m instead.

Memory size suffixes

Zabbix supports the following memory size suffixes:
 K - kilobytes
 M - megabytes
 G - gigabytes
 T - terabytes

Item value suffixes

Suffixes can also be used to display numeric item values in a human-readable format.

To enable this, use one of the following suffixes in the *Units* field when configuring an item:
 B - bytes;
 Bps - bytes per second;
 s - seconds, displayed using up to three largest non-zero time units;
 uptime - elapsed time in hh:mm:ss or N days, hh:mm:ss;
 unixtime - Unix timestamp, formatted as yyyy.mm.dd hh:mm:ss.

The following additional rules apply to how these suffixes are interpreted and displayed:

- For B and Bps, Zabbix uses base-2 conversion (1K = 1024B), following the JEDEC standard.
- For other units (such as Hz, W, etc.), Zabbix uses base-10 conversion (1K = 1000).
- For s (seconds):
  - The format includes yyy mmm ddd hhh mmm sss ms; only up to three largest non-zero time units are displayed (e.g., 1M 10d 4h).
  - If a unit is zero and between two non-zero units, it is omitted (e.g., 10d 56m instead of 10d 0h 56m).

When *Units* are used, the following multiplier suffixes are automatically applied to item values:
 K, **M**, **G**, **T** - kilo, mega, giga, tera;
 P, **E**, **Z**, **Y** - peta, exa, zetta, yotta *(these are applied in frontend only)*.

To prevent unit conversion, use the `!` prefix (e.g., !B or !s).

The following examples show how received item values are converted based on the specified units:

```
1 B → 1 B
1024 B → 1 KB
1536 B → 1.5 KB
881764 B → 881.76 KB
881764 !B → 881764 B

0.0000155 s → 0.016ms
3470400 s → 1M 10d 4h
2606400 s → 1M 4h
2592000 s → 1M
2592001 s → 1M
2592001 !s → 2592001 s

17764 uptime → 04:56:04
86400 uptime → 1 day, 00:00:00
881764 uptime → 10 days, 04:56:04
32417764 uptime → 375 days, 04:56:04
32417764 !uptime → 32417764 uptime

881764 unixtime → 1970-01-11 04:56:04 AM

17764 Hz → 17.76 KHz
86400 Hz → 86.4 KHz
881764 Hz → 881.76 KHz
32417764 Hz → 32.42 MHz

0 ! → 0
0 !! → 0 !
```

Note:
Before Zabbix 4.0, there was a hardcoded unit stoplist consisting of `ms`, `rpm`, `RPM`, `%`. This stoplist has been deprecated, thus the correct way to prevent converting such units is `!ms`, `!rpm`, `!RPM`, `!%`.

Command execution

Zabbix uses common functionality for external checks, user parameters, `system.run` items, custom alert scripts, remote commands and global scripts.

Note:
All commands and scripts are executed with the permissions of the OS user running the Zabbix agent. The user must have sufficient permissions to execute the command or script.

Execution steps

Note:
By default, all scripts in Zabbix are executed using the *sh* shell, and it is not possible to modify the default shell. To utilize a different shell, you can employ a workaround: create a script file and invoke that script during command execution.

The command or script is executed similarly on both Unix and Windows platforms:

1. Zabbix (the parent process) creates a pipe for communication.
2. Zabbix sets the pipe as the output for the to-be-created child process.
3. Zabbix creates the child process (runs the command/script).
4. A new process group (in Unix) or a job (in Windows) is created for the child process.
5. Zabbix reads from the pipe until timeout occurs or no one is writing to the other end (ALL handles/file descriptors have been closed). Note that the child process can create more processes and exit before they exit or close the handle/file descriptor.
6. If the timeout has not been reached, Zabbix waits until the initial child process exits or timeout occurs.
7. If the initial child process exits before the timeout is reached, Zabbix checks exit code of the initial child process and compares it with 0. A non-zero exit code is considered an execution failure, but only for custom alert scripts, remote commands, and user scripts executed on the Zabbix server or Zabbix proxy.
8. At this point it is assumed that everything is done and the whole process tree (i.e. the process group or the job) is terminated.

Attention:
Zabbix assumes that a command/script has done processing when the initial child process has exited AND no other process is still keeping the output handle/file descriptor open. When processing is done, ALL created processes are terminated.

All double quotes and backslashes in the command are escaped with backslashes and the command is enclosed in double quotes.

Exit code checking

Exit code are checked with the following conditions:

- Only for custom alert scripts, remote commands and user scripts executed on Zabbix server and Zabbix proxy.
- Any exit code that is different from 0 is considered as execution failure.
- Contents of standard error and standard output for failed executions are collected and available in frontend (where execution result is displayed).
- Additional log entry may be created for remote commands executed on Zabbix agent/proxy by enabling the LogRemoteCommands parameter on agent/proxy.

Possible frontend messages and log entries for failed commands or scripts:

- Contents of standard error and standard output for failed executions (if any).
- `Process exited with code: N.` (for empty output, and exit code not equal to 0).
- `Process killed by signal: N.` (for process terminated by a signal, on Linux only).
- `Process terminated unexpectedly.` (for process terminated for unknown reasons).

See also

- External checks
- User parameters
- system.run items
- Custom alert scripts
- Remote commands
- Global scripts

Escaping examples

Overview

This page provides examples of using correct escaping when using regular expressions in various contexts.

Note:
When using the trigger expression constructor, correct escaping in regular expressions is added automatically.

Examples

User macro with context

Regular expression: `\.+\"[a-z]+<br>` User macro with context: `{$MACRO:regex:"\.+\\"[a-z]+"}`

Notice:

- backslashes are not escaped;
- quotation marks are escaped.

Macro function inside item key parameter

Regular expression: `.+:(\d+)$<br>` Item key: `net.tcp.service[tcp,,"{{$ENDPOINT}.regsub(\".+:(\d+)$\",\1)}"]`

Notice:

- regular expression inside the `regsub` macro function is double-quoted (because of contains closing parenthesis);
- quotation marks around the regular expression are escaped (because the whole third item parameter is double-quoted);
- third item key parameter is double-quoted because it contains a comma.

LLD macro function

Regular expression: `\.+\"([a-z]+)<br>` LLD macro: `{{#MACRO}.iregsub("\.+\\"([a-z]+)", \1)}`

Notice:

- backslashes are not escaped;
- quotation marks are escaped.

LLD macro function inside user macro context

Regular expression: `\.+\"([a-z]+)<br>` LLD macro: `{{#MACRO}.iregsub("\.+\\"([a-z]+)", \1)}<br>` User macro with context: `{$MACRO:"{{#MACRO}.iregsub(\"\.+\\\"[a-z]+\", \1)}"}`

Notice:

- backslash escaping for LLD does not change;
- upon inserting the LLD macro into user macro context, we need to put it into string:

1. Quotation marks are added around the macro expression;
2. Quotation marks get escaped; in total, 3 new backslashes are introduced.

String parameter of function (any)

`concat` is used as example.

String content: `\.+\"[a-z]+<br>` Expression: `concat("abc", "\\.\\\"[a-z]+")`

Notice:

- String parameters require escaping both for backslashes and quotation marks.

LLD macro function inside string parameter of function

Regular expression: `\.+\"([a-z]+)<br>` LLD macro: `{{#MACRO}.iregsub("\.+\\"([a-z]+)", \1)}<br>` Expression: `concat("abc, "{{#MACRO}.iregsub(\"\\.+\\\\\"([a-z]+)\", \\1)}")`

Notice:

- String parameters require escaping both for backslashes and quotation marks;
- Another layer of escaping is added, because the macro will be resolved only after string is unquoted;

User macro with context inside string parameter of function

Regular expression: `\.+\"[a-z]+<br>` User macro with context: `{$MACRO:regex:"\.+\\"[a-z]+"}<br>` Expression: `concat("abc, "{$MACRO:regex:\"\\.+\\\\\"[a-z]+\"}")`

Notice:

- Same as in the previous example an additional layer of escaping is needed;
- Backslashes and quotation marks are escaped only for the top-level escaping (by virtue of it being a string parameter).

LLD macro function inside user macro context inside function

Regular expression: `\.+\"([a-z])+<br>` LLD macro: `{{#MACRO}.iregsub("\.+\\"([a-z]+)", \1)}<br>` User macro with context: `{$MACRO:"{{#MACRO}.iregsub(\"\.+\\\"([a-z]+)\", \1)}"}<br>` Expression: `concat("abc, "{$MACRO:\"{{#MACRO}.iregsub(\\\"\.+\\\\\\\"([a-z]+)\\\", \\1)}\"}")`

Notice the three layers of escaping:

1. For LLD macro function, without escaping of backslashes;
2. For User macro with context, without escaping of backslashes;
3. For the string parameter of a function, with escaping of backslashes.

User macro with context just inside string

Regular expression: `\.+\"[a-z]+<br>` User macro with context: `{$MACRO:regex:"\.+\\"[a-z]+"}<br>` Inside string of some expression, for example: `func(arg1, arg2, arg3)="{$MACRO:regex:\"\\.+\\\\\\"[a-z]+\"}"`

Notice:

- Strings also require backslash escaping;
- Strings also require quotation mark escaping;
- Again a case with 2 levels of escaping:

1. Escaping for user macro context without backslash escaping;
2. Escaping for it being a string with backslash escaping.

What is Zabbix

Zabbix was created by Alexei Vladishev, and currently is actively developed and supported by Zabbix SIA.

Zabbix is an enterprise-class open source distributed monitoring solution.

Zabbix is a software that monitors numerous parameters of a network and the health and integrity of servers, virtual machines, applications, services, databases, websites, the cloud and more. Zabbix uses a flexible notification mechanism that allows users to configure email-based alerts for virtually any event. This allows a fast reaction to server problems. Zabbix offers excellent reporting and data visualization features based on the stored data. This makes Zabbix ideal for capacity planning.

Zabbix supports both polling and trapping. All Zabbix reports and statistics, as well as configuration parameters, are accessed through a web-based frontend. A web-based frontend ensures that the status of your network and the health of your servers can be assessed from any location. Properly configured, Zabbix can play an important role in monitoring IT infrastructure. This is equally true for small organizations with a few servers and for large companies with a multitude of servers.

Zabbix is available at no cost. Zabbix is written and distributed under the AGPL-3.0 license. It means that its source code is freely distributed and available for the general public.

Commercial support is available and provided by Zabbix Company and its partners around the world.

Learn more about Zabbix features.

Users of Zabbix Many organizations of different size around the world rely on Zabbix as a primary monitoring platform.

Architecture Zabbix consists of several major software components. Their responsibilities are outlined below.

Server

Zabbix server is the central component to which agents report availability and integrity information and statistics. The server is the central repository in which all configuration, statistical and operational data is stored.

Database storage

All configuration information as well as the data gathered by Zabbix is stored in a database.

Web interface

For an easy access to Zabbix from anywhere and from any platform, the web-based interface is provided. The interface is part of Zabbix server, and usually (but not necessarily) runs on the same physical machine as the one running the server.

Proxy

Zabbix proxy can collect performance and availability data on behalf of Zabbix server. A proxy is an optional part of Zabbix deployment; however, it may be very beneficial to distribute the load of a single Zabbix server.

Agent

Zabbix agents are deployed on monitoring targets to actively monitor local resources and applications and report the gathered data to Zabbix server. Since Zabbix 4.4, there are two types of agents available: the Zabbix agent (lightweight, supported on many platforms, written in C) and the Zabbix agent 2 (extra-flexible, easily extendable with plugins, written in Go).

Data flow In addition it is important to take a step back and have a look at the overall data flow within Zabbix. In order to create an item that gathers data you must first create a host. Moving to the other end of the Zabbix spectrum you must first have an item to create a trigger. You must have a trigger to create an action. Thus if you want to receive an alert that your CPU load is too high on *Server X* you must first create a host entry for *Server X* followed by an item for monitoring its CPU, then a trigger which activates if the CPU is too high, followed by an action which sends you an email. While that may seem like a lot of steps, with the use of templating it really isn't. However, due to this design it is possible to create a very flexible setup.

Features Zabbix is a highly integrated network monitoring solution, offering a multiplicity of features in a single package.

Data gathering

- availability and performance checks
- support for SNMP (both trapping and polling), IPMI, JMX, VMware monitoring
- custom checks
- gathering desired data at custom intervals
- performed by server/proxy and by agents

Flexible threshold definitions

- you can define very flexible problem thresholds, called triggers, referencing values from the backend database

Highly configurable alerting

- sending notifications can be customized for the escalation schedule, recipient, media type
- notifications can be made meaningful and helpful using macro variables
- automatic actions include remote commands

Real-time graphing

- monitored items are immediately graphed using the built-in graphing functionality

Web monitoring capabilities

- Zabbix can follow a path of simulated mouse clicks on a web site and check for functionality and response time

Extensive visualization options

- ability to create custom graphs that can combine multiple items into a single view
- network maps
- slideshows in a dashboard-style overview
- reports
- high-level (business) view of monitored resources

Historical data storage

- data stored in a database
- configurable history
- built-in housekeeping procedure

Easy configuration

- add monitored devices as hosts
- hosts are picked up for monitoring, once in the database
- apply templates to monitored devices

Use of templates

- grouping checks in templates
- templates can inherit other templates

Network discovery

- automatic discovery of network devices
- agent autoregistration
- discovery of file systems, network interfaces and SNMP OIDs

Fast web interface

- a web-based frontend in PHP
- accessible from anywhere
- you can click your way through
- audit log

Zabbix API

- Zabbix API provides programmable interface to Zabbix for mass manipulations, third-party software integration and other purposes.

Permissions system

- secure user authentication
- certain users can be limited to certain views

Full featured and easily extensible agent

- deployed on monitoring targets
- can be deployed on both Linux and Windows

Binary daemons

- written in C, for performance and small memory footprint
- easily portable

Ready for complex environments

- remote monitoring made easy by using a Zabbix proxy

Glossary In this section you can learn the meaning of some terms commonly used in Zabbix.

host

- any physical or virtual device, application, service, or any other logically-related collection of monitored parameters.

host group

- a logical grouping of hosts. Host groups are used when assigning access rights to hosts for different user groups.

item

- a particular piece of data that you want to receive from a host, a metric of data.

value preprocessing

- a transformation of received metric value before saving it to the database.

trigger

- a logical expression that defines a problem threshold and is used to "evaluate" data received in items.

When received data is above the threshold, triggers go from *Ok* into a *Problem* state. When received data is below the threshold, triggers stay in/return to an *Ok* state.

template

- a set of entities (items, triggers, graphs, low-level discovery rules, web scenarios) ready to be applied to one or several hosts.

The job of templates is to speed up the deployment of monitoring tasks on a host; also to make it easier to apply mass changes to monitoring tasks. Templates are linked directly to individual hosts.

template group

- a logical grouping of templates. Template groups are used when assigning access rights to templates for different user groups.

event

- a single occurrence of something that deserves attention such as a trigger changing state or a discovery/agent autoregistration taking place.

event tag

- a pre-defined marker for the event. It may be used in event processing, permission granulation, etc.

complex event processing

- rule-based processing for Zabbix events, allowing to process or transform events.

problem

- a trigger that is in "Problem" state.

problem update

- problem management options provided by Zabbix, such as adding comment, acknowledging, changing severity or closing manually.

action

- a predefined means of reacting to an event.

An action consists of operations (e.g., sending a notification) and conditions (*when* the operation is carried out)

escalation

- a custom scenario for executing operations within an action; a sequence of sending notifications/executing remote commands.

macro

- a variable placeholder that resolves to a specific value depending on the context.

media

- a means of delivering notifications; delivery channel.

notification

- a message about some event sent to a user via the chosen media channel.

remote command

- a pre-defined command that is automatically executed on a monitored host upon some condition.

web scenario

- one or several HTTP requests to check the availability of a web site.

frontend

- the web interface provided with Zabbix.

dashboard

- customizable section of the web interface displaying summaries and visualizations of important information in visual units called widgets.

widget

- visual unit displaying information of a certain kind and source (a summary, a map, a graph, the clock, etc.), used in the dashboard.

Zabbix API

- Zabbix API allows you to use the JSON RPC protocol to create, update and fetch Zabbix objects (like hosts, items, graphs and others) or perform any other custom tasks.

Zabbix server

- a central process of Zabbix software that performs monitoring, interacts with Zabbix proxies and agents, calculates triggers, sends notifications; a central repository of data.

Zabbix proxy

- a process that may collect data on behalf of Zabbix server, taking some processing load from the server.

Zabbix agent

- a process deployed on monitoring targets to actively monitor local resources and applications.

Zabbix agent 2

- a new generation of Zabbix agent for actively monitoring local resources and applications, allowing the use of custom plugins for monitoring.

> **Attention:**
> Because Zabbix agent 2 shares much functionality with Zabbix agent, the term ”Zabbix agent” in documentation stands for both - Zabbix agent and Zabbix agent 2, if the functional behavior is the same. Zabbix agent 2 is only specifically named where its functionality differs.

encryption

- support of encrypted communications between Zabbix components (server, proxy, agent, zabbix_sender and zabbix_get utilities) using Transport Layer Security (TLS) protocol.

agent autoregistration

- automated process whereby a Zabbix agent itself is registered as a host and started to monitor.

network discovery

- *automated discovery of network devices*.

low-level discovery

- *automated discovery of low-level entities on a particular device* (e.g., file systems, network interfaces, etc).

low-level discovery rule

- *set of definitions for automated discovery of low-level entities* on a device.

item prototype

- *a metric with certain parameters as variables, ready for low-level discovery*. After low-level discovery the variables are automatically substituted with the real discovered parameters and the metric automatically starts gathering data.

trigger prototype

- *a trigger with certain parameters as variables, ready for low-level discovery*. After low-level discovery the variables are automatically substituted with the real discovered parameters and the trigger automatically starts evaluating data.

Prototypes of some other Zabbix entities are also in use in low-level discovery - graph prototypes, host prototypes, host group prototypes.

Zabbix Cloud

Zabbix Cloud

Discover the available configuration options and access controls for organizations and nodes in Zabbix Cloud.

Please note that Zabbix Cloud nodes currently run **Zabbix 7.0**.
 For in-depth information on Zabbix 7.0, see official documentation.

Set up and manage Zabbix Cloud

Start using Zabbix Cloud Guided instructions to get your Zabbix Cloud instance up and running quickly, including initial configuration and connecting your monitored devices.

Node configuration Learn how to set up your Zabbix Cloud node, focusing on access, data handling, and resource settings.

Adding users Find out how to add and assign roles to users within a Zabbix Cloud organization and its nodes.

Learn about Zabbix Cloud

Zabbix Cloud vs. on-premises Compare the differences in management, scalability, and maintenance between the two deployments, and learn which can support your workflow most seamlessly.

Explore Zabbix Cloud Get an overview of Zabbix Cloud, including its key features, perks, and what comes with a subscription. For common questions on setup and usage, see the FAQ.

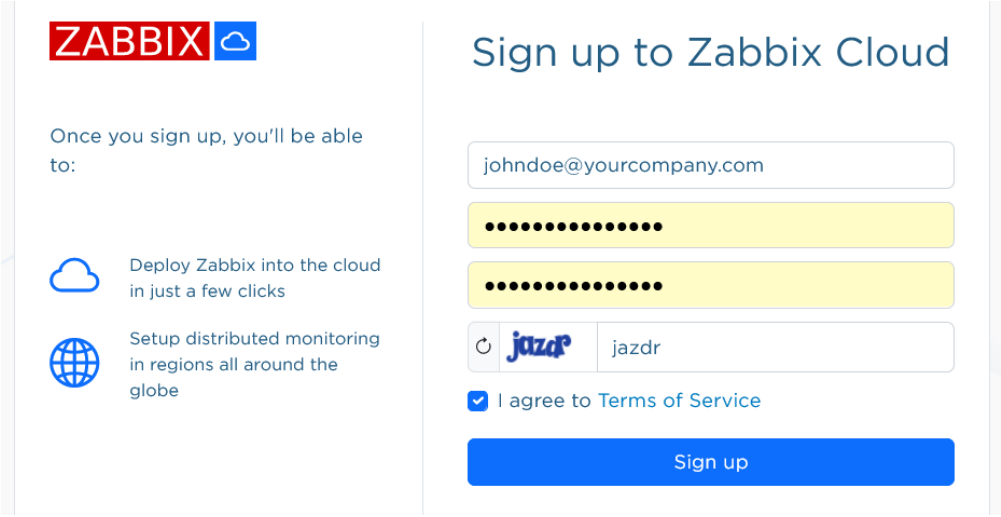
Deploy Zabbix in the cloud

To start Zabbix in the cloud, follow this process:

1. Log in or sign up at cloud.zabbix.com.
2. Create a cloud node.
3. Start using Zabbix in the cloud.

Log in or sign up You can log in at cloud.zabbix.com with an existing Google, GitHub, or Microsoft account.

Alternatively, go to the sign up page to register.



After signing up, you will receive a confirmation of your new account in the email.

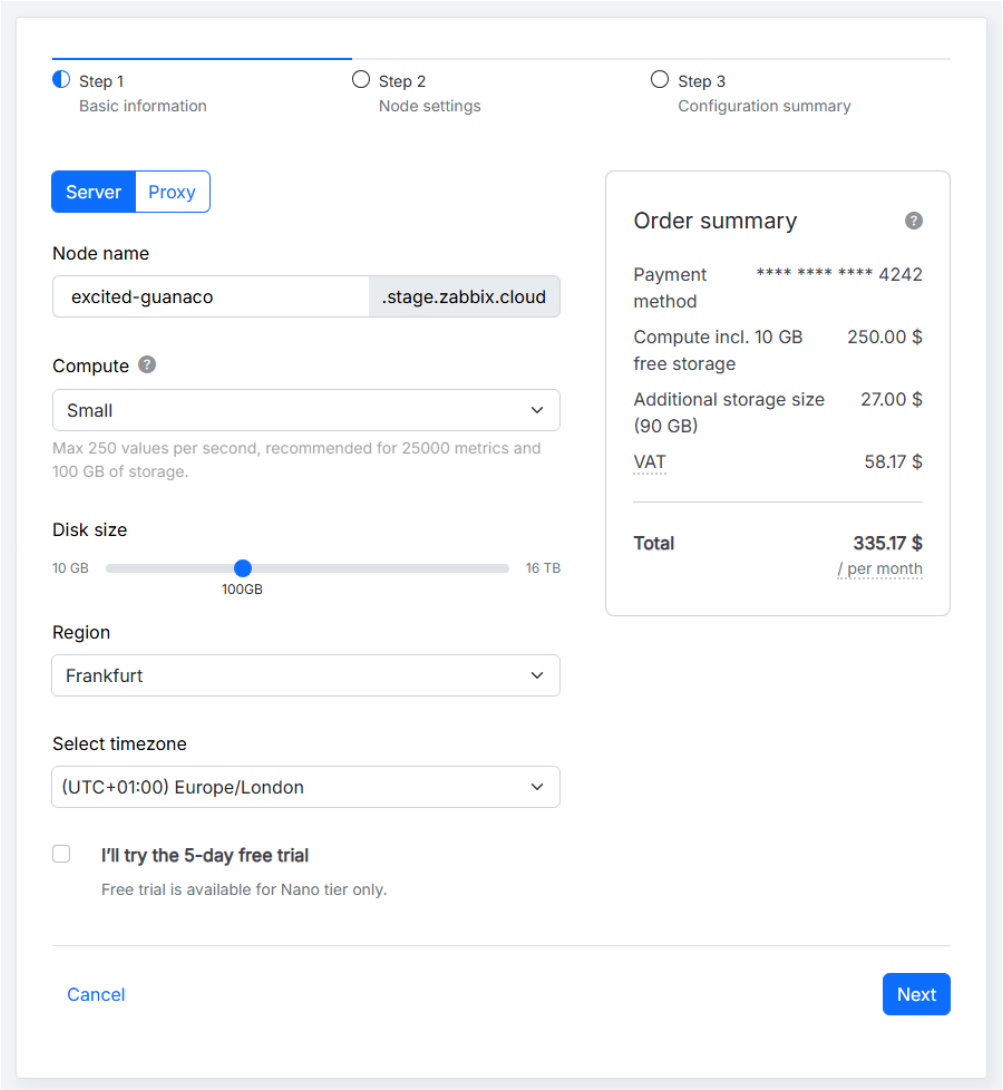
Go to cloud.zabbix.com to log in. You will be asked for a one-time password when logging in. This password will be sent to you in the email.

Create cloud node A cloud node is a Zabbix instance deployed in Zabbix Cloud. It can be deployed as a server node or proxy node.

When logged in, you will see that a *Private* organization already exists for your account by default. You can rename the default organization later.

Click *Create New Node*. A three-step node configuration form will open.

In the first step you will select basic information.



Parameter	Description
Server/Proxy	Select whether to deploy a Zabbix server instance or a Zabbix proxy instance.
Node name	The node name forms a unique subdomain in the format `<your-name>.zabbix.cloud`, which is used to access the node's services.
Compute	Select the computing power for the node. To determine what level of resources you need, see pricing details. You can always upgrade later if you need more power. A billing method will be required before creating a non-trial node. Add billing information via the *Billing* menu. A 5-day **free trial** is available for trying a *Nano* tier node. You will not be charged during the trial period, but once it ends, your node will be removed if not converted to a paid node.
Disk size	The slider defaults to a recommended minimum based on tier. Increase it if you store large volumes of history or trends. Be aware that larger disks incur higher monthly costs and that you cannot shrink a disk yourself—you'd need to contact support to reduce it.
Region	Select the region where the node will be deployed.
Select timezone	Select the timezone of the Zabbix server. This timezone is used by the Zabbix server and is independent of the maintenance and user profile timezones.

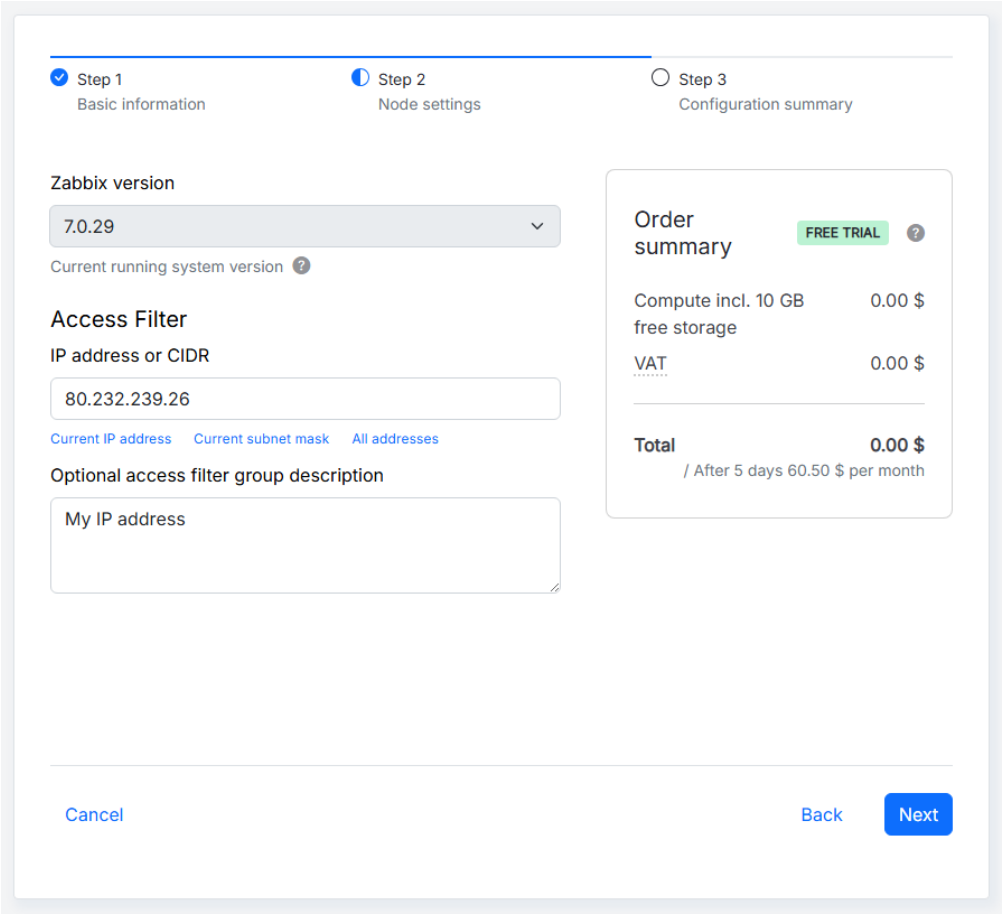
After completing the configuration, click *Next*.

In the second step you will select the node settings.

The configuration options in the second step depend on whether you are creating a server or proxy node.

Configure a server node

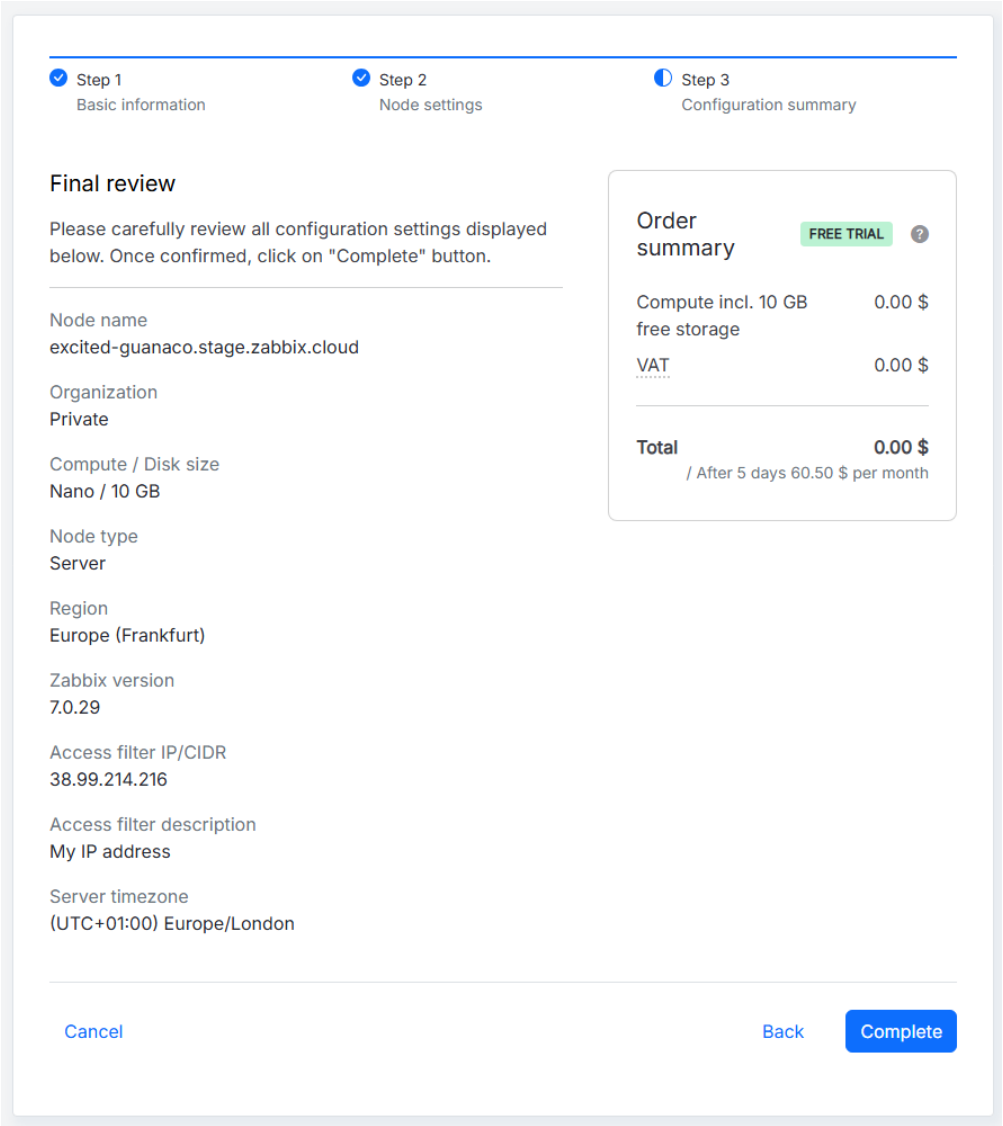
During server node creation, the following parameters are available:



Parameter	Description
Zabbix version	Select the Zabbix version. Zabbix Cloud supports the latest LTS version and automatically applies minor version updates.
IP address or CIDR	Specify the initial access filter for the node. Enter the IP address or CIDR manually, or complete the field automatically by selecting the relevant option. For server nodes, this filter applies to connections to the server, API, and frontend. The filter can be changed later in the Access filters tab.
Optional access filter group description	Add an optional description for the access filter group, which can contain multiple IP addresses or CIDRs.

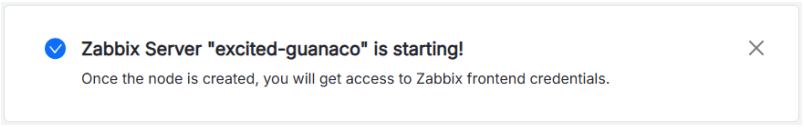
Warning:
During node creation, the *IP address or CIDR* field is prefilled with the IP address used to access Zabbix Cloud. You can change this value to specify a different initial access filter. After the node is created, you can add, edit, or delete access filters from the *Access filters* tab.

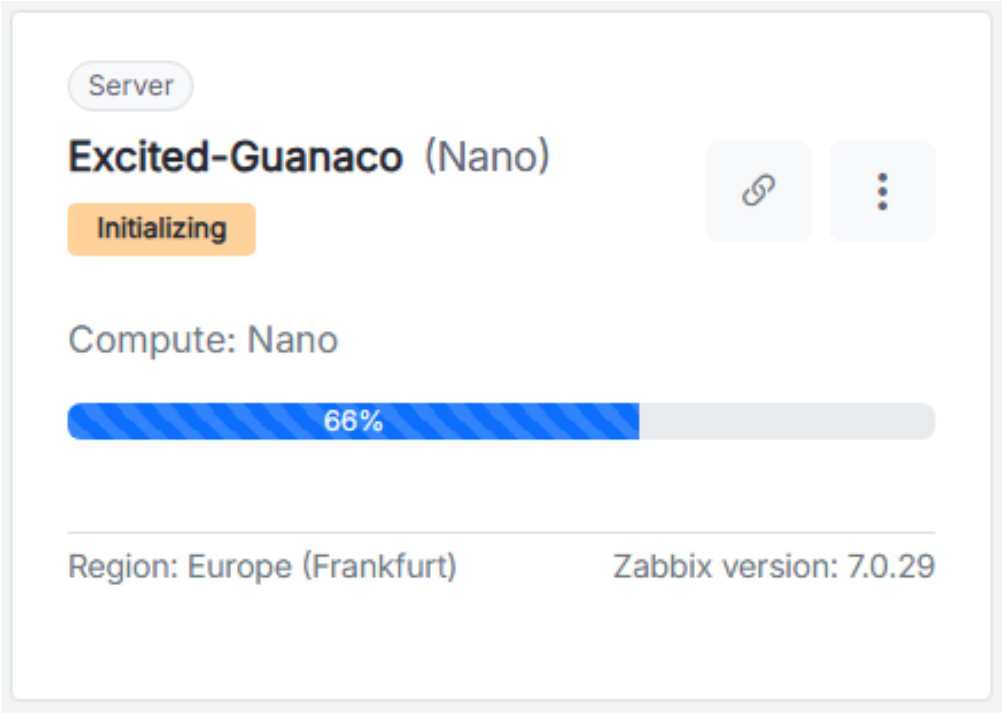
The third step displays a configuration summary for the node.



After reviewing configuration settings, and confirming everything is correct, click *Complete* for a trial node or *Proceed to payment* for a paid node. Otherwise, return to a previous step by clicking the *Back* button, or *Cancel* to start from beginning.

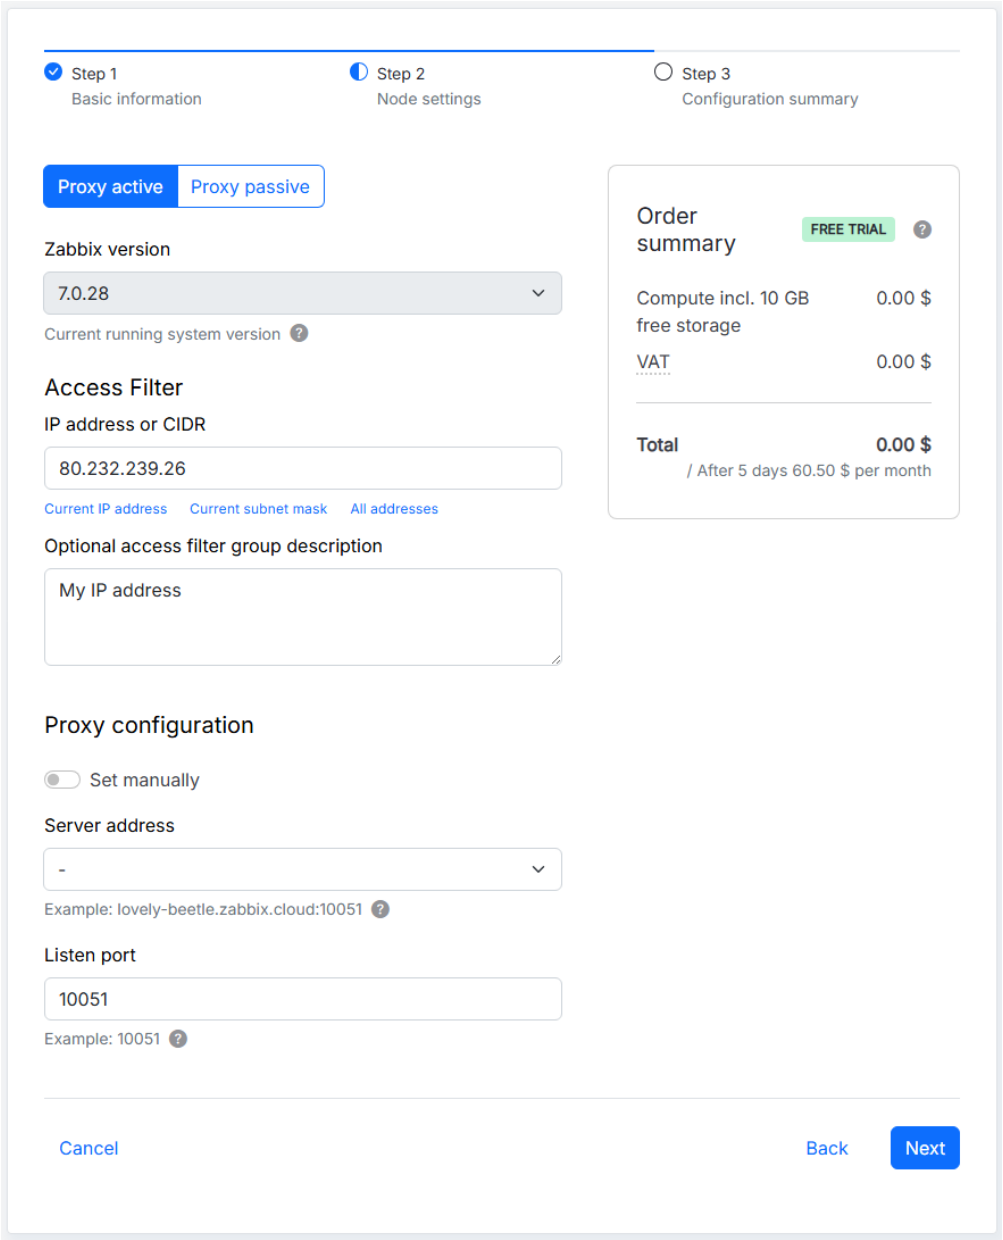
As the node is being initialized, a corresponding message is displayed and the node already appears under *Organizations*.





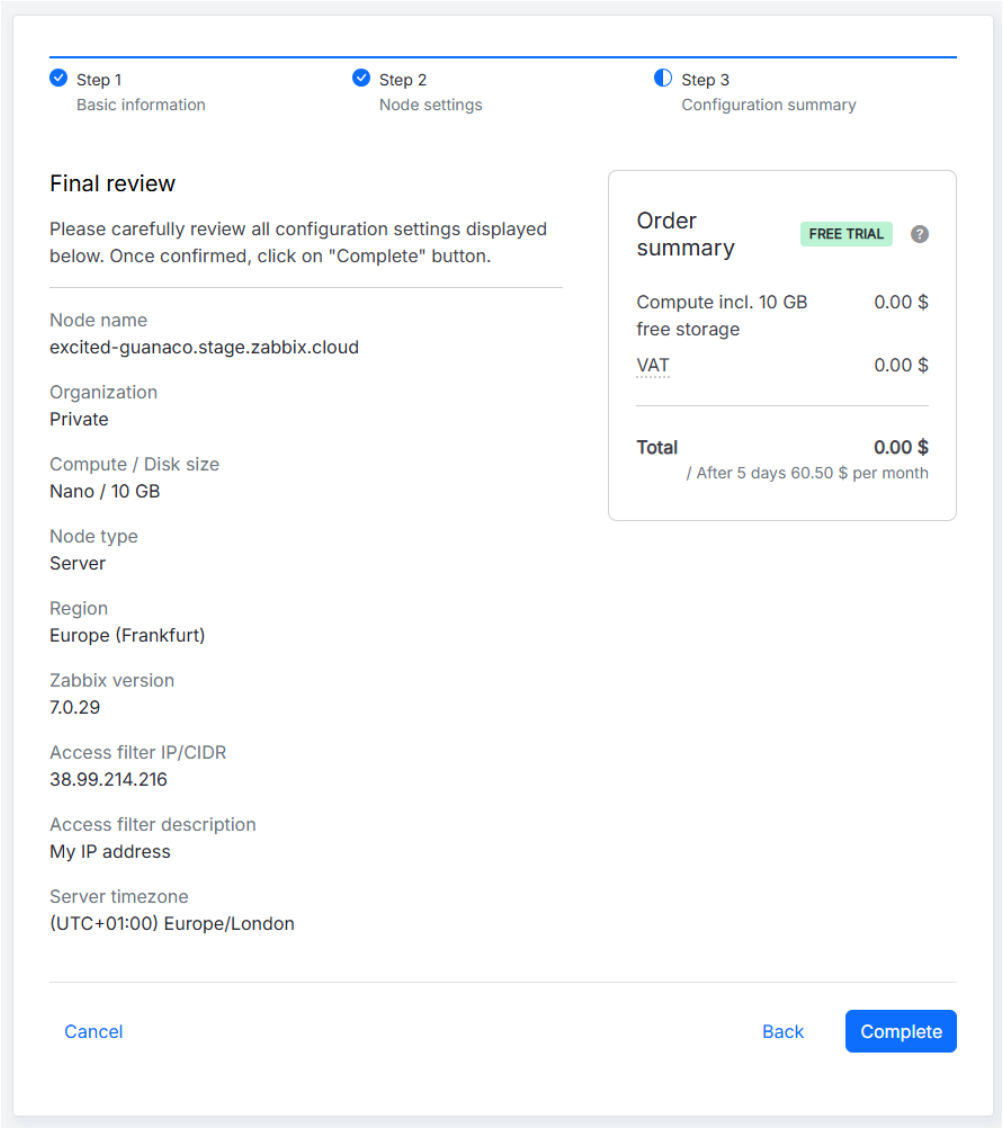
Configure a proxy node

During the proxy node creation the following parameters are available:



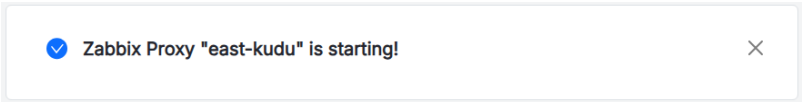
Parameter	Description
Proxy active/Proxy passive	Select the proxy mode. In active mode, the proxy initiates connections to Zabbix server; in passive mode, Zabbix server initiates connections to the proxy.
Zabbix version	Select the Zabbix version. Zabbix Cloud supports the latest LTS version and automatically applies minor version updates.
IP address or CIDR	Specify the initial inbound IP access filter for the proxy. Enter the IP address or CIDR manually, or complete the field automatically by selecting the relevant option. The filter defines which network requests are allowed to reach the proxy and can be changed later in the Access filters tab.
Optional access filter group description	Add an optional description for the access filter group, which can contain multiple IP addresses or CIDRs.
Set manually	Select how to specify the Zabbix server. When disabled, select a server node from the list of server nodes available to you. When enabled, specify server addresses manually.
Server address	Specify the Zabbix server address. When *Set manually* is disabled, select a server node from the list. When enabled, specify up to 10 server addresses. For an active proxy, specify the IP address or DNS name of the Zabbix server. For a passive proxy, specify the IP address, DNS name, or CIDR of the Zabbix server. Access filters are configured separately and are not updated automatically when a server address is specified.
Listen port	Specify the port on which the proxy listens for incoming connections. The default port is 10051. The listen port cannot be changed after node creation.

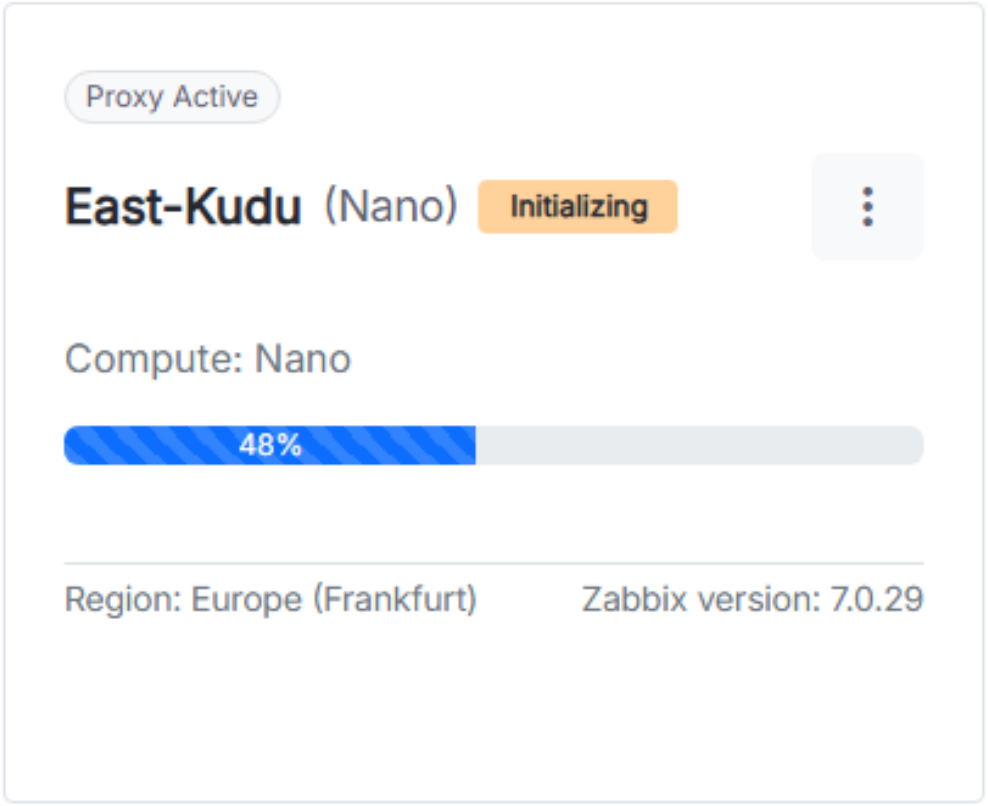
In the third step you will see the configuration summary for the node.



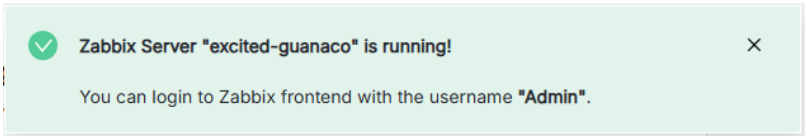
After reviewing the configuration settings and confirming that everything is correct, click *Complete* for a trial node or *Proceed to payment* for a paid node. Otherwise return to previous steps by clicking the *Back* button, or *Cancel* to start from beginning.

As the node is being initialized, a corresponding message is displayed and the node already appears under *Organizations*.





To configure a Zabbix Cloud proxy, see Set up a Zabbix Cloud proxy.

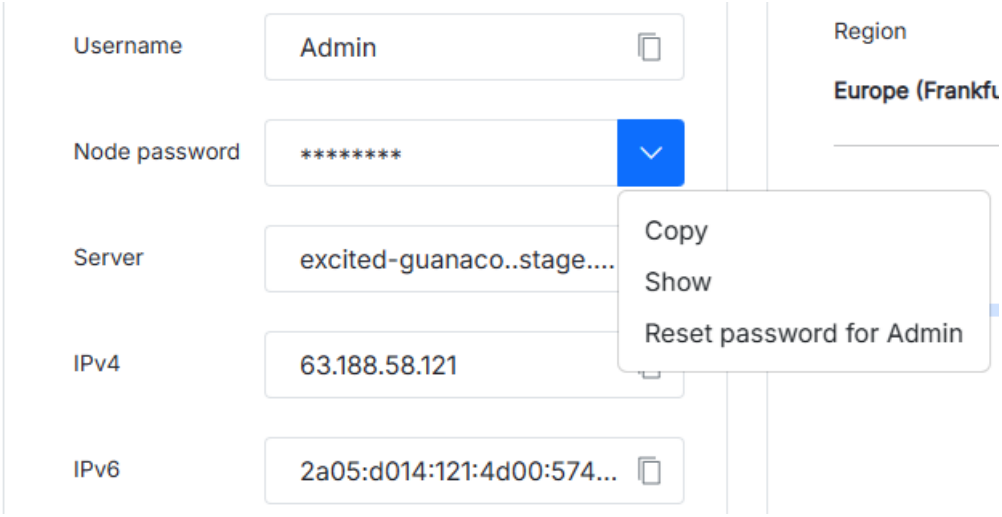


Start using Zabbix

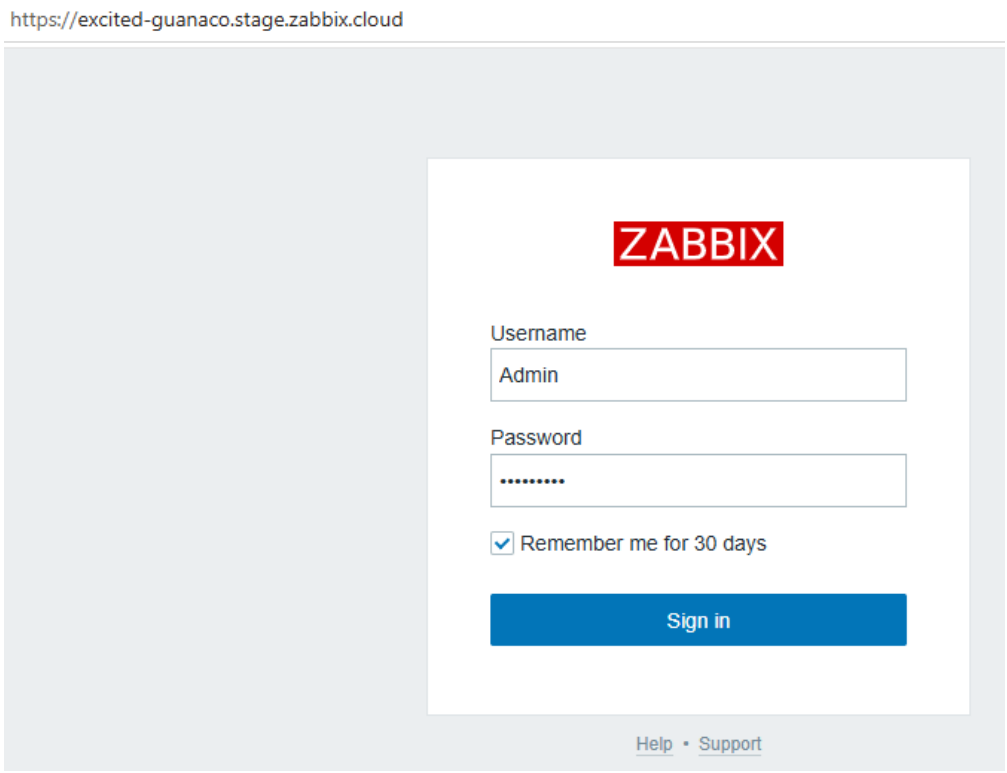
When the server node status changes from *Initializing* to *Running*, you can log in to the Zabbix frontend.

To access the initial frontend password, click the created node.

You will be taken to node settings (overview, access filters, encryption, backups, etc.). In **Connection info**, locate *Node password* field and select *Reset password for Admin* from the dropdown. After the password is reset, click the same dropdown and select *Copy*.



Click the frontend URL to open Zabbix frontend (`excited-guanaco.zabbix.cloud`).



Log in using the following credentials:

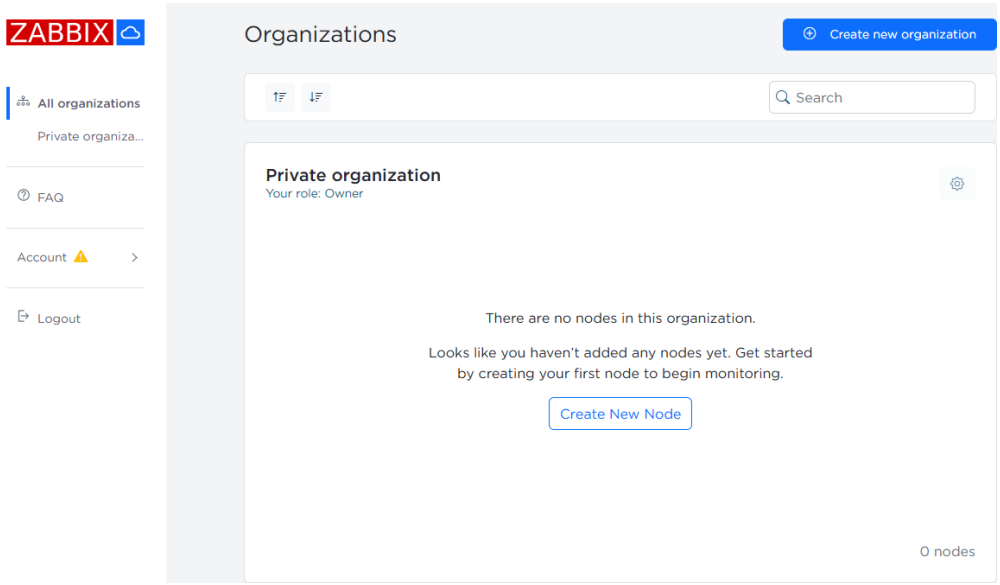
- *Username* - `Admin`
- *Password* - provide the copied password

To start monitoring, you need to define access filters - IP address or addresses that are allowed to connect to the Zabbix Cloud server. At least one must be defined.

Organization management An organization is a container for one or more nodes. Organizations also help you structure your resources and delegate responsibilities, such as:

- Node management
- Billing

When you log in to Zabbix Cloud, the *All organizations* page is displayed. By default, a private organization is created for your account.

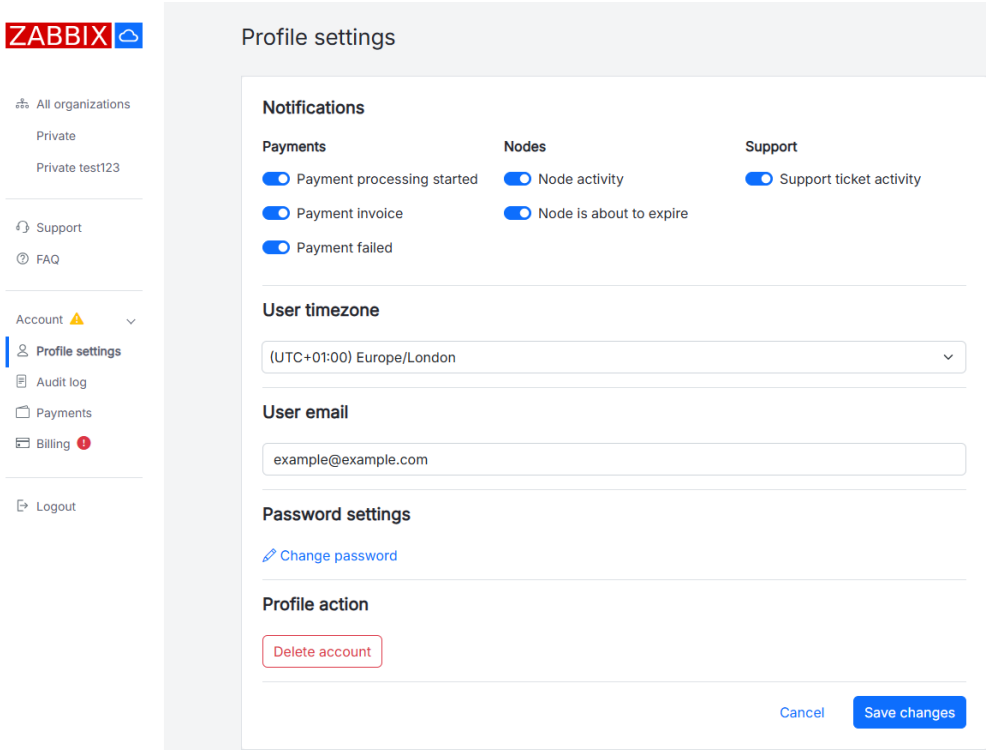


As the owner of this organization, you can manage users, admins, nodes, and the organization name. Click the  settings button to access management options.

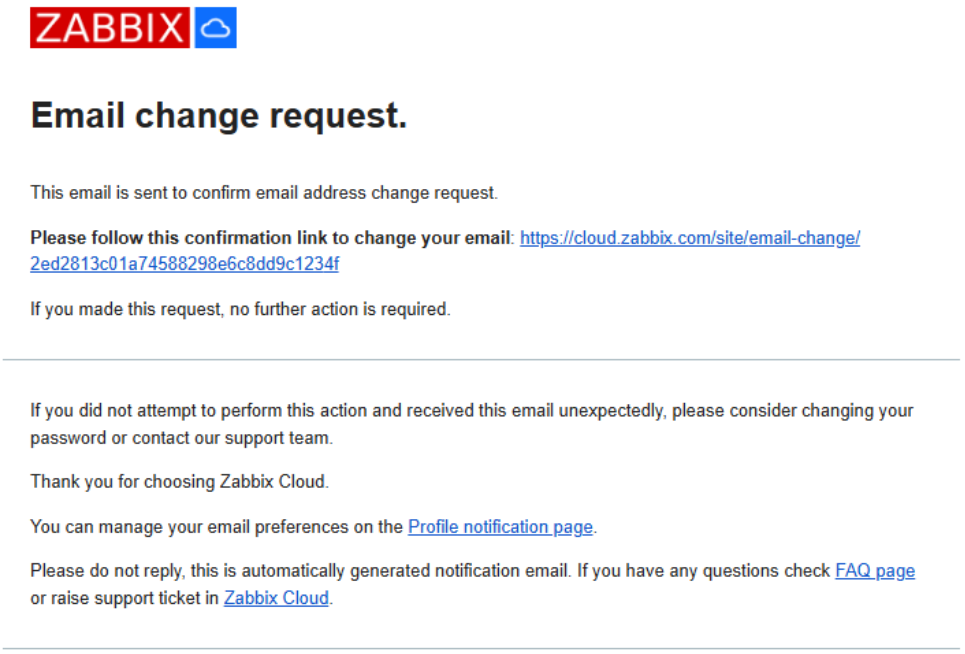
Account management Changing the account email address

To change the account email address, navigate to *Account* > *Profile settings*.

In the *User email* field, enter the new email address and click *Save changes*.



A confirmation message is sent to the new email address. Follow the link in the message to confirm the email address change.



The confirmation link opens the *Profile settings* page with a confirmation message and the updated email address in the User email field.

After the change is confirmed, an email confirming the change is sent to both the new and previous email addresses. The message sent to the previous address includes the previous and new email addresses.

Node configuration

Overview

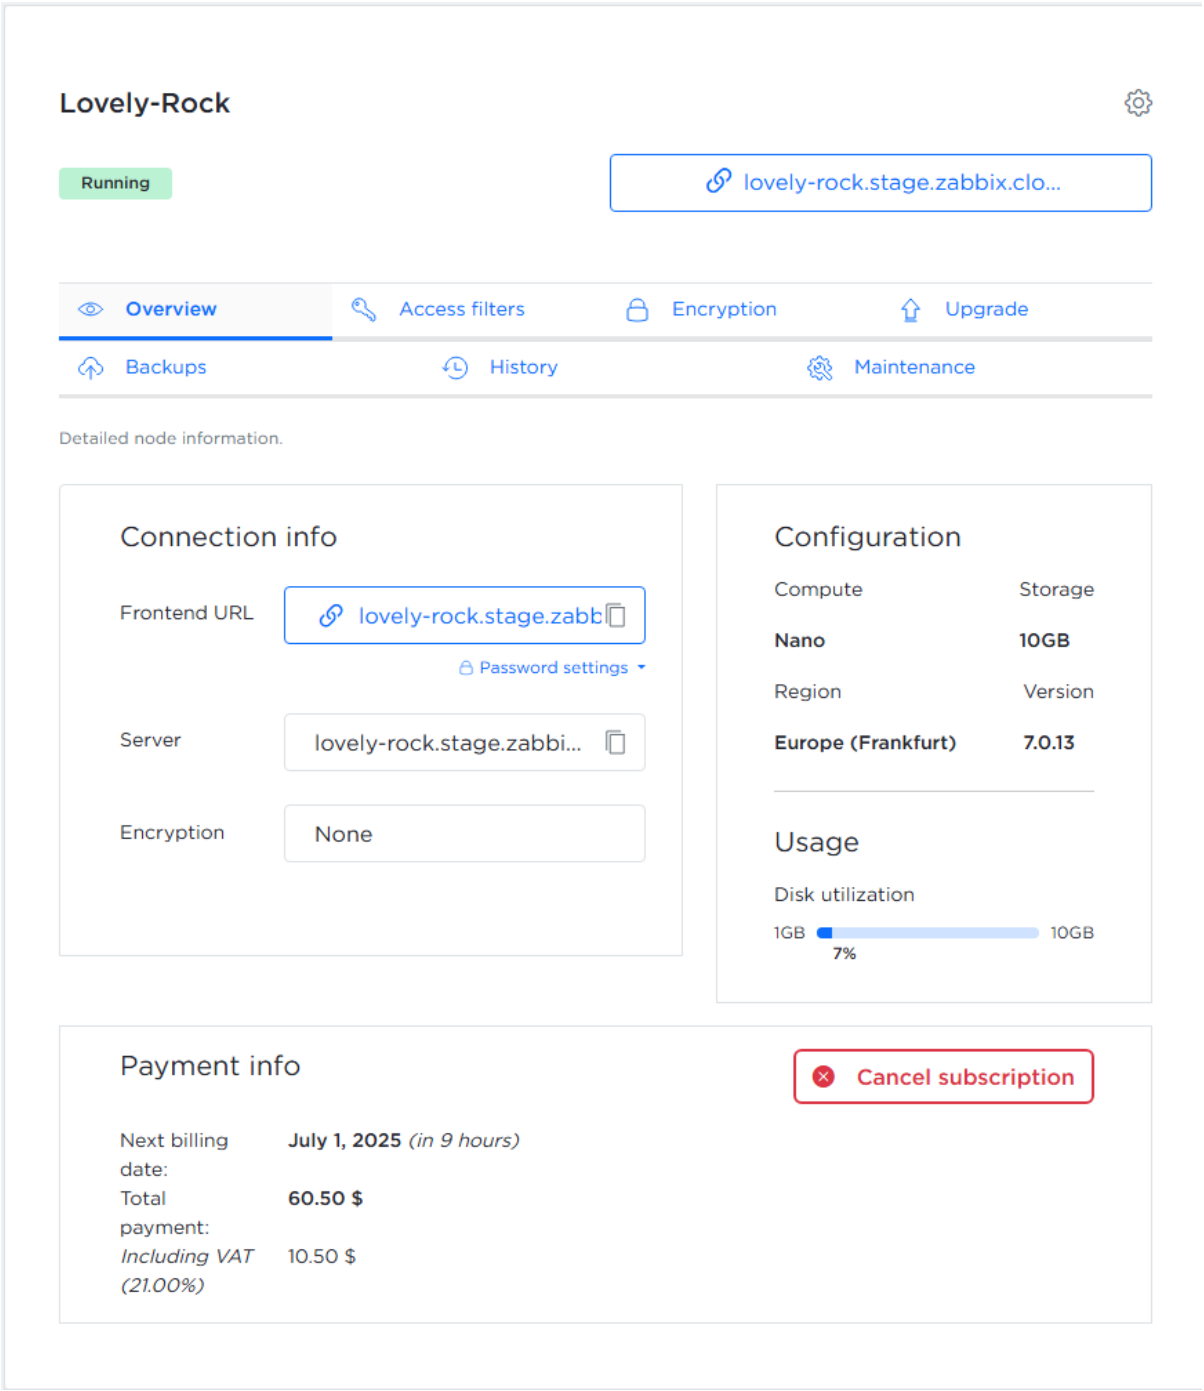
A **cloud node** is a Zabbix component deployed in the cloud. Server nodes include the Zabbix server, database, and web frontend, while proxy nodes include the Zabbix proxy and database. A cloud node has dedicated computing resources (CPU, memory) and disk storage.

Users may create one or more cloud nodes. Cloud nodes can be grouped into **organizations** (structural, regional, etc), creating a distributed monitoring setup.

Cloud nodes and organizations are created using the Zabbix Cloud user interface. To access the node configuration page, click the node from *All organizations*, your organization menu, or the *Node configuration page* link shown after the node is launched.

In the node configuration view, each tab lets you manage a specific aspect of your node:

- Overview - view node details, open its frontend, and manage node users.
- Access filters - control access by IP address or subnet.
- Encryption - configure TLS by uploading certificates, keys, and optional CRLs.
- Backups - create, restore, and delete node snapshots; view the next scheduled backup.
- History - set retention periods for audit logs, history, and trends, with preset units and a 3x/day update limit.
- Maintenance - define weekly maintenance windows by time zone, day, and hour.
- Upgrade - scale compute/storage resources or manage your trial/payment status.

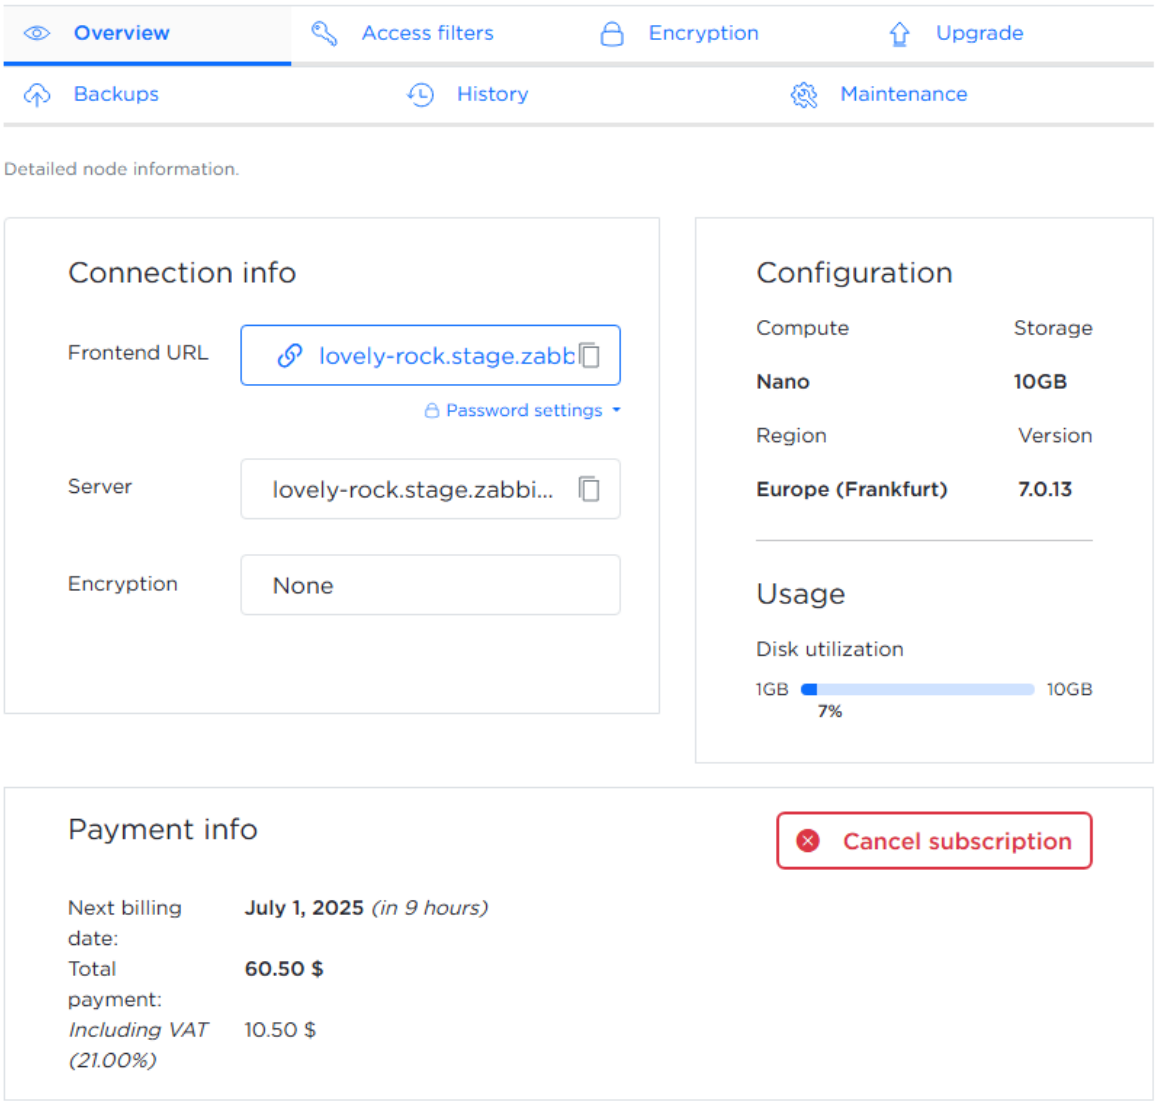


Note:
Some tabs may be unavailable depending on the node type (server or proxy), type of user who views the page and node status (trial/non-trial).

Overview tab

The *Overview* tab provides detailed information about the node. The information displayed depends on whether the node is a server or proxy.

Server node The *Overview* tab for the server node contains the following sections:



Connection info:

- Clickable frontend URL with a copy-to-clipboard option
- Username
- Node password and its settings
- Server hostname with a copy-to-clipboard option
- IPv4 address with a copy-to-clipboard option
- IPv6 address with a copy-to-clipboard option
- Current encryption method

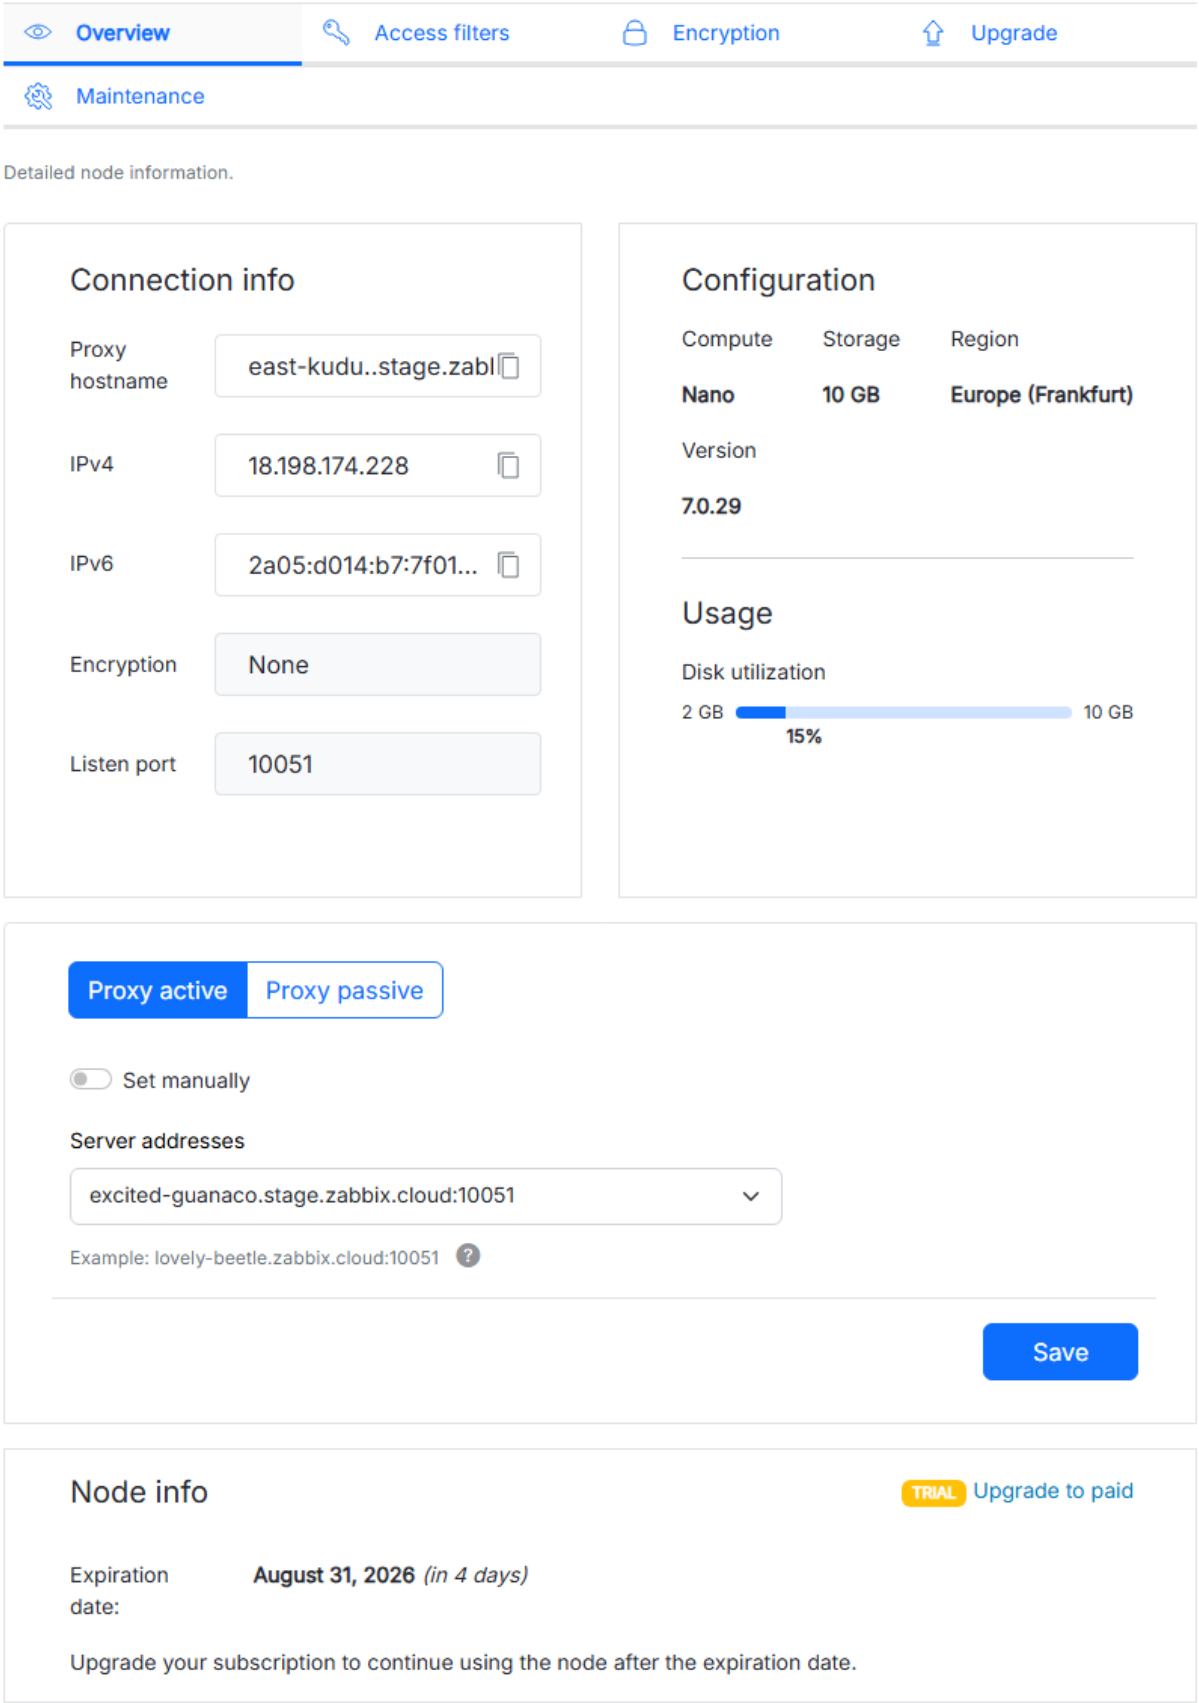
Configuration and usage:

- Compute tier, maximum storage, region
- Zabbix components version
- Disk utilization

Node info/Payment Info:

- Expiration date (for trial and prepaid nodes)
- A context-specific button is shown depending on the node type after the node completes initialization: - *Upgrade to paid* - available for nodes on a free trial - *Cancel subscription* - available for paid nodes with an active subscription - *Resubscribe to paid* - shown if a subscription has been canceled - No action button is displayed for prepaid nodes, which have long validity and cannot be canceled

Proxy node The *Overview* tab for the proxy node contains the following sections:



Connection info:

- Proxy hostname with a copy-to-clipboard option
- IPv4 address with a copy-to-clipboard option
- IPv6 address with a copy-to-clipboard option
- Current encryption method
- Current listen port

Configuration and usage:

- Compute tier, maximum storage, region
- Zabbix components version
- Disk utilization

Proxy mode:

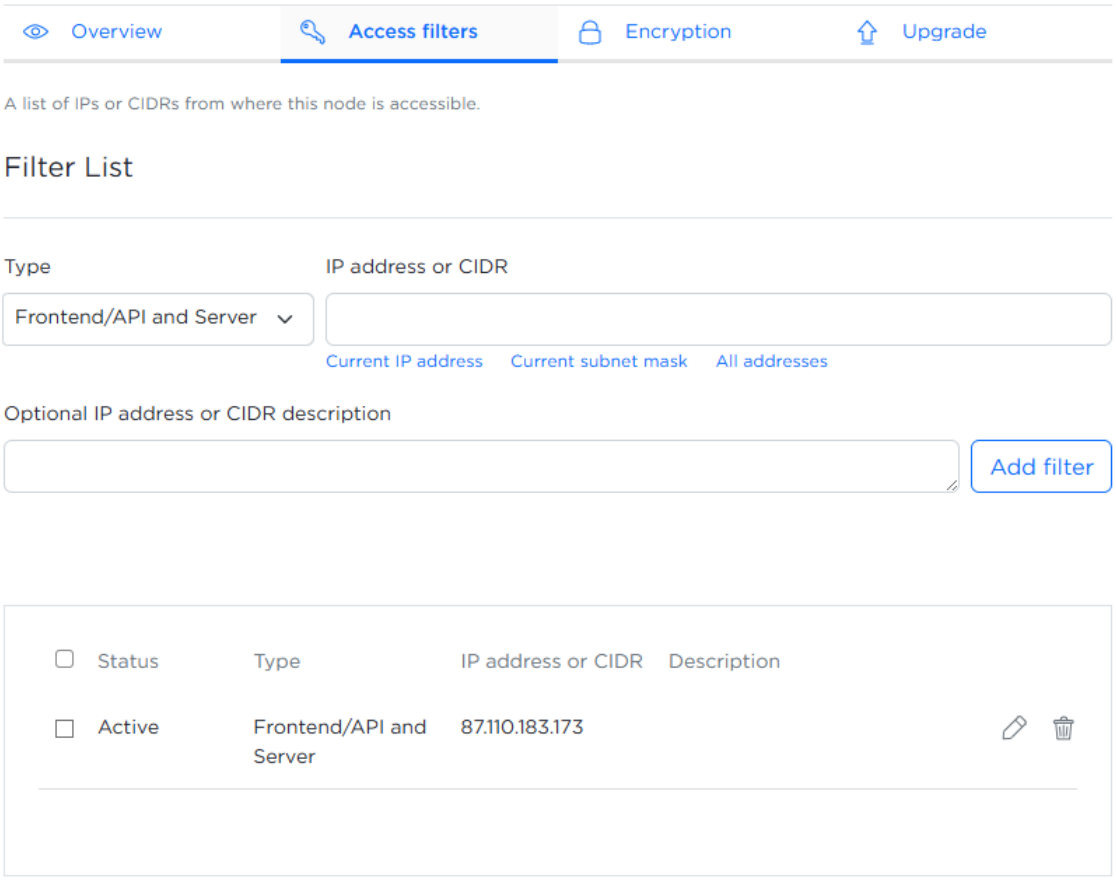
- Switch between active and passive proxy modes
- Select the server address from the *Server addresses* dropdown, or enable *Set manually* to enter the server address manually

Node info/Payment Info:

- Expiration date (for trial and prepaid nodes)
- A context-specific button is shown depending on the node type after the node completes initialization: - *Upgrade to paid* - available for nodes on a free trial - *Cancel subscription* - available for paid nodes with an active subscription - *Resubscribe to paid* - shown if a subscription has been canceled - No action button is displayed for prepaid nodes, which have long validity and cannot be canceled

Access filters tab

The *Access filters* tab lets you restrict access to your Zabbix node by IP address or subnet.



To configure an access filter, complete the following fields:

- *Type* - for server nodes, use the dropdown to choose which Zabbix component to limit: frontend/API only, server only, or both. For proxy nodes, only *Inbound IP* is available. It defines which incoming network requests are allowed to reach the proxy.
- *IP address or CIDR* - enter the IP address or CIDR manually, or complete the field automatically by selecting the relevant option.
- *Optional IP address or CIDR description* - add an optional description for each IP or subnet.

After completing the fields, click *Add filter*. Added filters appear in the table below, showing status, type, IP address or subnet, and description.

Warning:
During node initialization, an access filter is automatically created for the IP address you used. If you later connect from a different IP or network, you must add that address here or you will be blocked.

You can:

-  - edit a filter.
-  - delete a filter or multiple filters using the checkboxes.

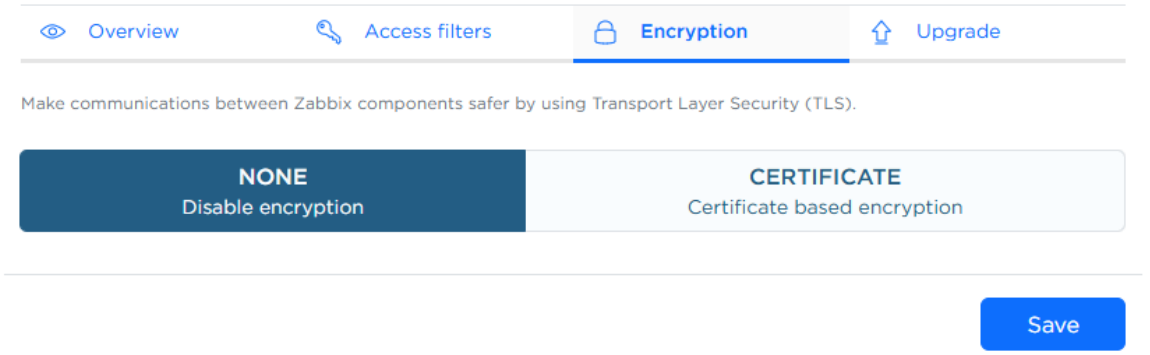
Encryption tab

The *Encryption* tab lets you configure TLS encryption for communication between Zabbix components by uploading the required certificate files.

To enable secure communication using certificates, select *CERTIFICATE* and upload the following files from your computer:

- *Root CA certificates* - trusted CA certificates used to verify peer certificates. Click *Choose file* and select the file containing trusted Root CA certificates (typically a `.crt` or `.pem` file, PEM-encoded).
- *Certificate chain* - represents the chain of trust for your own certificate. Click *Choose file* and select the PEM-encoded file (e.g., `.crt` or `.pem`) that includes your certificate and any intermediate CA certificates.
- *Private key* - the private key corresponding to your certificate. Click *Choose file* and select the PEM-encoded `.key` file.
- *Revoked certificates* (optional) - if you maintain a Certificate Revocation List (CRL), click *Choose file* and select the CRL file to specify certificates that should no longer be trusted.

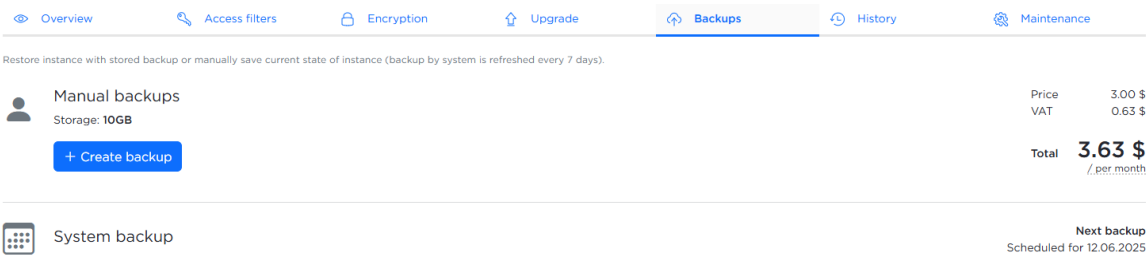
After selecting the files, click *Save*. The files will be uploaded, validated, and applied to configure TLS encryption between Zabbix components.



Backups tab

The *Backups* tab lets you view, create, restore, and delete backups of your Zabbix node.

You can restore your instance from any saved backup or create a new manual backup at any time. The tab also shows the date of the next scheduled backup.



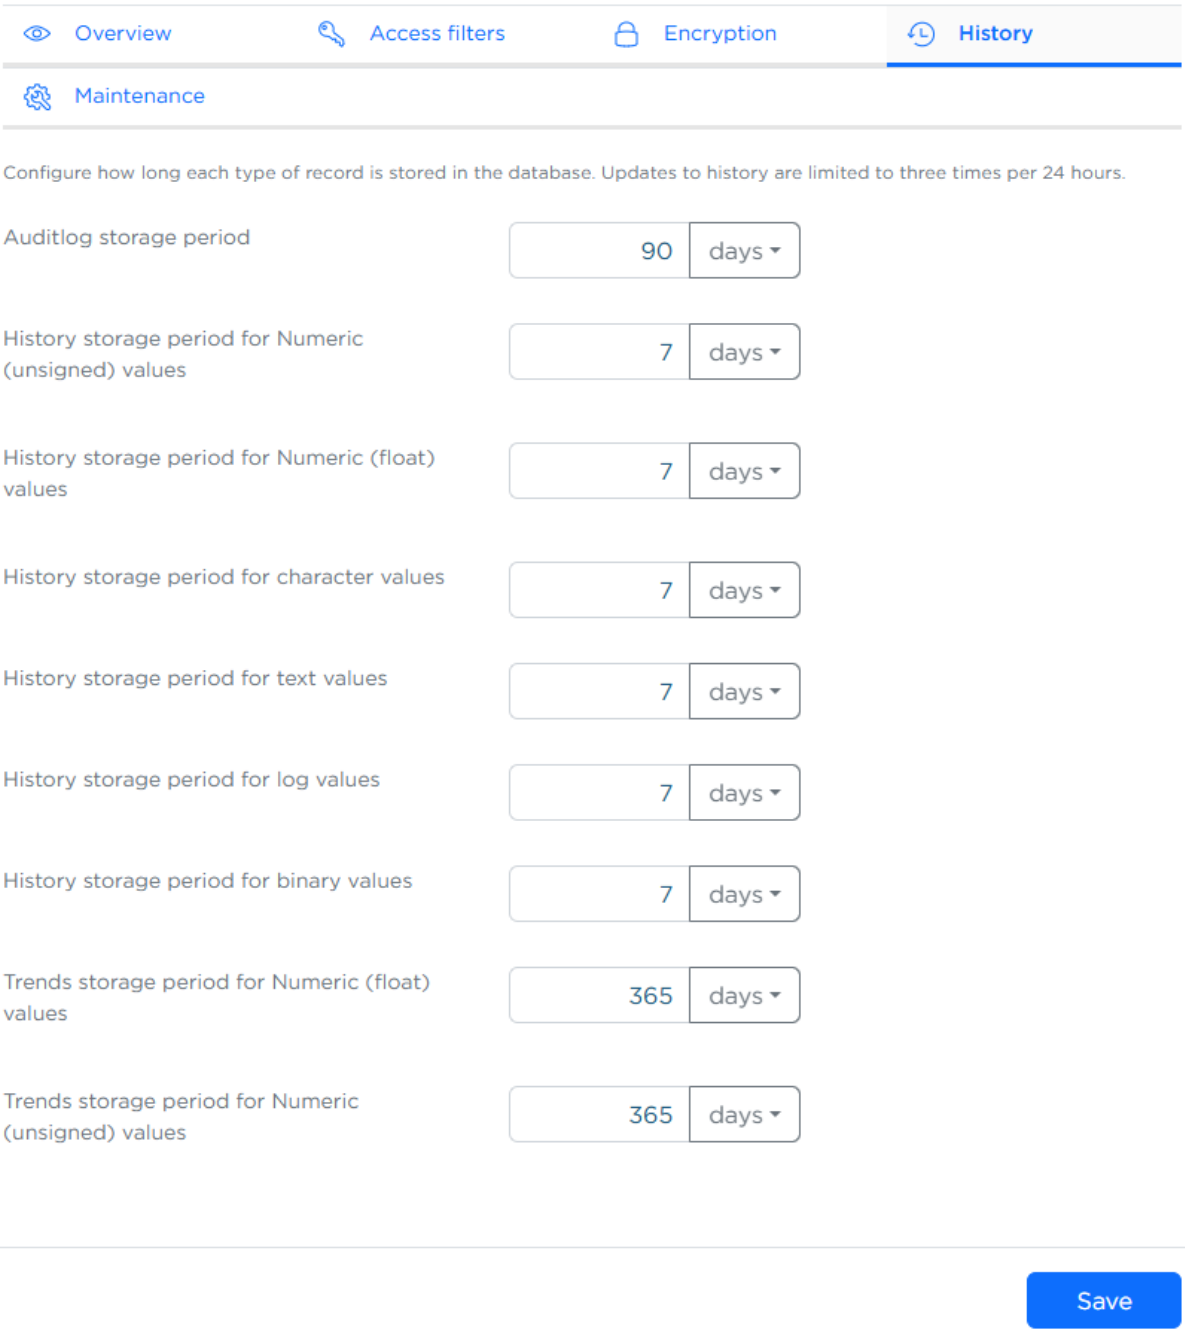
To create a backup:

1. Click *Create backup* button.
2. Enter an optional description in the popup window.
3. Click *Create* to confirm.

Warning:
Additional backups are charged monthly based on storage usage.

History tab

The *History* tab lets you configure how long different types of historical data are stored.



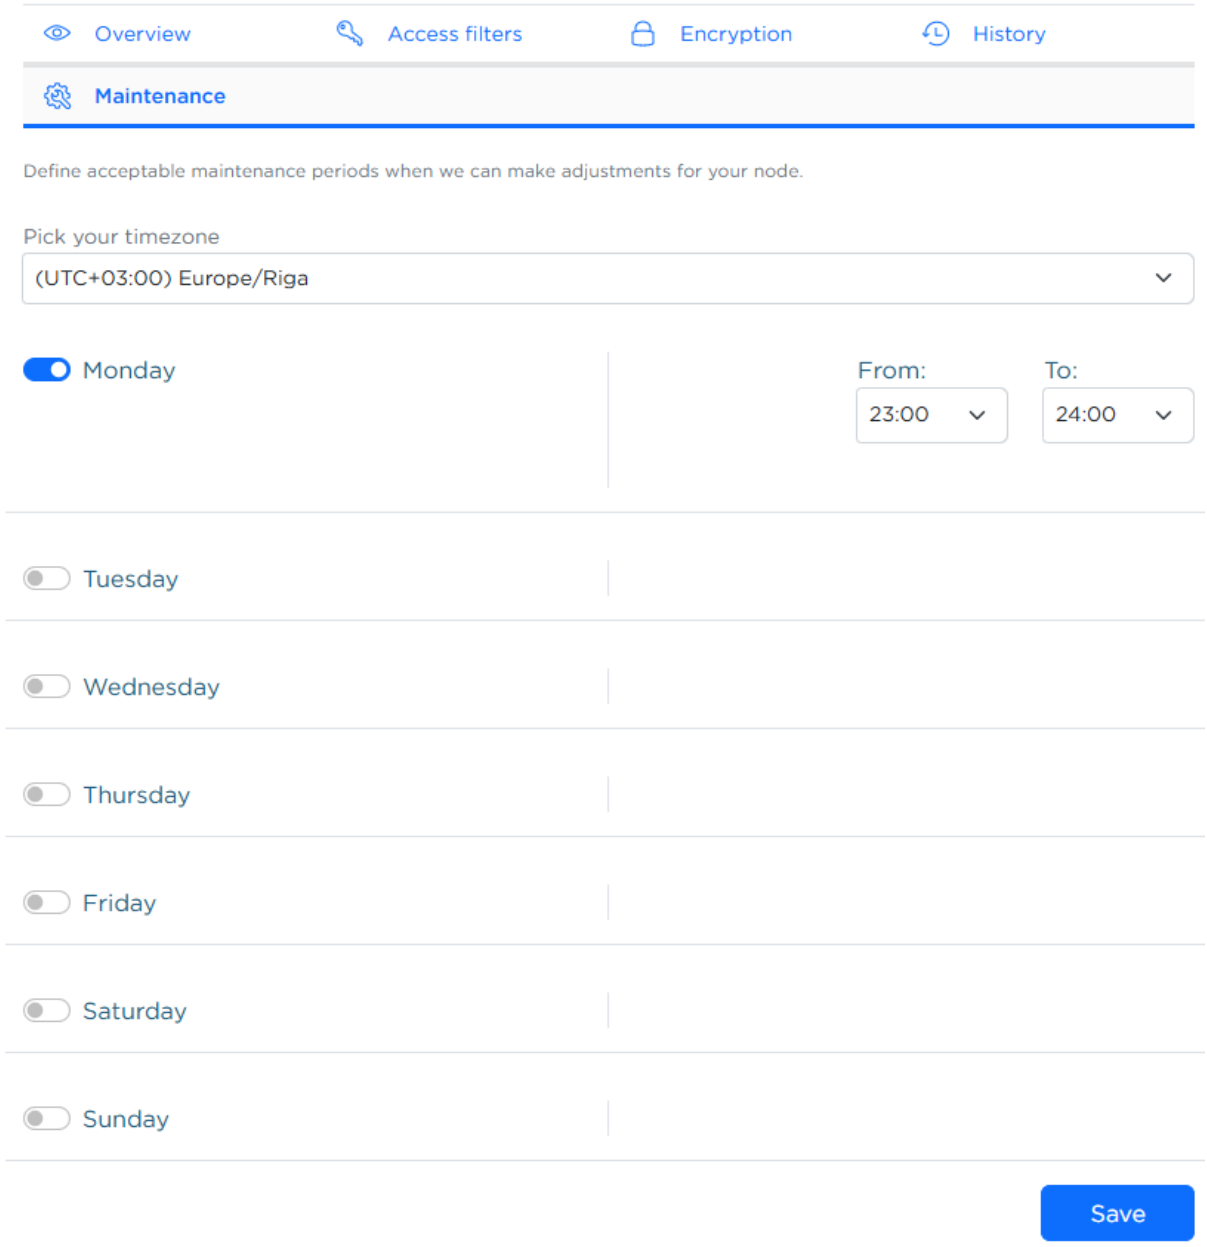
You can set the storage period for each data type (e.g., numeric, text, logs) and choose the unit (days, months, years). Changes can be made up to three times every 24 hours.

After updating the values, click *Save* to apply the changes.

Warning:
Longer retention periods can increase disk usage and may impact performance. It is recommended to configure retention periods based on your actual data and analysis needs.
 Zabbix Cloud uses database partitioning to manage data storage instead of traditional Zabbix housekeeping. This enables faster performance and more predictable retention handling.

Maintenance tab

The *Maintenance* tab lets you define acceptable maintenance windows when updates or adjustments can be performed on your node.

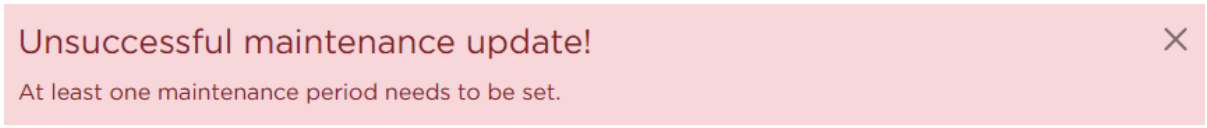


You must:

- Select your time zone from the dropdown.
- Choose at least one day and one hourly window for maintenance.

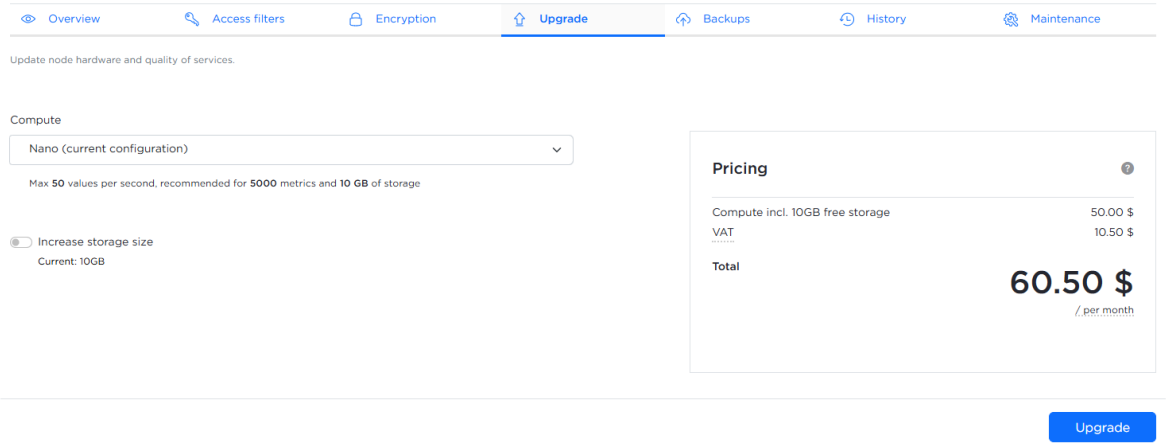
Click *Save* to apply changes.

If no day or time window is selected, an error will be displayed.



Upgrade tab

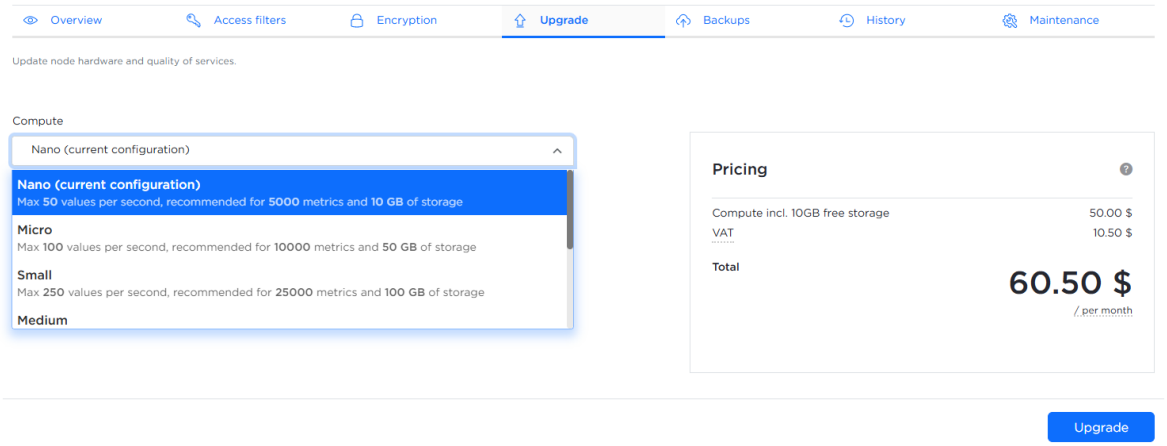
The *Upgrade* tab lets you scale your node's resources or manage its subscription.



You may increase the compute tier or disk size. Note that disk size can only be increased once every 24 hours and cannot be reduced via the UI (contact Zabbix support for any shrink requests). Free-trial nodes can be converted to paid subscriptions.

To increase compute tier:

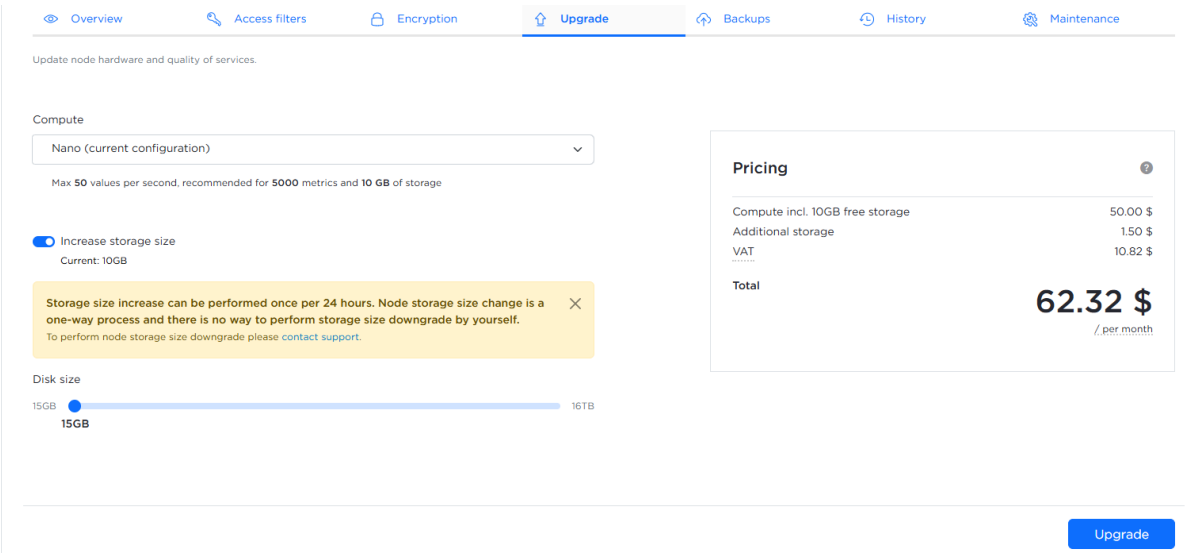
1. Click the *Compute* dropdown.



2. Select desired configuration tier.
3. Click *Upgrade* button.

To increase disk size:

1. Click *Increase storage size* radiobutton.
2. Adjust the *Disk size* slider to the desired value.



3. Click *Upgrade*.

Note:
The updated total monthly price is shown in the *Pricing* table before confirming changes.

Adding users

Introduction

The following examples describe how to add users to your organization and nodes.

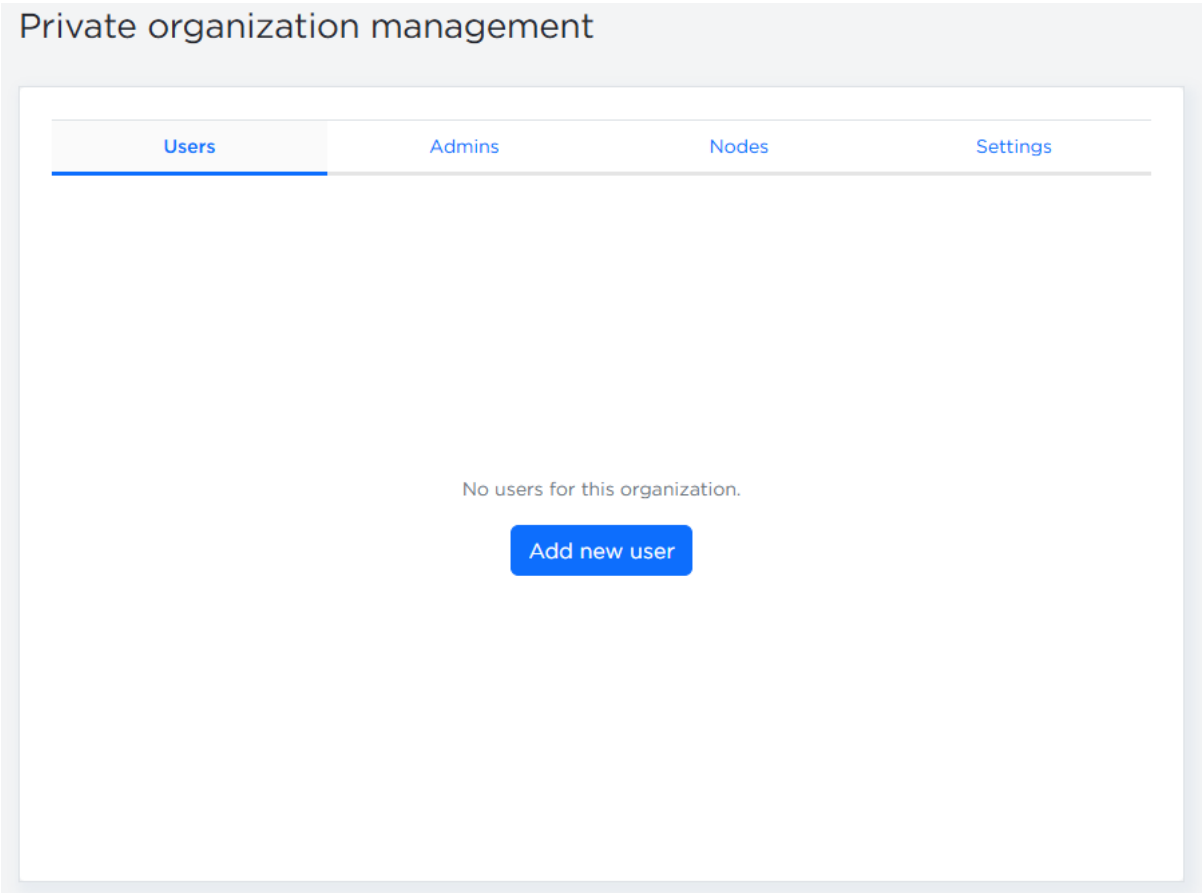
Note:
Copying, clearing, and resetting frontend password can be only performed by the owner.

Adding users to organizations

Similar to users in Zabbix, you can create users in Zabbix Cloud and assign them roles to control what they can do within an organization and its nodes.

To add users to your organization:

1. Open the *All organizations* page.
2. Click the settings button of the particular organization to enter its management menu.



3. In the *Users* or *Admins* tab, click *Add new user*.
4. Enter the user's email and select a role: *Administrator* or *User*.

Note that selecting the *User* role adds the user to the organization, but does not grant access to any nodes.

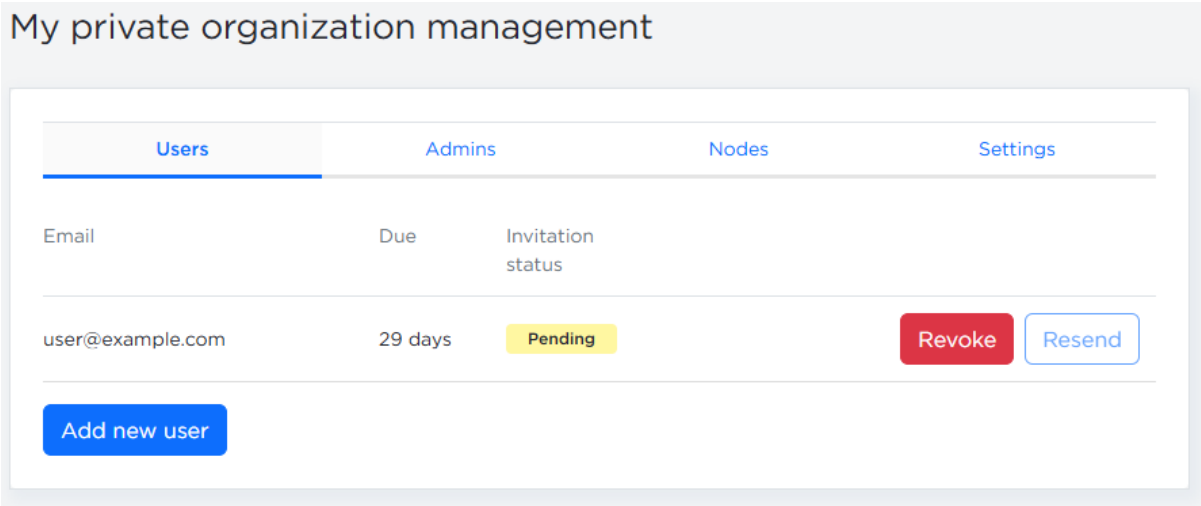
Each role grants a different level of access:

- **User** - can access nodes they're assigned to and can perform configuration tasks that do not generate charges, for example:
  - Manage access filters

  - Configure encryption
  - Set history/trend retention periods
  - Schedule maintenance
- **Administrator** - has all User permissions and can also perform actions that may generate charges, for example:
  - Create and restore backups
  - Upgrade or downgrade node resources (compute tier and disk size)
  - Cancel subscriptions
- **Owner** - has full control over the organization. In addition to all Administrator permissions, the Owner can:
  - Create and destroy nodes
  - Add or remove Users and Administrators
  - Rename or delete the organization
  - Manage billing details

> **Attention:**
> The Owner's payment method is charged for all services. If bank confirmation is required, the Owner must authorize the transaction, even if it was initiated by an Administrator.

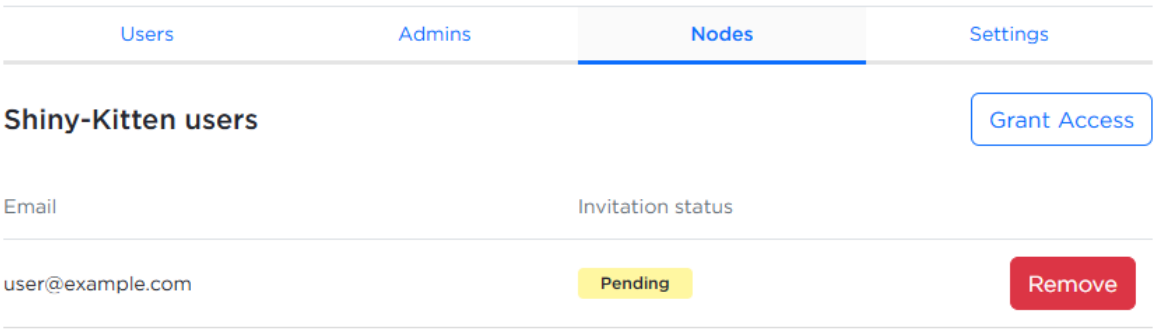
The invited user will be added with a *Pending* status. You can revoke the invite or resend it (resending is available once every 10 minutes—the button will be disabled if this limit is reached).



Adding users to nodes

To add the user to the node:

1. Go to the *Nodes* tab.
2. Click *Grant Access*.
3. Select a user added to the organization from the dropdown.
4. Click *Grant access*.



Alternatively, you can grant access from the *Node configuration* window:

1. Click the  settings button.

2. Select *Add user*.
3. Choose a user and click *Grant access*.

Invited users must accept the invitation by creating a Zabbix Cloud account using the same email address where the invitation was sent. If the user already has a Zabbix Cloud account under that email, they must sign in and confirm the invitation. Invitations are tied to a specific email address and cannot be accepted from a different account.

Key differences of Zabbix Cloud

Introduction

Zabbix Cloud is a cloud-based platform allowing its users to run self-contained Zabbix instances in the cloud (cloud nodes).

Compared to a regular, on-premise Zabbix deployment:

- No installation or maintenance is required.
- Cloud nodes are fully managed by Zabbix.
- Cloud nodes can be started within minutes.

This service is offered on a pay-as-you-go basis (a **free trial** is available).

Comparison table While most of the functionality overlaps between on-premise Zabbix installations and cloud nodes, there are also differences. A summary comparison table is provided below. Subsequent sections describe selected items in detail.

Feature / Aspect	Zabbix Cloud	Zabbix on-premises
Release cycle	In step with LTS releases	You choose your Zabbix version (stable, beta, custom builds)
Ease of use	Ready-to-use platform, fully managed by Zabbix	Requires installation, configuration and maintenance
Vulnerability patching	Automatic OS/Zabbix patching	Manual patching by user
HTTPS certificate	Valid certificate out-of-the-box (no Let's Encrypt)	Self-signed by default (Let's Encrypt or other CA can be configured)
DB performance (INSERT/UPDATE/SELECT)	Auto-partitioning for max speed; no long-running DELETEs	Manual partitioning & cleanup (long DELETE transactions possible)
Performance tuning	Pre-tuned by Zabbix Cloud team (no changelogs)	Full visibility & control over tuning parameters
Firewall control	Managed by Zabbix Cloud team; GUI/trapper ports not exposed publicly	Full control of firewall rules and port exposure
Outbound traffic on port 25	Blocked by default to prevent spam and abuse; use SMTP port 587 or 465 instead	Not restricted by Zabbix; depends on your network and hosting provider
DB isolation & protection	Highly isolated, no direct DB/SSH access	Depends on your network and host security
High availability (HA)	Runs on a highly-available cloud platform, Zabbix proxies can be used for higher data-collection resilience	User-managed HA
SNMP traps	Only via dedicated proxy (no HA/load-balanced trap receiver)	Direct or via proxy, HA/load-balance possible
SNMP polling	Hard without proxies (each device needs NAT/custom port)	Native SNMP polling; proxy optional
Retention settings	History/Trends/Audit only via UI (no API or configuration file)	Configurable in Zabbix server configuration file or via API
Custom scripts	`AlertScriptsPath`, `ExternalScripts`, frontend & community modules not supported	Fully supported (scripts path, modules, integrations)
ODBC monitoring	Only PostgreSQL (driver `{postgresql}`); MariaDB plugin present as a placeholder	ODBC for PostgreSQL, MySQL, Oracle, etc. - fully configurable
Limiting ODBC calls	Not possible (`StartODBCPollers=1`) only; heavy sync queries block other work)	Fully tunable poller count & scheduling
SAML certificate upload	Not supported	Supported via UI or API
Scheduled reports	Must create own Email media (script media unsupported)	Supports both script and Email media out-of-the-box

Feature / Aspect	Zabbix Cloud	Zabbix on-premises
Host interface for active checks	Platform-issued IP address (created automatically)	Interfaces managed by you; IPs under your control

Feature differences SNMP traps

SNMP traps are only supported through a dedicated Zabbix proxy. If SNMP trap monitoring is required, automated load balancing or high availability for the proxy cannot be used because SNMP traps must be sent to a single IP address.

SNMP polling

Without proxies, SNMP polling requires each device to be exposed via NAT and custom ports.

Retention settings

History, trend, and audit log retention periods can be configured only via the web interface in the cloud. These settings cannot be configured through `zabbix_server.conf` or the API.

Manual per-item history overrides are not supported (partitioning is controlled globally). Cloud node URLs do not accept the same query parameters as an on-premise installations.

Custom scripts

The following custom script types are not supported by Zabbix Cloud:

- Alert scripts (`AlertScriptsPath`)
- External scripts (`ExternalScripts`)
- Frontend scripts on Zabbix Server

Community-developed frontend modules cannot be installed.

ODBC monitoring

Zabbix Cloud supports ODBC monitoring only for PostgreSQL. Use the official Zabbix ODBC template and define the connection string in the template using:

```
Driver={postgresql}
```

A MariaDB plugin is also installed but currently not functional. If used, define:

```
Driver={mysql}
```

> **Attention:**
> There is no known working configuration for Oracle MySQL 8.0 in Zabbix Cloud. While simple queries (such as `SELECT 1`) may succeed, more complex queries return `SQL_ERROR`.

 It is not possible to limit ODBC calls (`StartODBCPollers=1` only). Heavy synchronous reporting may impact performance, and only a single `SELECT` can run at a time.

Infrastructure access

Zabbix Cloud does not provide SSH access to the underlying nodes, nor does it permit direct database connections (for example, to port 3306). All configuration, monitoring, and troubleshooting activities must be carried out via the cloud UI or API, ensuring that the OS and database layers remain isolated and secured.

Outbound traffic on port 25

To prevent spam and abuse, outbound connections from Zabbix Cloud nodes to TCP port 25 are blocked by default. This affects:

- Email media types that send through an SMTP server on port 25. Since 25 is the default *SMTP server port* for a new Email media type, change it to 587 (with *Connection security* set to STARTTLS) or 465 (SSL/TLS), depending on what your SMTP server supports. The *Office365 relay* email provider cannot be used, as it connects to Microsoft's relay endpoint on port 25.
- Scheduled reports, which require an Email media type.
- Checks that the cloud node itself performs against port 25, such as `net.tcp.service[smtp]` simple checks or SMTP checks in network discovery rules. These report the SMTP service as down even when it is running; perform such checks through a Zabbix proxy or Zabbix agent in your own network instead.

The default *Cloud Email* media type is not affected.

Active check interface

Active checks in the cloud automatically create a host interface with an IP unrelated to your on-premise network. By default, this IP may be external. To maintain consistency, you can manually delete or adjust the host interface after creation.

SAML certificate upload

SAML authentication is not supported because few providers accept unencrypted or unsigned SAML messages.

Scheduled reports

The default *Cloud Email* media type uses a script transport and cannot be used for scheduled reports. To send reports by email:

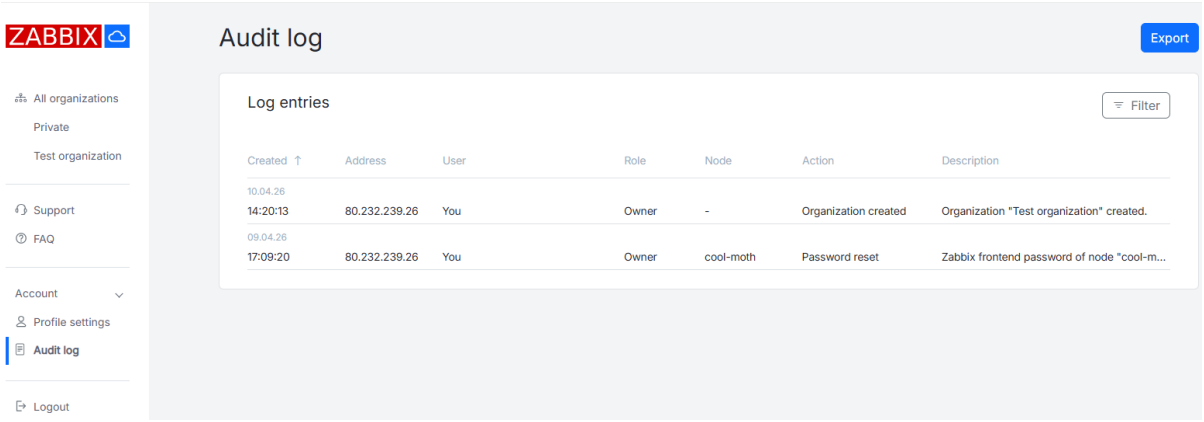
1. Create a new Email media type.
2. Assign the email media to your user in *User settings* > *Media*.
3. Select that media type when configuring scheduled reports.

Attention:
Do not use port 25 as the *SMTP server port* for this media type; see Outbound traffic on port 25.

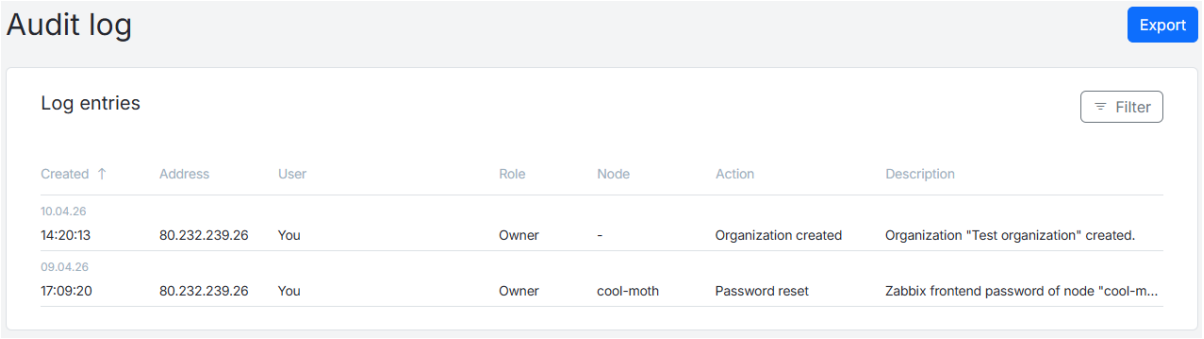
Audit log

Overview In Zabbix Cloud, the Audit log records user and system activity.

After logging in to a Zabbix Cloud account, go to *Account* > *Audit log* to view audit log records. The audit log entries available to a user depend on the user's organization role.



Viewing audit log records Audit log records are displayed in a table.



The table contains the following information:

Column	Description
Created	Date and time when the log entry was created. The displayed time corresponds to the user's timezone selected in the *User timezone* dropdown in *Account* > *Profile settings*.
User	Name of the organization user who performed the action.
Role	The role of the user who performed the action.
Node	Name of the node related to the action.
Action	Name of the performed action.
Description	Description of the performed action.

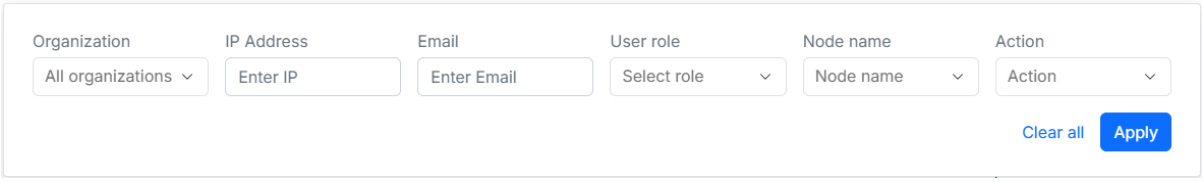
Column names can be clicked to sort records in ascending or descending order.

Filtered audit log records can be exported in `.csv` format using the *Export* button.

Using filters

Use the *Filter* button to open the filter.

To display the filtered results, click *Apply* after selecting one or more filtering parameters. To clear the filter, click *Clear all*.



It is possible to filter records by one or more parameters. The following filtering parameters are available:

Parameter	Description
Organization	Select all organizations or a specific one.
IP Address	Specify an IP address.
Email	Specify an email address attached to the user.
User role	Select all roles or a specific one: Owner, Admin, or User.
Node name	Select all nodes or a specific one.
Action	Select a specific action type.

Audit log tracks the following actions:

- *Node created*
- *Node deleted*
- *Node upgraded*
- *Node version updated*
- *Password reset*
- *Maintenance updated*
- *History setting updated*
- *Node expiration date extended*
- *Proxy configuration updated*
- *Backup created*
- *Backup deleted*
- *Backup restored*
- *Backup updated*
- *Access filter added*
- *Access filter updated*
- *Access filter deleted*
- *Encryption enabled*
- *Encryption disabled*
- *Subscription cancelled*
- *Subscription renewed*
- *Subscription payment failed*
- *Payment requires confirmation*
- *Organization created*
- *Organization deleted*
- *Admin invited*
- *User invited*
- *Admin joined*
- *User joined*
- *Admin removed from organization*
- *User removed from organization*
- *Node added to organization*
- *Node removed from organization*
- *Node access granted*
- *Node access revoked*
- *User settings updated*
- *Payment created*

- *Payment completed*
- *Payment failed*
- *VAT charge applied*
- *VAT validation failed*
- *Billing info updated*
- *Payment method updated*
- *Payment method removed*
- *Prepaid node created*
- *Prepaid node updated*
- *Incident started*
- *Incident ended*

Set up a Zabbix Cloud proxy

Overview This page describes how to configure a Zabbix proxy deployed on Zabbix Cloud to collect monitoring data from hosts in a remote environment.

The procedure is primarily intended for a Zabbix server and proxy created on Zabbix Cloud.

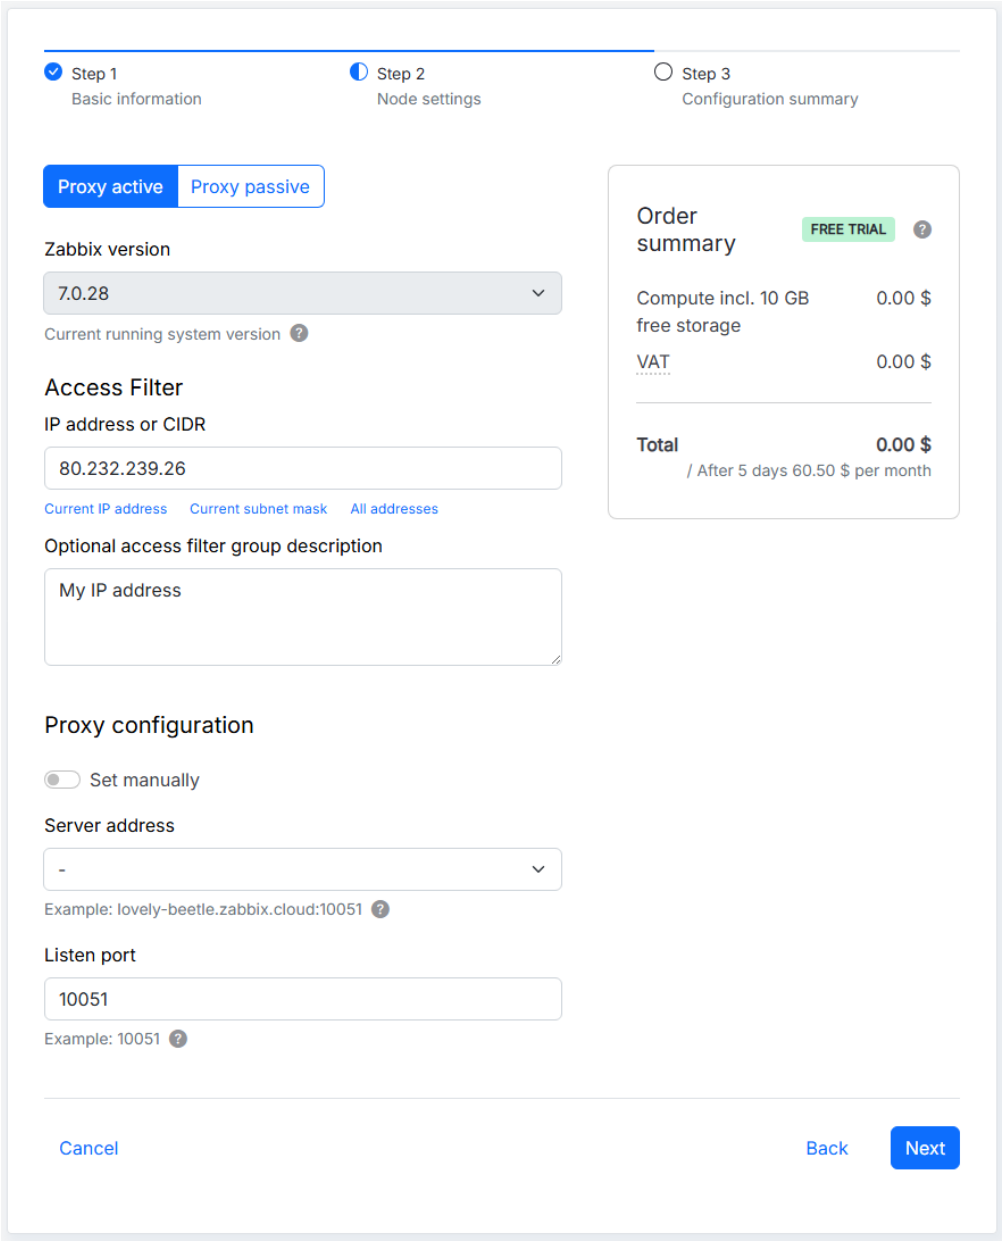
The procedure also describes the configuration required when the Zabbix server is not hosted on Zabbix Cloud.

Prerequisites A Zabbix server is installed, configured, and running, or a Zabbix server node is running on Zabbix Cloud.

Create a proxy node Create a Zabbix Cloud proxy node.

During node creation:

1. Select *Proxy*.
2. Select *Proxy active* or *Proxy passive* depending on how the proxy should communicate with the Zabbix server.
3. Configure the node's IP address or CIDR access filters (can be done after node creation).
4. Specify the Zabbix server to which the proxy will connect.



5. Complete the node creation process.

For more information about creating a proxy node, see Deploy Zabbix in the cloud.

Configure access filters Zabbix Cloud uses access filters to control which IP addresses or subnets can connect to a cloud node.

During node initialization, an access filter is automatically created for the IP address used to access the node. This does not necessarily include the addresses of monitored hosts or other Zabbix components. Add the required addresses manually.

See Node configuration — Access filters for more information.

Proxy node

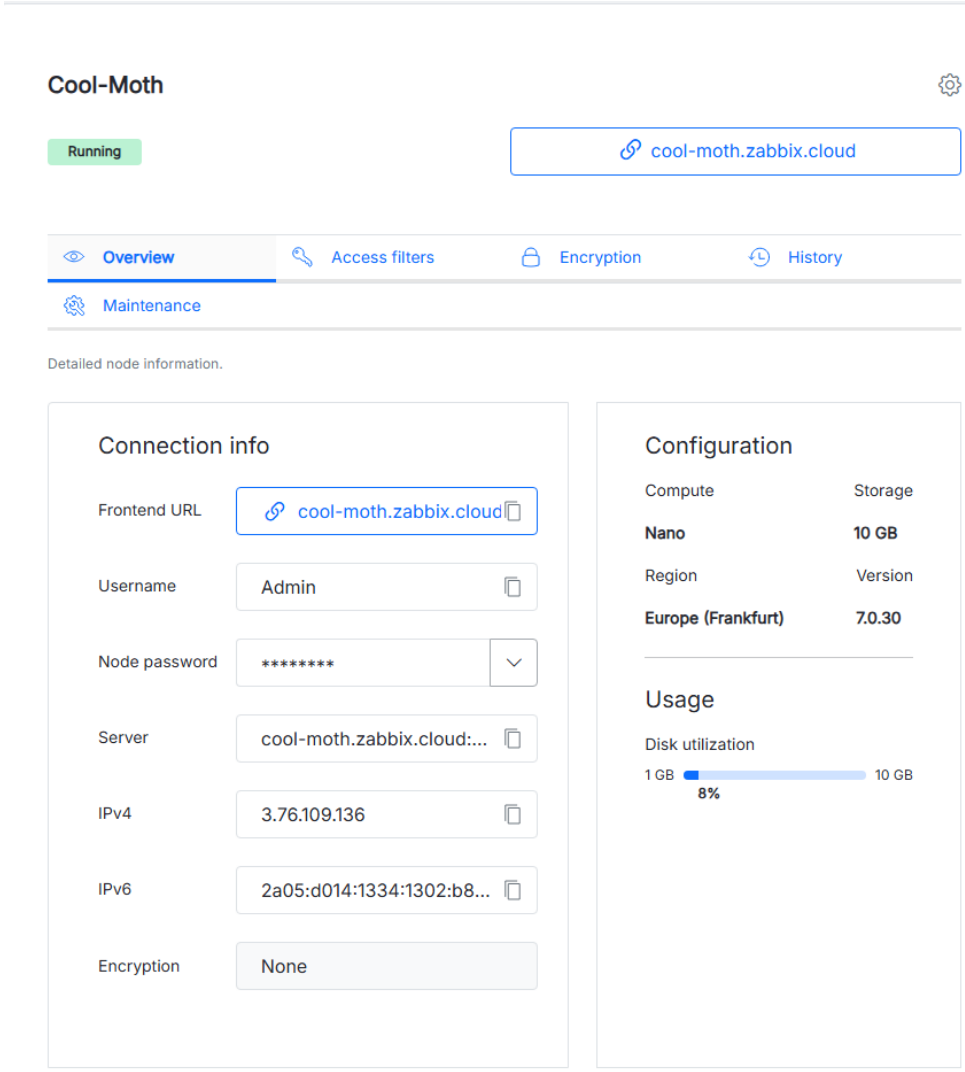
Add the IP addresses or subnets from which connections to the proxy are expected.

If monitored hosts use active agent checks, the agents initiate connections to the proxy, so the IP address or subnet of the monitored hosts must be allowed by the proxy node's access filters.

If *Proxy Passive* was set as a proxy mode, the Zabbix server connects to the proxy, so the IP address or subnet of the Zabbix server must be allowed by the proxy node's access filters.

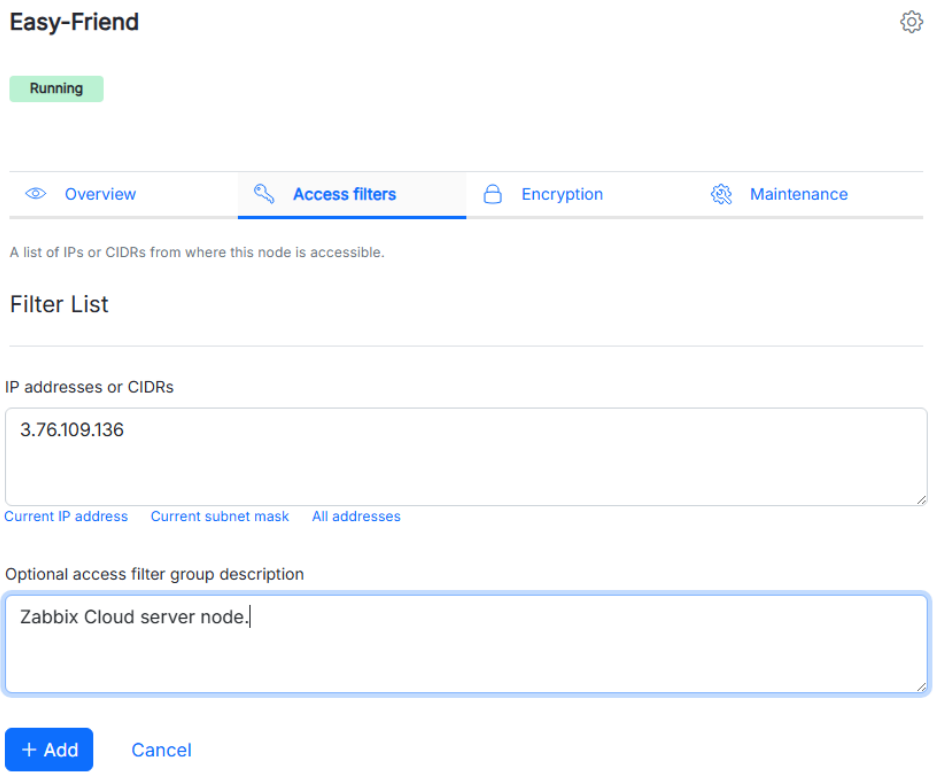
If *Proxy passive* was selected and a server node is used:

1. Navigate to the *Overview* tab of your server node, and copy the IPv4 or IPv6 address of the node.

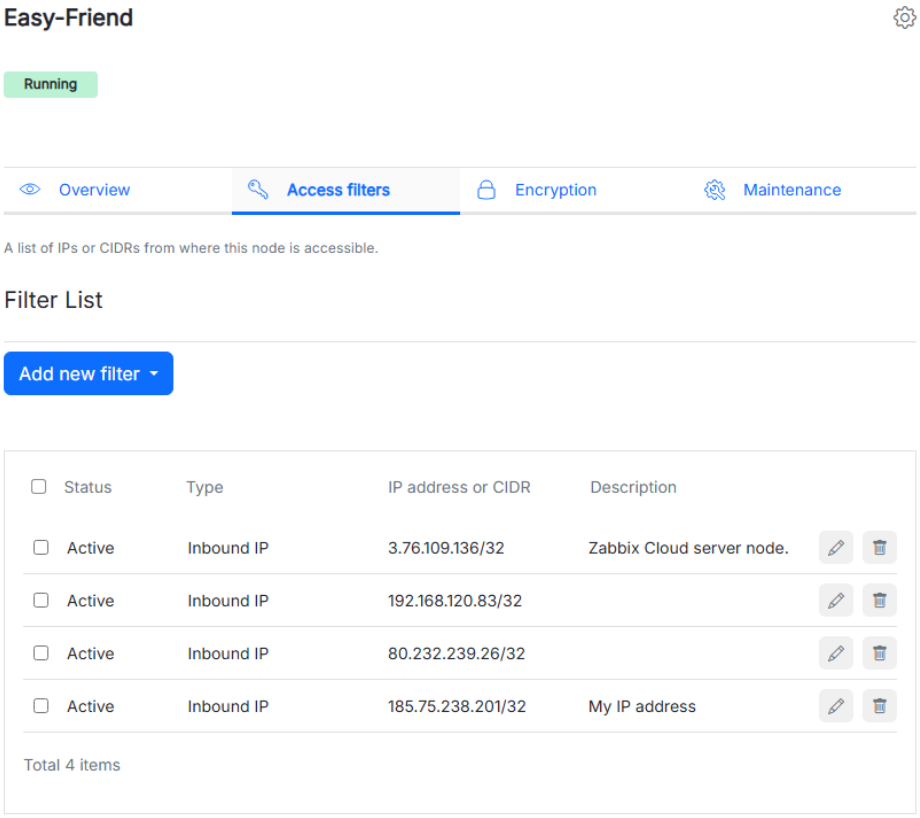


2. Navigate to the *Access filters* tab of your proxy node, click *Add new filter* > *Add*.

3. Paste the copied address to the *IP addresses or CIDRs* field, and add description to the *Optional access filter group description* field.



4. Click *Add*.



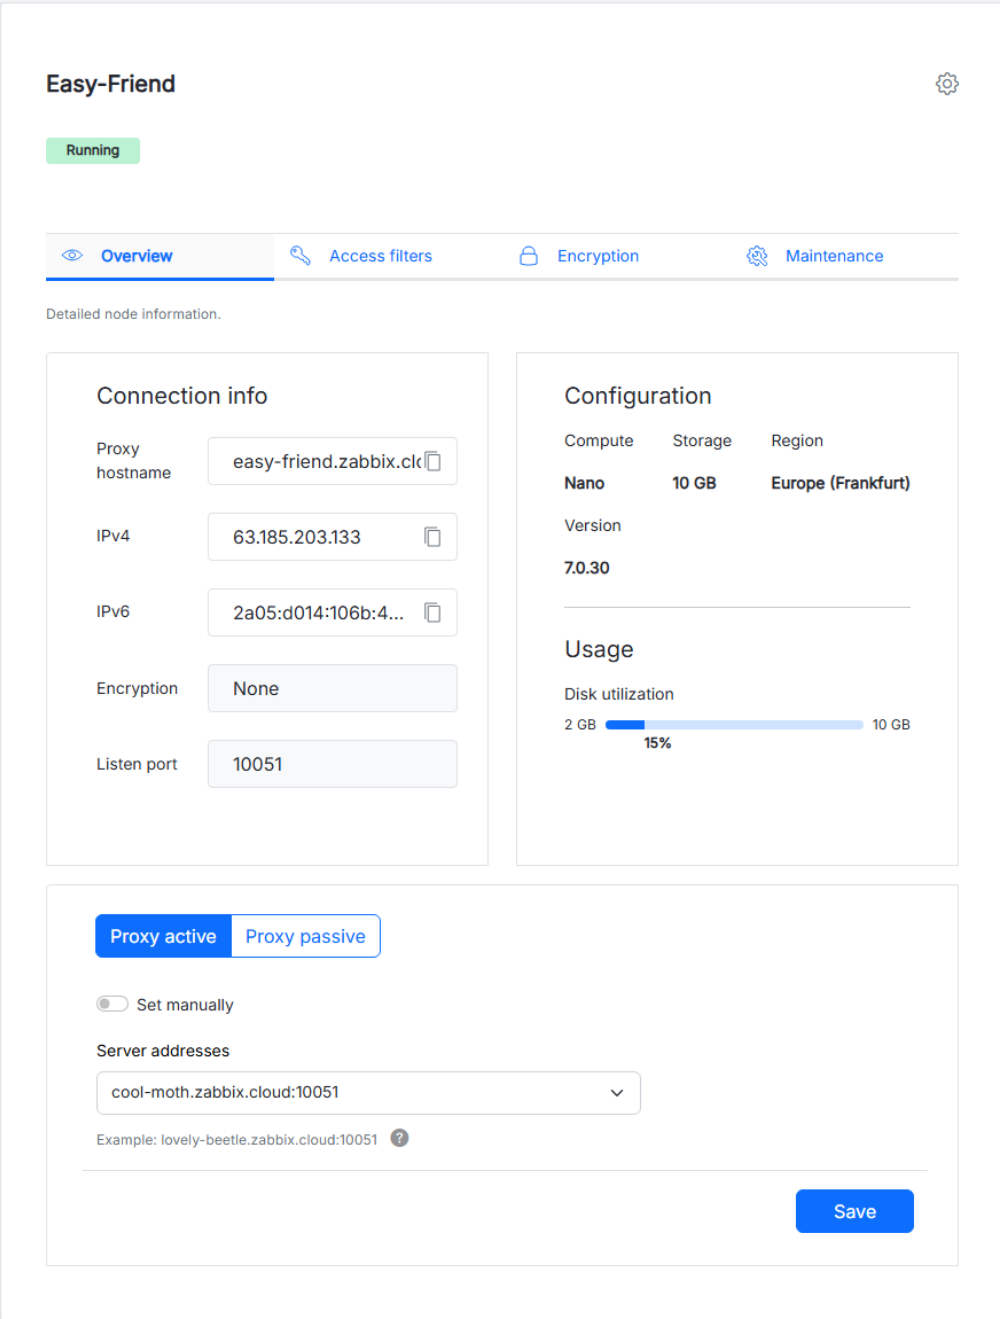
Do steps 2-4 for the IP address of a Zabbix agent using active checks.

Server node

If the Zabbix server is hosted on Zabbix Cloud, configure access filters on the server node as well.

For a proxy in the active mode, add the IP address or CIDR of the proxy node to the server node's access filters, because the proxy initiates the connection to the server.

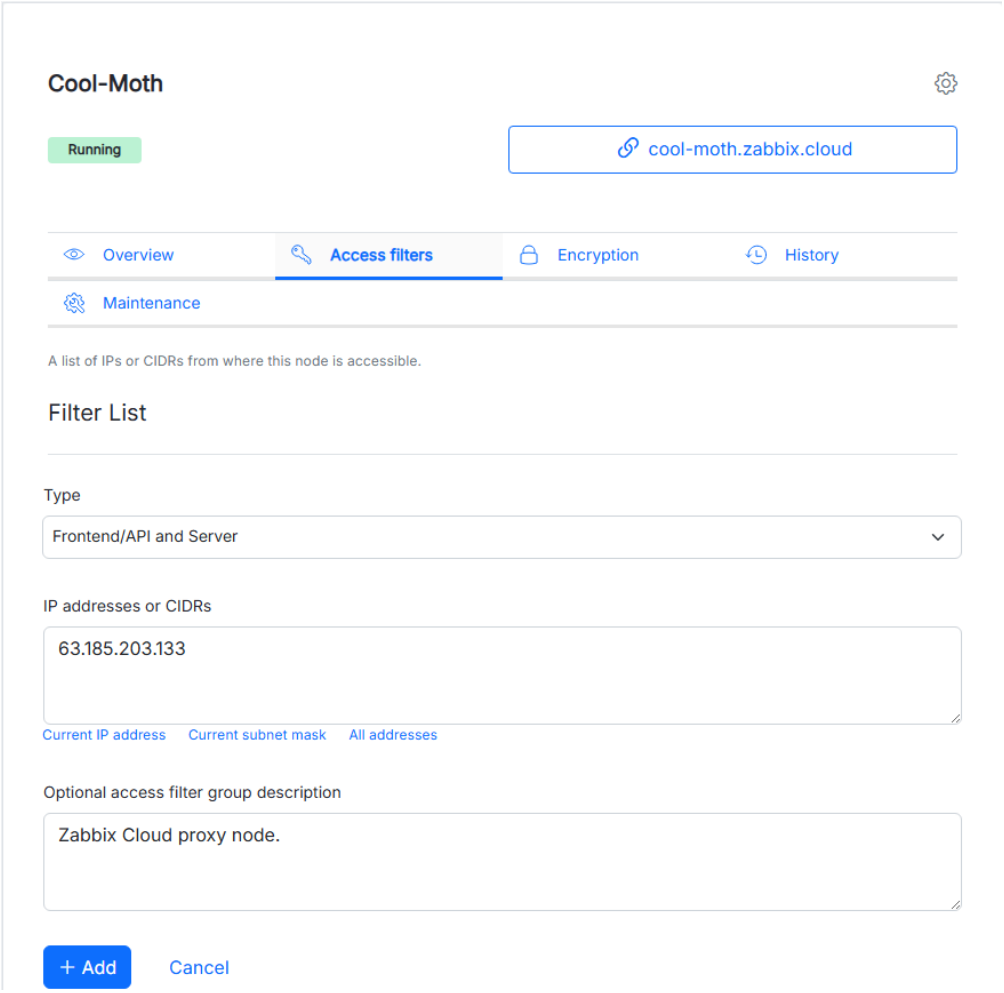
1. Navigate to the *Overview* tab of your proxy node, and copy the IPv4 or IPv6 address of the node.



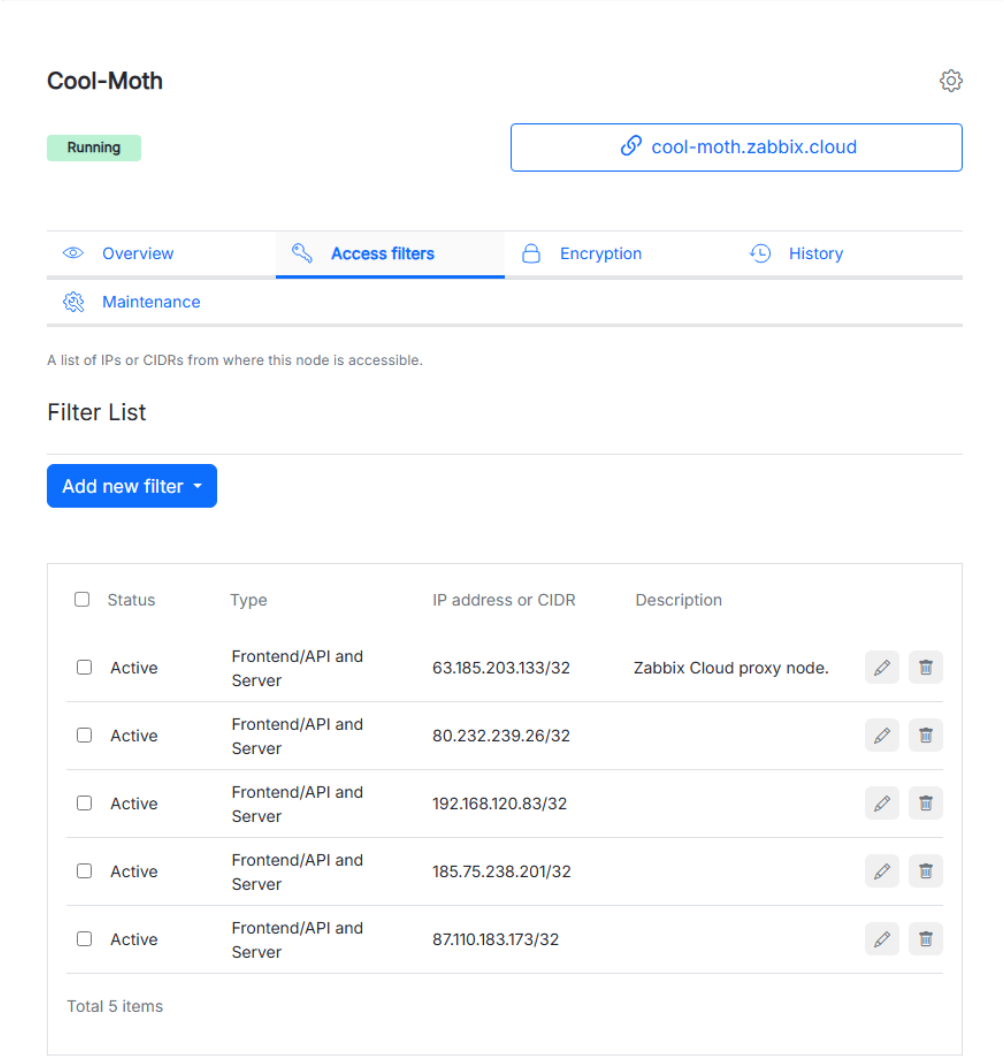
2. Navigate to the *Access filters* tab of your server node, click *Add new filter* > *Add*.

3. Paste the copied address to the *IP addresses or CIDRs* field, and add description to the *Optional access filter group description* field.

4. Select the appropriate filter type.


Cool-Moth
Running
cool-moth.zabbix.cloud
Overview
Access filters
Encryption
History
Maintenance
A list of IPs or CIDRs from where this node is accessible.
Filter List
Type
Frontend/API and Server
IP addresses or CIDRs
63.185.203.133
Current IP address
Current subnet mask
All addresses
Optional access filter group description
Zabbix Cloud proxy node.
+ Add
Cancel

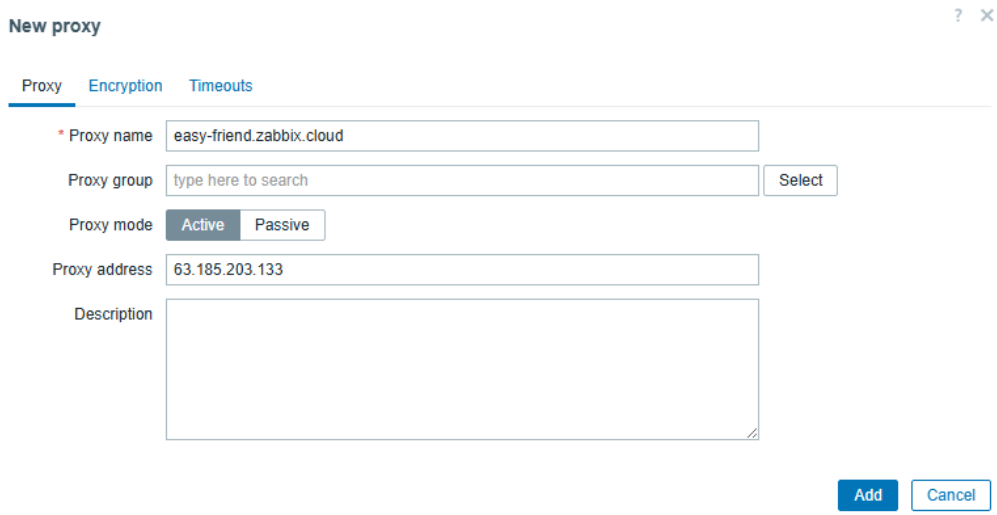
5. Click *Add*.



Configure the proxy The proxy must be configured on the Zabbix frontend to communicate with the Zabbix server.

Open *Administration* > *Proxies* in the Zabbix frontend, click *Create proxy*, and configure the proxy as required:

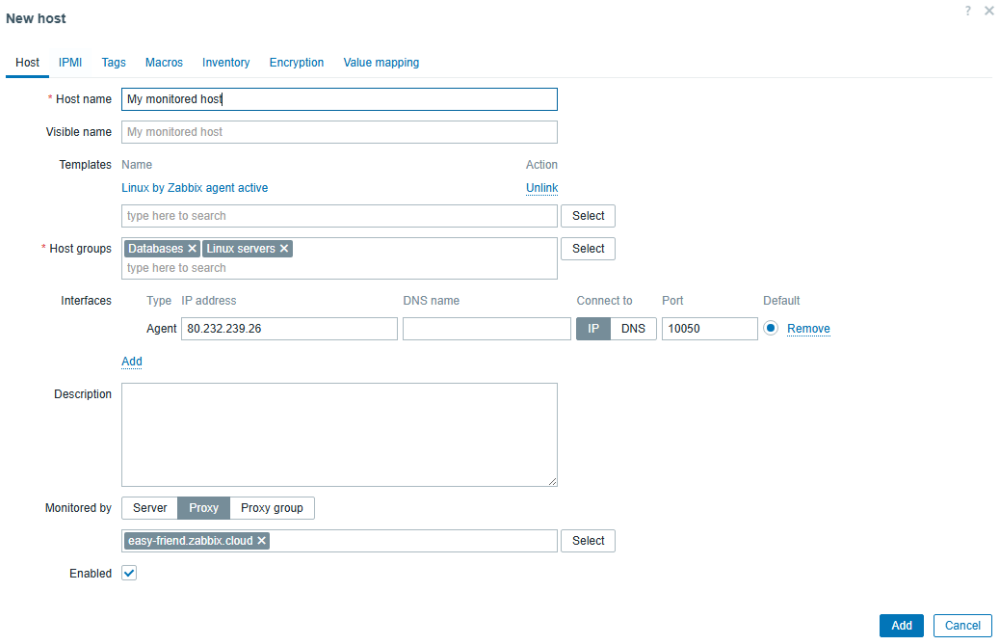
1. Add the proxy name to the *Proxy name* field. It must match the proxy hostname configured for the proxy. Type it in, or copy it from the *Overview* tab of your proxy.
2. Select the Proxy mode corresponding to the one selected during proxy creation.
3. For an active proxy, enter the IP address or DNS name of the proxy.
4. Click *Add*.



For more information, see Distributed monitoring using proxies.

Configure monitored hosts Add the hosts that should be monitored through the proxy.

1. Go to *Data collection* > *Hosts*.
2. Create a host or open an existing host.
3. In *Monitored by*, select the *Proxy* and either type in the proxy name to search and select, or press the *Select* button to open a list of proxies to select from.
4. Configure the host interface and other parameters.
5. Click *Add* or *Update*.



A host configured to be monitored by a proxy sends or receives monitoring data through that proxy, depending on the item type and agent mode.

Configure Zabbix agent When a monitored host uses Zabbix agent or Zabbix agent 2, configure the agent to communicate with the appropriate Zabbix component.

The configuration depends on whether the agent uses *passive* or *active* checks.

Passive agent checks

With passive checks, the Zabbix proxy initiates connections to the agent and requests monitoring data.

Open the Zabbix agent configuration file and update the `Server` parameter to allow connections from the IP address of the Zabbix Cloud proxy.

For example:

```
Server=<proxy_IP>
```

Make sure that the monitored host allows incoming TCP connections from the proxy on the agent port, which is 10050 by default.

The proxy must be able to reach the monitored host over the network.

Active agent checks

With active checks, the agent initiates connections to the Zabbix proxy and sends collected monitoring data.

Open the Zabbix agent configuration file and update the `ServerActive` parameter with the address of the Zabbix Cloud proxy.

For example:

`ServerActive=<proxy_IP>:<proxy_listen_port>`

Replace `<proxy_IP>` with the IP address of the Zabbix Cloud proxy and `<proxy_listen_port>` with the listen port configured when the proxy node was created. The default listen port is 10051.

If a proxy is specified in `ServerActive`, the agent obtains its active-check configuration from the proxy rather than directly from the Zabbix server.

Make sure that the monitored host can establish an outgoing TCP connection to the proxy on the configured listen port.

For more information about passive and active agent checks, see Passive and active agent checks.

External Zabbix server A Zabbix Cloud proxy can also communicate with a Zabbix server deployed outside Zabbix Cloud.

In this case, the external Zabbix server must be able to communicate with the cloud proxy according to the selected proxy mode.

For a proxy in the active mode, the proxy initiates the connection to the external Zabbix server. Configure the external Zabbix server to accept connections from the proxy and make sure that the proxy can reach the server over the network.

For a passive proxy, the Zabbix server initiates the connection to the proxy. Add the IP address or subnet of the external Zabbix server to the proxy node's access filters and make sure that the server can reach the proxy over the network.

The external Zabbix server does not use Zabbix Cloud access filters. Instead, configure the required firewall and network rules on the external environment.

The monitored hosts can be located in the same environment as the external Zabbix server, in another network, or in a network accessible by the proxy.

When agents use active checks, configure their `ServerActive` parameter to point to the Zabbix Cloud proxy. When agents use passive checks, configure their `Server` parameter to allow connections from the proxy.

Verify the configuration After configuring the proxy, agents, access filters, and network connectivity, verify that the proxy is communicating with the Zabbix server.

1. Go to *Administration* > *Proxies*.
2. Check the proxy *State*.
3. Go to *Monitoring* > *Latest data* and check that recent item values are available.

If the proxy is not communicating with the server, verify:

- the proxy mode;
- the server address and port;
- proxy access filters;
- server access filters, when the server is hosted on Zabbix Cloud;
- agent configuration;
- firewall and routing rules;
- the address configured in `Server` or `ServerActive`.

Mobile app

Mobile app

Zabbix mobile app lets you monitor your Zabbix environment from your mobile device. View problems, check their details, update problems, and get push notifications when something needs attention.

Get the app (coming soon)

Set up Zabbix for the app

Set up Zabbix bridge adapter Install Zabbix bridge adapter to let mobile devices connect to Zabbix.

Add devices Add your mobile device in the Zabbix web interface to connect it to your Zabbix server.

Set up push notifications Get alerts on your mobile device when a problem occurs.

Learn about the app

App overview Learn the basics of using Zabbix mobile app.

Secure mobile access Secure the app, your mobile devices, and Zabbix bridge adapter, or revoke access from lost or compromised devices.

Set up Zabbix for the app

To use Zabbix mobile app (coming soon), you must first:

1. Set up Zabbix bridge adapter.
2. Set up push notifications.
3. Add your mobile device in the Zabbix web interface.

The sections below cover each step.

After setting up, see Secure mobile access to keep the app and Zabbix bridge adapter secure.

Set up Zabbix bridge adapter To allow your Zabbix server to communicate with the mobile app, you must set up Zabbix bridge adapter.

The bridge adapter can be installed on the same system as your Zabbix server, or on a different system, using packages:

> **Attention:**
> The steps below are for Ubuntu 24.04. For other operating systems, take the repository and package installation commands (e.g., `wget`, `dpkg`, `apt`) from the Get Zabbix page for your OS, and use them to install the `zabbix-bridge-adapter` package.

1. Start a new shell session with root privileges:

```
sudo -s
```

2. Install the Zabbix repository:

```
wget https://repo.zabbix.com/zabbix/8.2/release/ubuntu/pool/main/z/zabbix-release/zabbix-release_latest_8.
dpkg -i zabbix-release_latest_8.2+ubuntu24.04_all.deb
apt update
```

3. Install Zabbix bridge adapter:

```
apt install zabbix-bridge-adapter
```

4. Edit the Zabbix server configuration file:

```
EnableMobileDevices=1
BridgeAdapterURL=https://bridge-adapter.example.com:10054/rpc
```

5. Restart Zabbix server, for example (see other options):

```
systemctl restart zabbix-server
```

6. Copy the Zabbix server ID from *Reports* > *System information*; you will use it in the next step.

7. Edit the Zabbix bridge adapter configuration file:

```
ZabbixAPIURL=https://example.com/zabbix/api_jsonrpc.php
ServerId=019fb1f9-22bd-74ea-bc03-51efb8b76aba
DataDir=/var/lib/zabbix-bridge-adapter
```

> **Attention:**
> The directory specified in `DataDir` must have 0700 permissions, and all files stored in it must have 0600 permissions.

8. Review and accept the end-user license agreement (EULA) of Zabbix bridge adapter:

```
zabbix_bridge_adapter license show
sudo -u zabbix-bridge-adapter zabbix_bridge_adapter -c /etc/zabbix/zabbix_bridge_adapter.conf license acce
```

9. Start Zabbix bridge adapter:

```
systemctl start zabbix-bridge-adapter
```

Attention:
When started, Zabbix bridge adapter uses the server ID to enroll your Zabbix server with the centrally hosted Zabbix Bridge. The enrollment state is stored in a state file in `DataDir`. If you reinstall the bridge adapter and the state file is lost, enrolling the same Zabbix server may fail. To avoid this, back up the state file before you reinstall the bridge adapter and restore it afterwards.

Set up push notifications To get alerts on your mobile device when something needs attention in your monitored environment:

1. Create a push notification media type (or enable the default one).
2. Create an action (or enable the default *Report problems to Zabbix administrators* trigger action).
3. Add a device to automatically add the media type to your user (or add user media manually).

Add devices To use Zabbix mobile app on your device, you must add that device in the Zabbix web interface. For instructions, see:

- *Users* > *Devices* (as *Super admin*, add your own or other user mobile devices)
- *User settings* > *Devices* (as any user, add your own mobile devices)

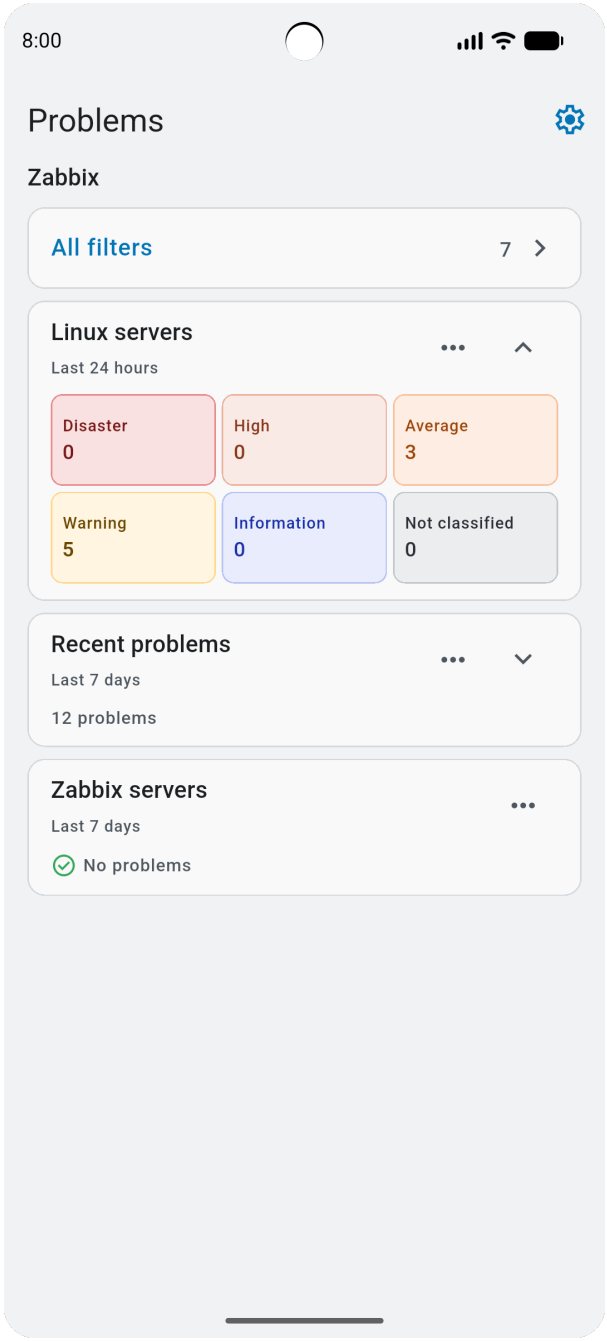
App overview

Zabbix mobile app (coming soon) lets you monitor your Zabbix environment from your mobile device. View problems, check their details, update problems, and get push notifications when something needs attention.

This page gives you an overview of the mobile app and how to use it.

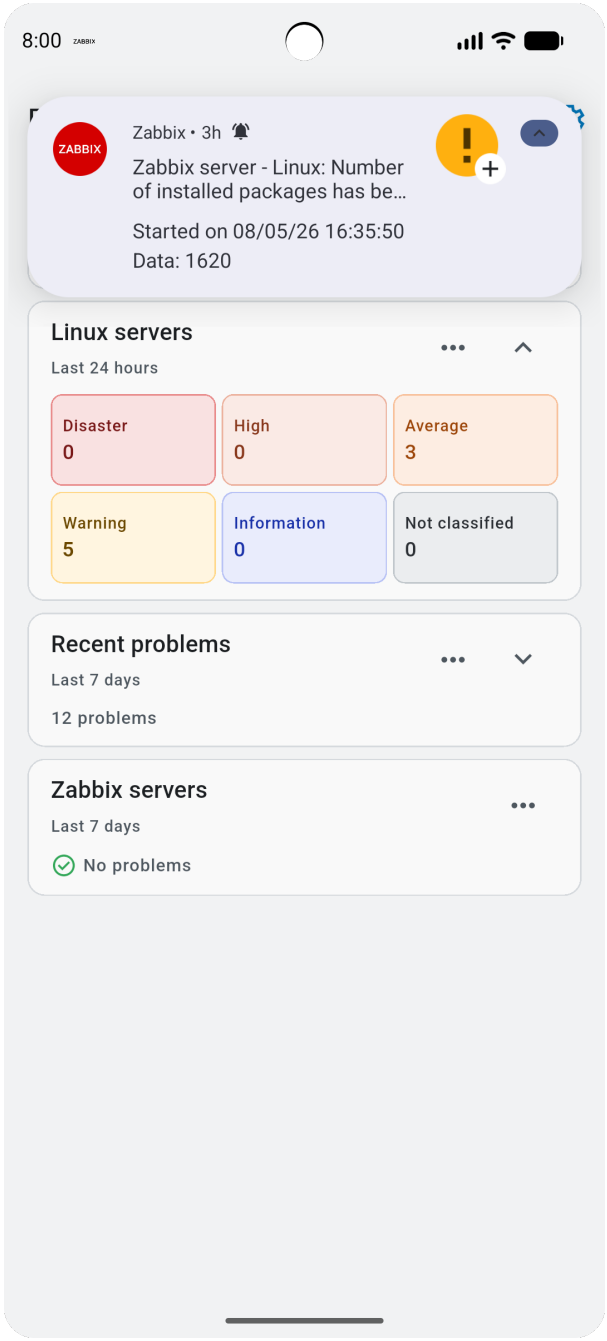
Problems The *Problems* screen shows you an overview of problems based on your pinned filters. It contains:

- *Settings* icon - tap to see app settings.
- *All filters* - tap to see and manage filters.
- Zabbix server name (or *Zabbix*, if not configured when installing the Zabbix web interface).
- Pinned filters - tap their name or a severity to open the *Problem list* screen. Tap the three-dot menu to manage (edit, duplicate, unpin, or delete) the filter, or tap the arrow to expand or collapse it. Note that pinned filters are ordered alphabetically, and the *Recent problems* filter cannot be edited and removed.



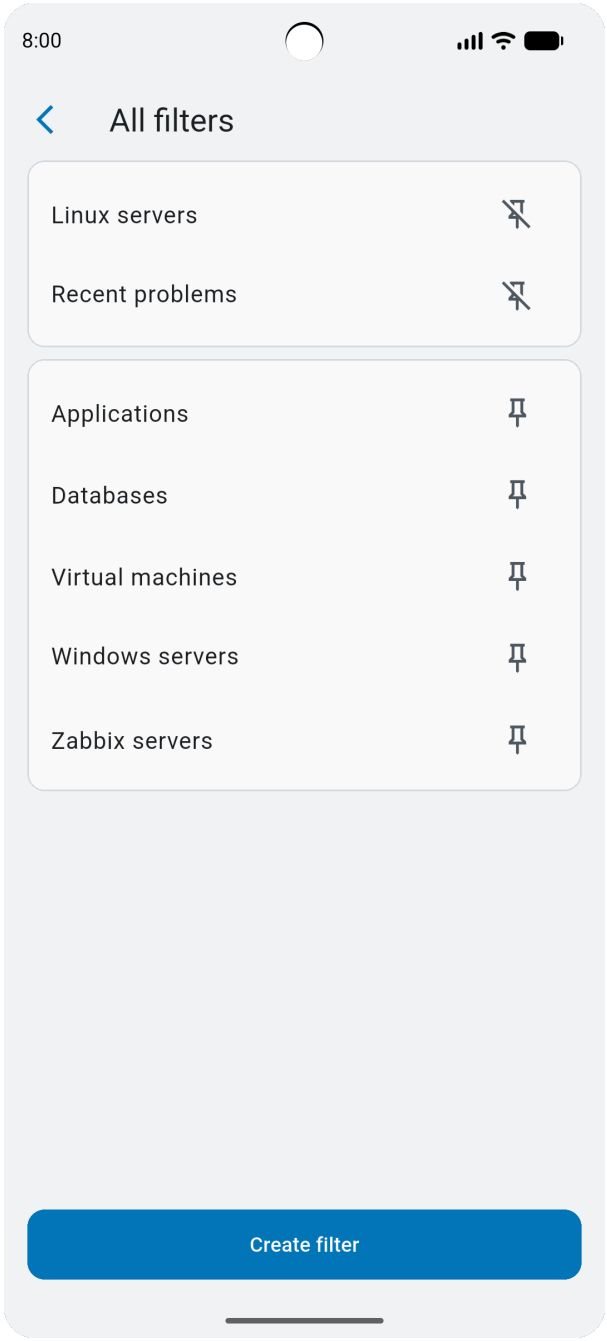
Notifications When Zabbix server detects a problem, or when a problem is updated or resolved, it can send a notification to your mobile device. The notification contains the problem host, name, and a message generated from the message templates configured for the push notification media type.

Tap the notification to open the *Problem details* screen.



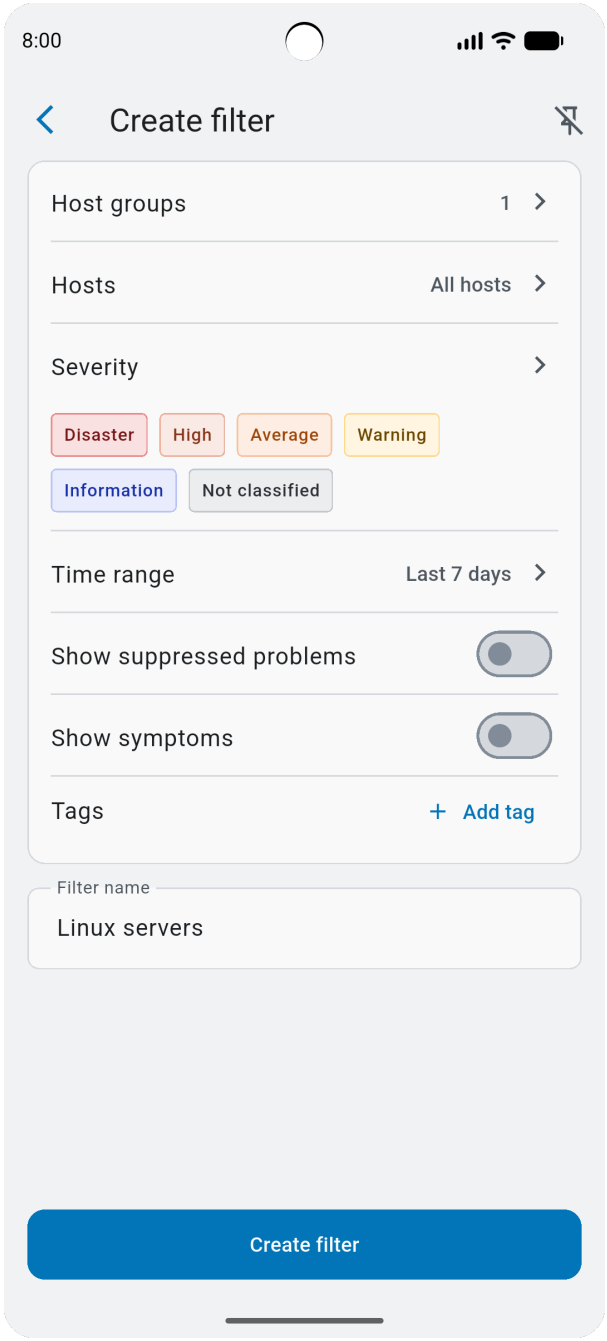
All filters The *All filters* screen lets you manage filters. It contains:

- Filter names - tap to edit the filter.
- Pin icons - tap to pin or unpin the filter to the *Problems* screen.
- *Create filter* - tap to create a new filter.



Create filter The *Create filter* screen lets you create a filter, with options similar to those in *Monitoring* > *Problems* (see details). You can also set a *Time range* for how long to show unresolved problems.

Tap *Create filter* to save the filter. New filters are pinned to the *Problems* screen; tap the pin icon to change this.



Problem list The *Problem list* screen shows problems that match the filter (e.g., *Recent problems*).

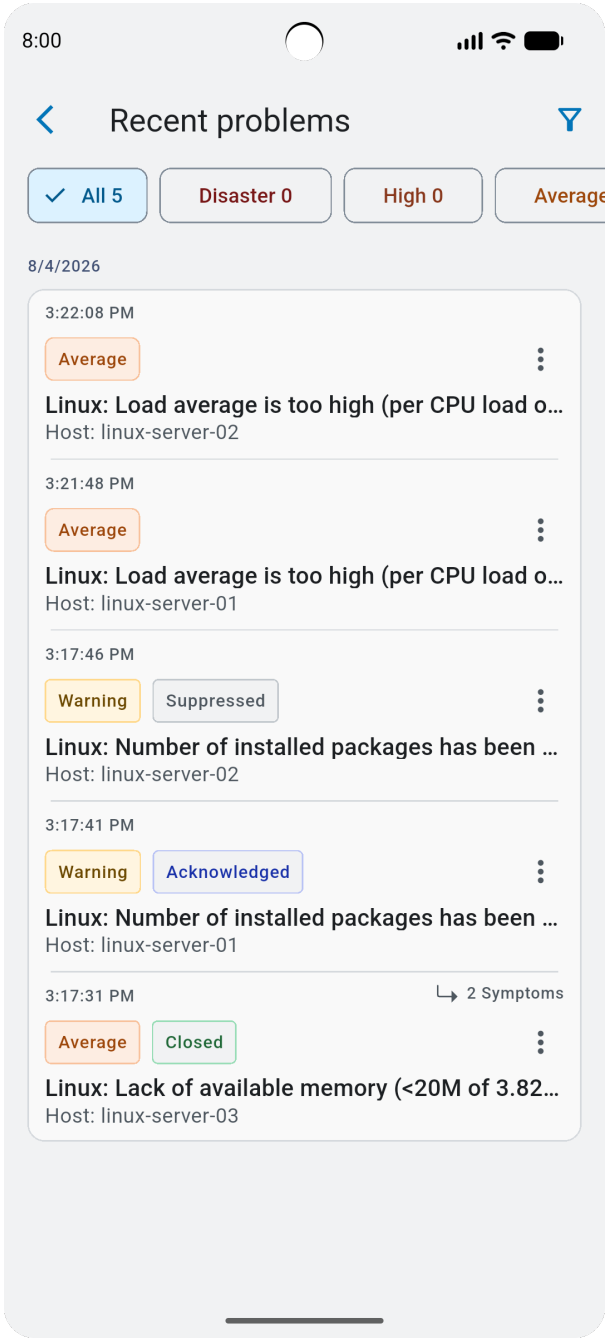
Tap the filter icon or a problem severity to narrow the list further.

For each problem, you can see:

- Start time and date.
- Severity, name, and host.
- Whether it was suppressed (displayed if *Show suppressed problems* is enabled).
- Whether it was acknowledged or recently closed (resolved automatically or closed manually).
- Whether it has symptoms - tap to see the symptom problem list.

Tap a problem to open the *Problem details* screen.

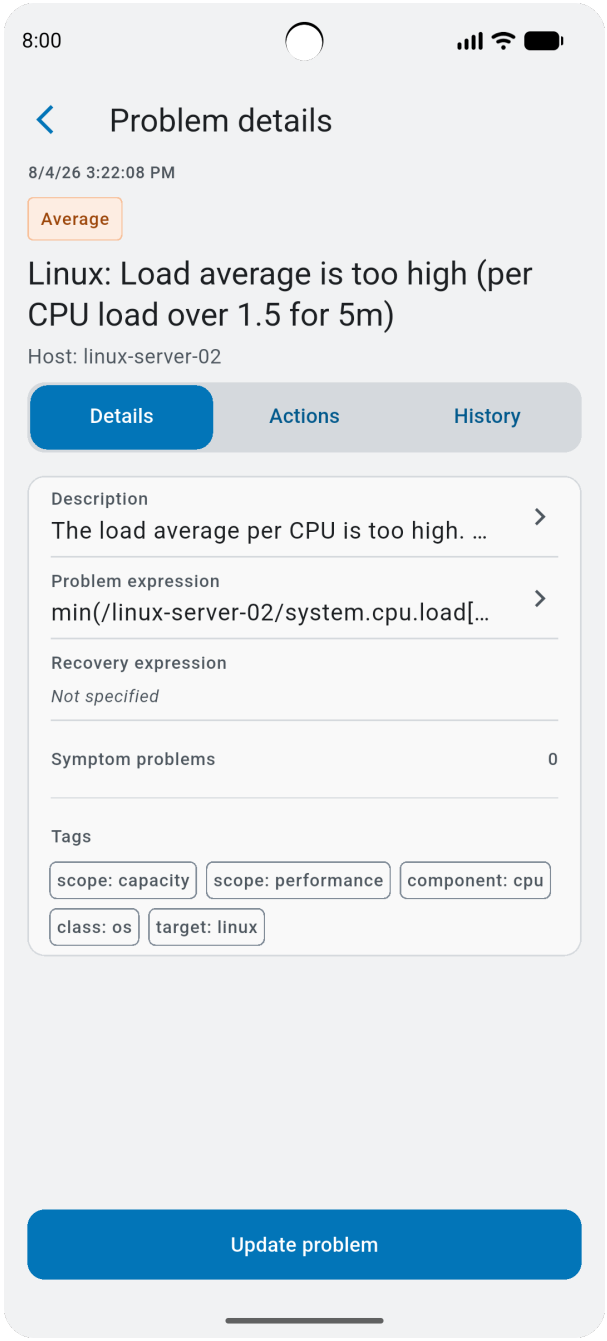
Tap the three-dot menu to quickly acknowledge the problem, or suppress it for a time you choose.



Problem details The *Problem details* screen shows you problem details, actions, and history.

- *Details* tab - shows the problem description, expressions, symptom problems (tap to see the problem list), and tags.
- *Actions* tab - shows the problem timeline: when it started or was acknowledged, whether a notification was sent, and other actions.
- *History* tab - shows the problem age, duration, and recovery time (or *Not available* if the problem is still ongoing).

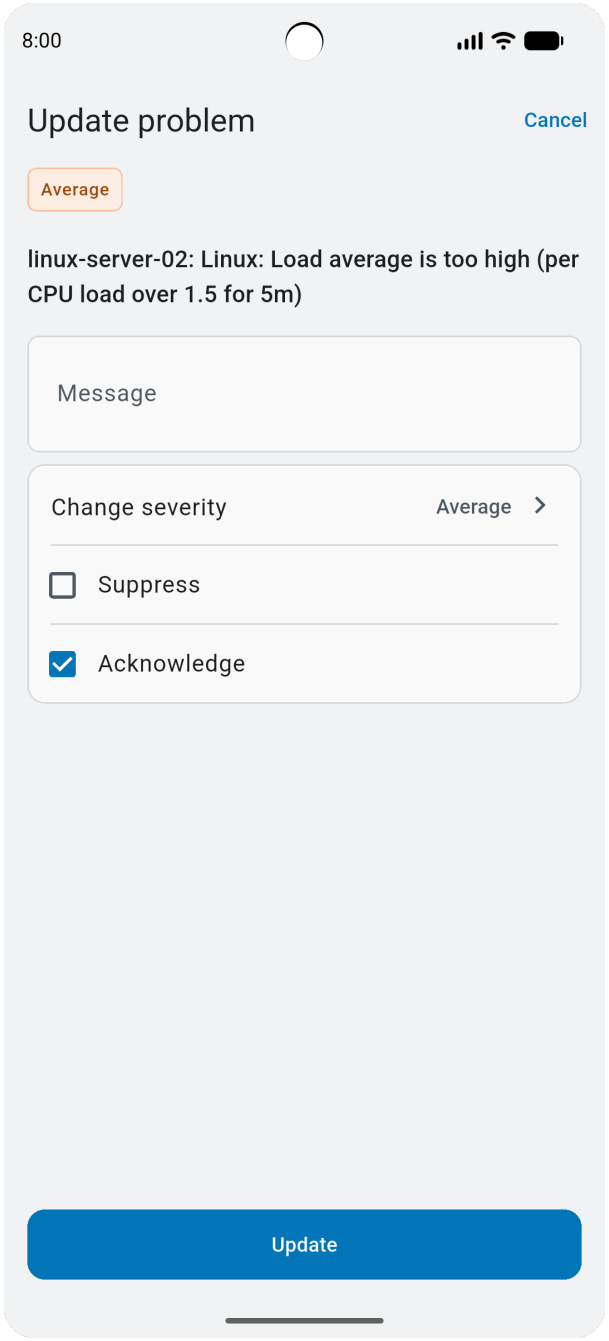
Tap *Update problem* to update the problem.



Update problem The *Update problem* screen lets you update a problem:

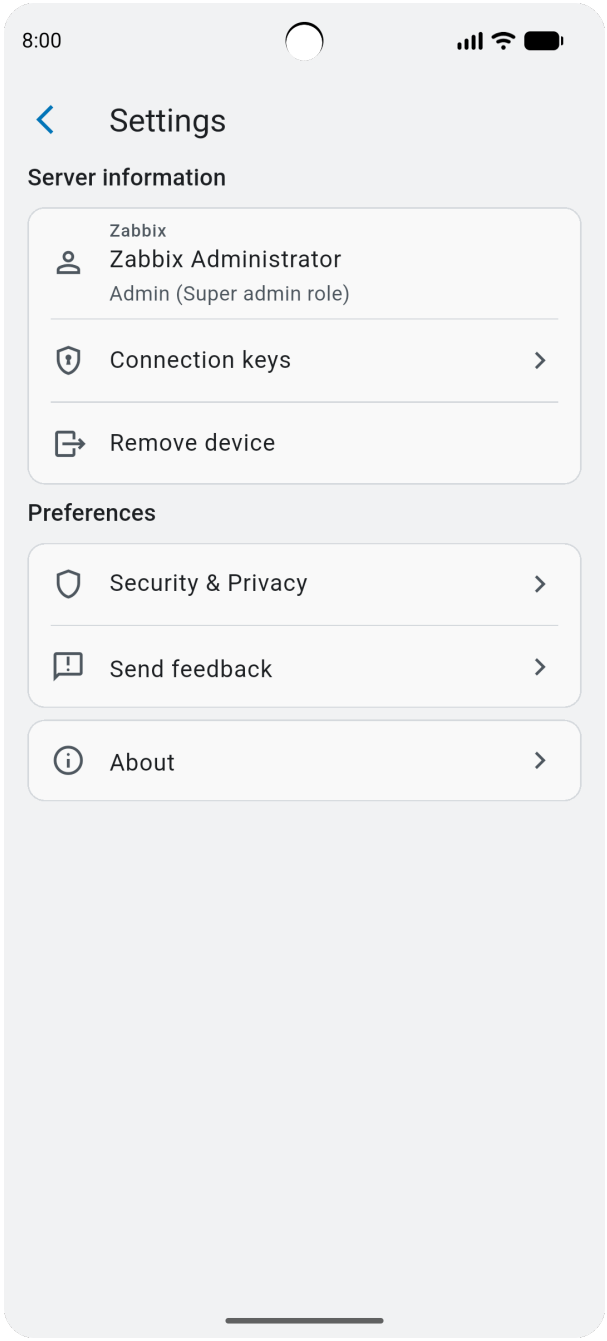
- Add a message.
- Change the problem severity.
- Suppress the problem for a time you choose.
- Acknowledge or unacknowledge the problem.
- Close the problem.

Tap *Update* to save your changes.



Settings The *Settings* screen lets you access app settings. It contains:

- Zabbix server name (or *Zabbix*, if not configured when installing the Zabbix web interface).
- Name, type, and role of the user the mobile device is added to.
- *Connection keys* - open to see and copy the server ID, device ID, and other details.
- *Remove device* - tap to remove the mobile device from Zabbix.
- *Security & Privacy* - open to enable app lock and manage diagnostic data sharing.
- *Send feedback* - open to send feedback to Zabbix.
- *About* - open to see the app version and other details about the app.



Secure mobile access

This page contains best practices for keeping Zabbix mobile app and Zabbix bridge adapter secure.

Protect mobile devices To prevent other people from opening the app on your mobile device, enable app lock in the app *Settings*.

If a device has been lost or compromised, remove it in the Zabbix web interface. This revokes and removes its access credentials, so the device can no longer access Zabbix. Depending on your user role permissions, you can remove devices in the following sections:

- *User settings > Devices* - any user can remove their own devices (requires the *Manage own devices* permission).
- *Users > Devices* - *Super admin* can remove any user's devices (requires the *Manage user devices* permission).

You can also limit which parts of Zabbix can be accessed from a mobile device. To do this, assign users a user role that grants only the permissions the app requires. The minimal permissions are as follows:

Section	Configuration
Access to UI elements	Clear all checkboxes.
Access to services	Set *Read-write access to services* and *Read-only access to services* to `None`.
Access to modules	Clear all checkboxes.
Access to API	Clear the *Enabled* checkbox.
Access to actions	Clear all checkboxes. To let users update problems in the app, also mark the actions they need: - *Add problem comments* - *Change severity* - *Acknowledge problems* - *Suppress problems* - *Close problems* - *Create and edit own media*
Devices	Note that the problems that users can see in the app also depend on their access to hosts. Mark the *Enabled* checkbox. To let users add and remove their own mobile devices, also mark the *Manage own devices* checkbox; otherwise, a *Super admin* must manage their devices in *Users > Devices*.

Protect Zabbix bridge adapter Zabbix bridge adapter stores its data (state file and encryption keys) in the directory specified in the `DataDir` parameter of Zabbix bridge adapter configuration file. Make sure that the directory has 0700 permissions, and all files stored in it have 0600 permissions.

Zabbix bridge adapter passes requests from mobile devices to Zabbix API. To restrict which API methods can be called through Zabbix bridge adapter, use its `AllowedAPIMethods` parameter. For example, allow only the minimal set of methods that the mobile app requires, or even fewer, if you do not want some mobile app features to be used (e.g., do not allow `event.acknowledge` if mobile app users should not update problems).

Developer center

Developer center

Start developing custom extensions for Zabbix with guidance on plugins, modules, and widgets.

Build on Zabbix

Plugins An overview on how to develop and manage Zabbix plugins to extend functionality or integrate with external systems.

Python library for Zabbix Resources for using the official `zabbix_utils` library to interact with the Zabbix API and build integrations.

Frontend modules Guides and references for building custom modules that extend or modify the Zabbix frontend to suit specific use cases.

Widgets A breakdown of widget structure and logic, with instructions for creating custom dashboard elements tailored to your needs.

Creating a plugin

This page is intended for those who are interested in developing their own plugin (independent data collector) for Zabbix agent 2.

A custom plugin extends Zabbix agent 2 functionality beyond the provided built-in plugins and loadable plugins.

Each plugin is a *Go* package that defines the structure and implements one or several interfaces (*Exporter*, *Configurator*, *Runner*). For more information, see Plugin interfaces and Connection diagram.

This guide will help you create a custom loadable plugin.

For additional guidance, see also the repository of:

- Example plugin
- Existing loadable plugins

What you'll create This is ste-by-step tutorial for creating a simple loadable plugin named **MyIP**. This plugin will implement a single item key (`myip`) returning the external IP address of the Zabbix agent host.

Step 1: Setup

1. A plugin is a standard Go module. Start by initializing the `go.mod` file in the plugin directory to track plugin dependencies:

```
cd path/to/plugins/myip # Switch to your plugin directory
go mod init myip
```

2. Install the mandatory dependency Zabbix Go SDK (`golang.zabbix.com/sdk`):

```
go get golang.zabbix.com/sdk@$LATEST_COMMIT_HASH
```

Replace `$LATEST_COMMIT_HASH` with the latest `HEAD` commit hash from the `golang.zabbix.com/sdk` repository in the appropriate release branch. For example:

```
go get golang.zabbix.com/sdk@af85407
```

Note that `golang.zabbix.com/sdk` versioning is currently not supported, but this may change in the future.

Additional dependencies can be installed as needed using `go get`.

3. Create an empty `main.go` file for the plugin source code:

```
touch main.go
```

Now the initial setup is complete, and the plugin is ready for development.

Step 2: Plugin structure The `golang.zabbix.com/sdk` module, installed in the previous step, provides the necessary framework for plugin development and ensures all plugins have a consistent structure.

1. Set up basic execution flow.

Start by defining the main execution flow of the plugin. Add the following code to `main.go`:

```
package main

func main() {
    err := run()
    if err != nil {
        panic(err)
    }
}

func run() error {
    return nil
}
```

This establishes the basic execution flow for the plugin. The run function will later contain the core logic of the plugin.

2. Explore the plugin interfaces.

A Zabbix agent 2 plugin shall be represented by a struct that implements interfaces from the `golang.zabbix.com/sdk/plugin` package:

- *Accessor* - defines essential methods all plugins must implement, such as setting the plugin name and handling item key timeouts.
- One or more of the following functional plugin interfaces:
  - *Exporter* - performs a poll and returns a value (or values), nothing, or an error; often used alongside the *Collector* interface.
  - *Collector* - manages the periodic collection of data.
  - *Runner* - defines plugin startup and shutdown procedures.
  - *Watcher* - enables independent metric polling, bypassing the agent's internal scheduler; useful for trap-based or event-driven monitoring.
  - *Configurator* - defines how the plugin reads and applies its configuration settings.

You can either implement these interfaces yourself or use the default ones provided by the Zabbix Go SDK, modifying them as needed. This tutorial uses the default implementations.

3. Create the plugin struct.

Now, import the *plugin* package and create a `myIP` struct that embeds the `plugin.Base` struct:

```
import "golang.zabbix.com/sdk/plugin"

type myIP struct {
    plugin.Base
}
```

The myIP struct currently satisfies the Accessor interface. A method for implementing one of the functional plugin interfaces, the `Exporter`, will be added later in the tutorial.

Step 3: Define item keys Your plugin needs the item keys to collect data and provide it to Zabbix server or proxy.

1. Import *errs* package for error handling:

```
import "golang.zabbix.com/sdk/errs"
```

2. Register item keys using the `plugin.RegisterMetrics()` function within the `run()` function:

```
func run() error {
    p := &myIP{}

    // Register the `myip` item key.
    err := plugin.RegisterMetrics(
        p,
        "MyIP",                          // Plugin name
        "myip",                          // Item key name
        "Returns the host's IP address.", // Item key description
    )
    if err != nil {
        return errs.Wrap(err, "failed to register metrics")
    }

    return nil
}
```

To register several item keys, repeat the parameters *metric name* and *description* for each metric. For example:

```
plugin.RegisterMetrics(&impl, "Myip", "metric.one", "Metric one description.", "metric.two", "Metric two d
```

Step 4: Set up the handler The handler facilitates communication between the agent and the plugin.

1. Import the *container* package:

```
import "golang.zabbix.com/sdk/plugin/container"
```

2. Indise the `run()` function add code to create and set up a handler:

```
func run() error {
    p := &myIP{}

    // Register the `myip` item key.
    err := plugin.RegisterMetrics(
        p,
        "MyIP",                          // Plugin name
        "myip",                          // Item key name
        "Returns the host's IP address.", // Item key description
    )
    if err != nil {
        return errs.Wrap(err, "failed to register metrics")
    }

    // Create a new handler.
    h, err := container.NewHandler("MyIP") // Plugin name
```

```
    if err != nil {
        return errs.Wrap(err, "failed to create new handler")
    }

    // Setup logging to forward logs from the plugin to the agent.
    // Available via p.Logger.Infof, p.Logger.Debugf, etc.
    p.Logger = h

    // Start plugin execution.
    // Blocks until a termination request is received from the agent.
    err = h.Execute()
    if err != nil {
        return errs.Wrap(err, "failed to execute plugin handler")
    }

    return nil
}
```

Step 5: Implement data collection Data collection is done via the Exporter interface, which describes the `Export` method:

```
func Export(
  key string,             // The item key to collect.
  params []string,        // Arguments passed to the item key (`myip[arg1, arg2]`).
  context ContextProvider // Metadata about the item key data collection.
) (any, error)
```

1. Import the required packages for HTTP requests and response reading:

```
import (
    "io"
    "net/http"
)
```

2. Implement the `Export` method for the `myIP` struct:

```
func (p *myIP) Export(
    key string, params []string, context plugin.ContextProvider,
) (any, error) {
    // The plugin can use different data collection logic based on the `key` parameter.
    // This implementation only verifies that the provided `key` is supported.
    if key != "myip" {
        return nil, errs.Errorf("unknown item key %q", key)
    }

    // The log will get forwarded to the agent 2 log.
    p.Infof(
        "received request to handle %q key with %d parameters",
        key,
        len(params),
    )

    // Collect the data and return it.

    resp, err := http.Get("https://api.ipify.org")
    if err != nil {
        return nil, errs.Wrap(err, "failed to get IP address")
    }

    defer resp.Body.Close()

    body, err := io.ReadAll(resp.Body)
    if err != nil {
        return nil, errs.Wrap(err, "failed to read response body")
    }
```

```
    return string(body), nil
}
```

Step 6: Build and configure the plugin

1. To build the plugin, run:

```
go mod tidy
go build
```

This should create an executable `myip` in the current directory.

2. Configure Zabbix agent 2 to use the plugin:

```
echo "Plugins.MyIP.System.Path=$PATH_TO_THE_MYIP_PLUGIN_EXECUTABLE" > /etc/zabbix/zabbix_agent2.d/plugins.
```

Replace `$PATH_TO_THE_MYIP_PLUGIN_EXECUTABLE` with the path to the `myip` created in step 5.

The plugin name in the configuration parameter name (*MyIP* in this tutorial) must match the plugin name defined in the *plugin.RegisterMetrics()* function.

3. To test the plugin and its `myip` item, run:

```
zabbix_agent2 -c /etc/zabbix/zabbix_agent2.conf -t myip
```

The output should contain an external IP address of your host and look similar to this:

```
myip                                          [s|192.0.2.1]
```

With that, you have created a simple loadable plugin for Zabbix agent 2. Congrats!

```
package main

import (
    "io"
    "net/http"

    "golang.zabbix.com/sdk/errs"
    "golang.zabbix.com/sdk/plugin"
    "golang.zabbix.com/sdk/plugin/container"
)

var _ plugin.Exporter = (*myIP)(nil)

type myIP struct {
    plugin.Base
}

func main() {
    err := run()
    if err != nil {
        panic(err)
    }
}

func run() error {
    p := &myIP{}

    // Register the `myip` item key.
    err := plugin.RegisterMetrics(
        p,
        "MyIP",                           // Plugin name
        "myip",                           // Item key name
        "Returns the host's IP address.", // Item key description
    )
    if err != nil {
```

```
        return errs.Wrap(err, "failed to register metrics")
    }

    // Create a new handler.
    h, err := container.NewHandler("MyIP") // Plugin name
    if err != nil {
        return errs.Wrap(err, "failed to create new handler")
    }

    // Setup logging to forward logs from the plugin to the agent.
    // Available via p.Logger.Infof, p.Logger.Debugf, etc.
    p.Logger = h

    // Start plugin execution.
    // Blocks until a termination request from the agent is received.
    err = h.Execute()
    if err != nil {
        return errs.Wrap(err, "failed to execute plugin handler")
    }

    return nil
}

func (p *myIP) Export(
    key string, params []string, context plugin.ContextProvider,
) (any, error) {
    // The plugin can use different data collection logic based on the `key` parameter.
    // This implementation only verifies that the provided `key` is supported.
    if key != "myip" {
        return nil, errs.Errorf("unknown item key %q", key)
    }

    // The log will get forwarded to the agent 2 log.
    p.Infof(
        "received request to handle %q key with %d parameters",
        key,
        len(params),
    )

    // Collect the data and return it.

    resp, err := http.Get("https://api.ipify.org")
    if err != nil {
        return nil, errs.Wrap(err, "failed to get IP address")
    }

    defer resp.Body.Close()

    body, err := io.ReadAll(resp.Body)
    if err != nil {
        return nil, errs.Wrap(err, "failed to read response body")
    }

    return string(body), nil
}
```

Complete source code

Plugin interfaces

This section describes available plugin interfaces.

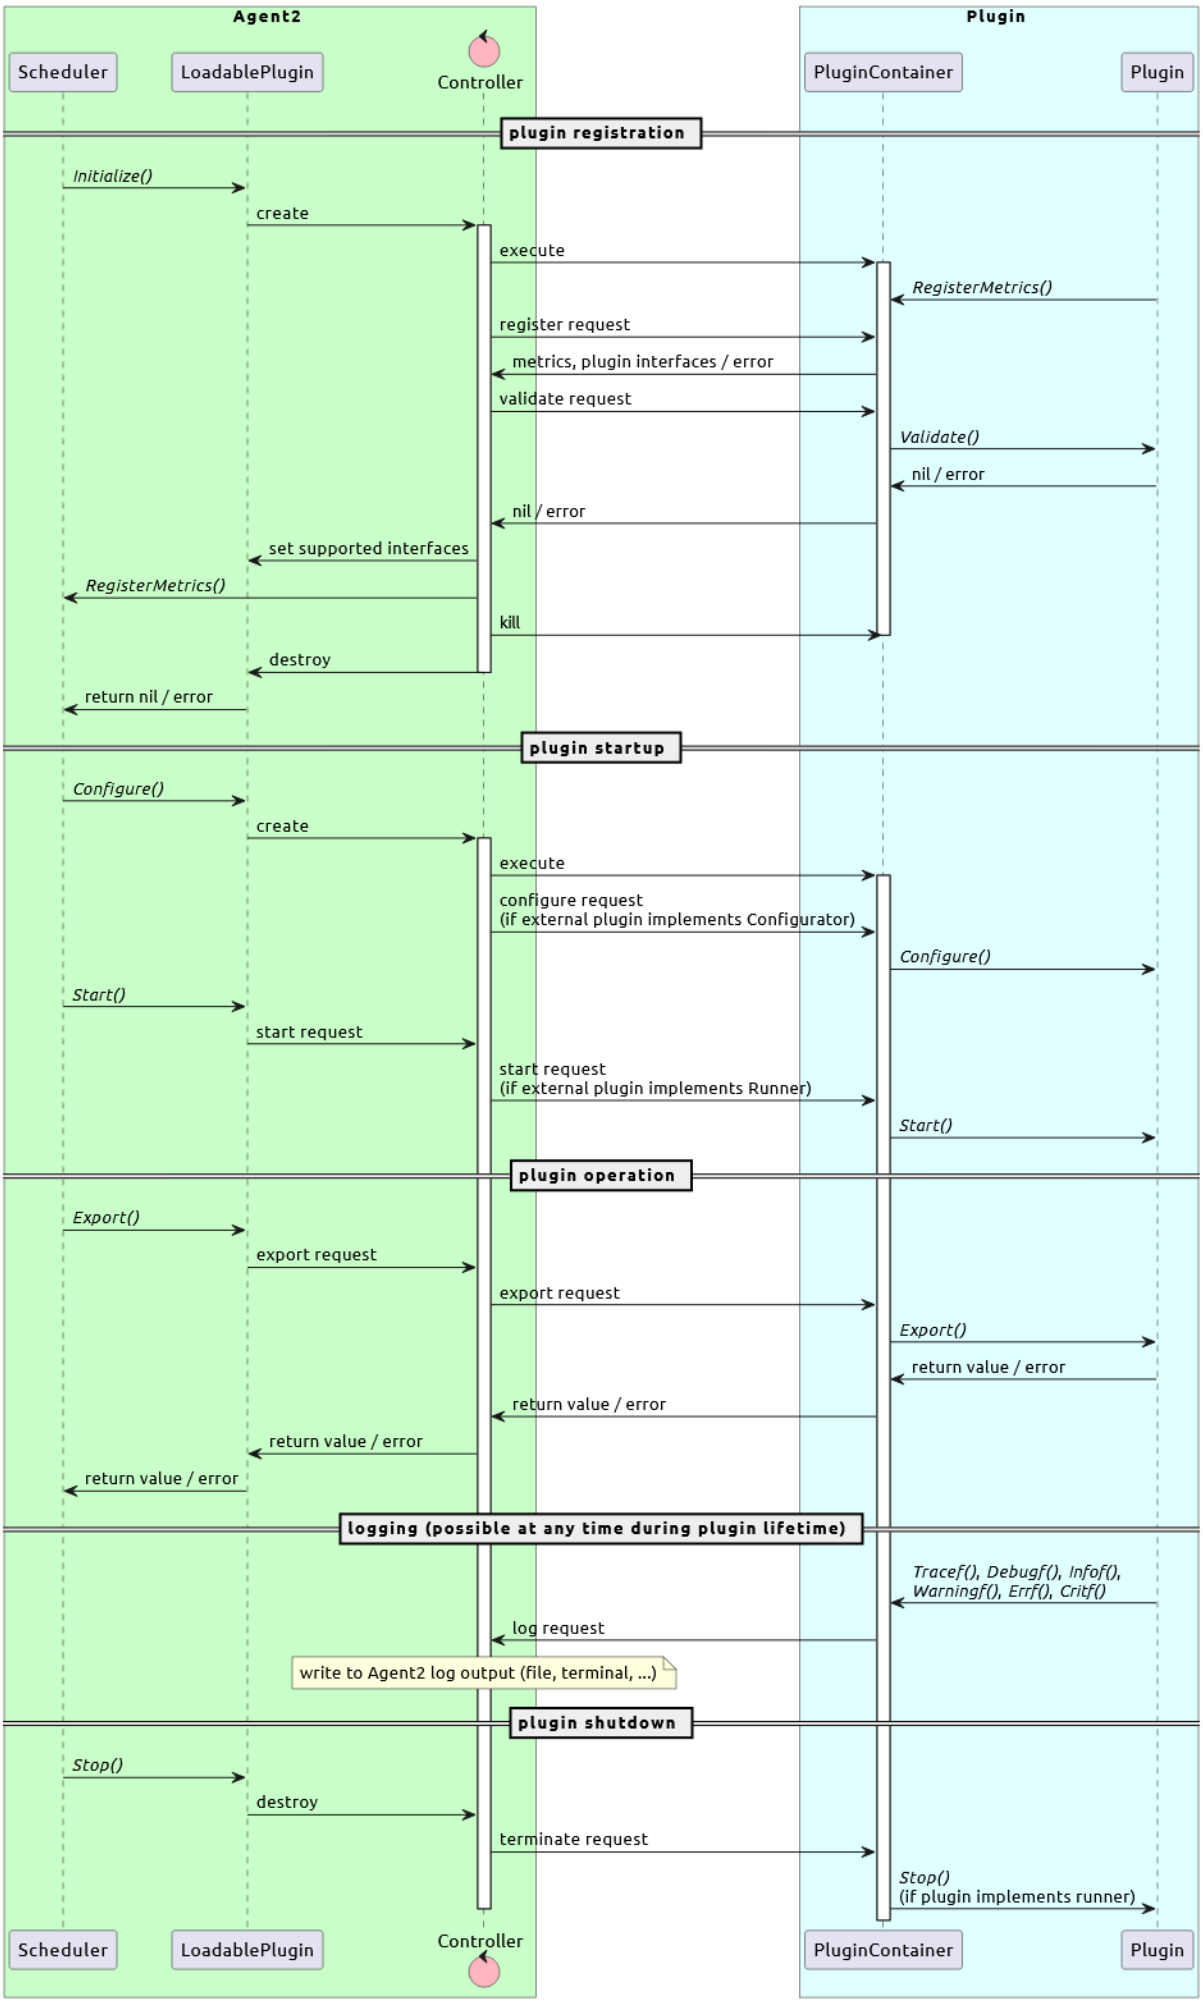
plugin.Exporter *Exporter* is the simplest interface that performs a poll and returns a value (values), nothing, or error. It accepts a preparsed item key, parameters, and context. Access to all other plugin interfaces is exclusive and no method can be called if a plugin is already performing a task. Also, there is a limit of 100 maximum concurrent *Export()* calls per plugin, which can be reduced according to the requirements for each plugin.

plugin.Configurator *Configurator* interface provides plugin configuration parameters from Zabbix agent 2 configuration files.

plugin.Runner *Runner* interface provides the means for performing initialization when a plugin is started (activated) and deinitialization when a plugin is stopped (deactivated). For example, a plugin can start/stop some background *goroutine* by implementing the Runner interface.

Connection diagram Zabbix agent 2 connects bidirectionally to the plugins using UNIX sockets on Linux and Named Pipes on Windows.

The connection diagram below illustrates the communication process between Zabbix agent 2 and a loadable plugin and the metrics collection process.


Agent2
Plugin
Scheduler
LoadablePlugin
Controller
PluginContainer
Plugin
plugin registration
Initialize()
create
execute
RegisterMetrics()
register request
metrics, plugin interfaces / error
validate request
Validate()
nil / error
nil / error
set supported interfaces
RegisterMetrics()
kill
destroy
return nil / error
plugin startup
Configure()
create
execute
configure request
(if external plugin implements Configurator)
Configure()
Start()
start request
start request
(if external plugin implements Runner)
Start()
plugin operation
Export()
export request
export request
Export()
return value / error
return value / error
return value / error
return value / error
logging (possible at any time during plugin lifetime)
Tracef(), Debugf(), Infof(),
Warningf(), Errf(), Critf()
log request
write to Agent2 log output (file, terminal, ...)
plugin shutdown
Stop()
destroy
terminate request
Stop()
(if plugin implements runner)
Scheduler
LoadablePlugin
Controller
PluginContainer
Plugin

Python library for Zabbix

zabbix_utils is the official Python library for Zabbix.

With zabbix_utils you can:

- Use Zabbix API to manage Zabbix objects (create hosts, update items, get events, etc.).
- Collect data from Zabbix agent (similarly to Zabbix get).
- Send data to Zabbix server or proxy (similarly to Zabbix sender).

The library supports both synchronous and asynchronous script execution, allowing it to handle both simple workflows and large-scale parallel operations such as mass data collection, bulk exports, alert processing, and efficient system scaling.

The library is compatible with supported Zabbix server versions and modern Python releases.

To get started, install the library and then follow the quickstart guide.

Installation

zabbix_utils can be installed using any of the methods described below.

Requirements:

- Zabbix 6.0 or later (tested on 6.0, 7.0, 7.2, 7.4)
- Python 3.8 or later (tested on 3.8–3.13)
- aiohttp async HTTP framework (required only for asynchronous mode)

Python Package Index (PyPI)

This is the most common method for most environments:

```
pip install zabbix_utils
pip install zabbix_utils[async] # Only for async support
```

Zabbix repository

Use this method if you prefer managing dependencies through your system's package manager.

Download Zabbix packages for your Linux distribution and then run the following commands.

On RHEL and derivatives:

```
dnf install python3-zabbix-utils
dnf install epel-release # Only for async support
dnf install python3-aiohttp # Only for async support
```

On Debian/Ubuntu and derivatives:

```
apt install python3-zabbix-utils
apt install python3-aiohttp # Only for async support
```

Source (GitHub)

Use this method if you prefer the latest development version:

```
git clone https://github.com/zabbix/python-zabbix-utils
cd python-zabbix-utils/
pip install -r requirements.txt # Only for async support
python3 setup.py install
```

To install dependencies required for asynchronous mode, you can also use one of the methods described above.

Quickstart guide

After you have installed zabbix_utils, you can use it in your script.

This quickstart guide shows you how to:

- Retrieve the hostname from Zabbix agent.
- Create a host and a trapper item using Zabbix API.
- Send a value to the item.

The guide presents the script step by step, explaining each part as it is introduced. The complete script is provided at the end of the page.

The guide also assumes your Zabbix server, agent, and API are running locally.

Retrieve hostname from Zabbix agent

Start by retrieving the hostname of the system where Zabbix agent is running. You will need this hostname to create a host using Zabbix API.

1. Import the `Getter` class from zabbix_utils. This class works like Zabbix get and lets you request data from Zabbix agent.

2. Create a `Getter` instance that connects to the local Zabbix agent at `127.0.0.1` on port `10050`.

3. Call the `get()` method on the `Getter` instance to request the hostname from Zabbix agent.

- The `get()` method takes an item key as its parameter and sends a request to Zabbix agent for that item.
- The `get()` method returns an object, and the hostname is stored inside the `value` attribute of that object.
- `system.hostname` is a built-in Zabbix agent key that returns the host's name.

```
from zabbix_utils import Getter

agent = Getter(host='127.0.0.1', port=10050)
hostname = agent.get('system.hostname').value
```

Connect and log in to Zabbix API

Next, connect to the Zabbix API. This lets your script manage Zabbix objects such as hosts and items.

1. Import the `ZabbixAPI` class from zabbix_utils.

2. Create a `ZabbixAPI` instance and provide the URL of your Zabbix web interface.

3. Call the `login()` method on the `ZabbixAPI` instance and provide your username and password. The account must have permission to access Zabbix API.

4. Call the `logout()` method to close the session after API operations.

```
from zabbix_utils import ZabbixAPI

api = ZabbixAPI(url='127.0.0.1/zabbix')
api.login(user='Admin', password='zabbix')

### API operations go here

api.logout()
```

Create host in Zabbix

Now that you have the hostname from Zabbix agent and you're connected to Zabbix API, you can create a new host in Zabbix.

1. Call the `host.create()` API method on the `ZabbixAPI` instance and provide the host details:

- `host` - set to the `hostname` variable containing the hostname you retrieved from Zabbix agent.
- `interfaces` - contains connection details for Zabbix agent that's running on the host.
- `groups` - contains at least one host group the host should belong to.

2. Since Zabbix API returns the ID of the newly created host, save this ID in the `host_id` variable for later use.

3. Print a message to confirm the host was created.

```
api.host.create(
    host=hostname,
    interfaces=[{
        'type': 1,
        'main': 1,
        'useip': 1,
        'ip': '127.0.0.1',
        'dns': '',
        'port': '10050',
```

```
    }],
    groups=[{'groupid': '2'}]
)

host_id = host['hostids'][0]

print(f"Host '{hostname}' created with ID {host_id}")
```

Create new item in Zabbix

After creating a host, you can add an item to it.

1. Define a unique key for the item to be created.

2. Call the `item.create()` method on the `ZabbixAPI` instance and provide the item details:

   - `hostid` - set to the `host_id` variable containing the ID of the host you just created.
   - `name` - a name for the item.
   - `key_` - set to the `item_key` variable you just defined.
   - `type` - set to 2 (trapper item), required to receive values sent from your script in the following steps.
   - `value_type` - set to 3 (numeric unsigned), the type of data this item stores.

3. Since Zabbix API returns the ID of the newly created item, save this ID in the `item_id` variable for later use.

4. Print a message to confirm the item was created.

```
item_key = 'app.myservice.heartbeat'

item = api.item.create(
    hostid=host_id,
    name='App heartbeat',
    key_=item_key,
    type=2,
    value_type=3,
)

item_id = item['itemids'][0]

print(f"Item '{item_key}' created with ID {item_id}")
```

Send value to host

Now that you've created your host and item, you can send data to it.

1. Import the `time` class and wait a few seconds before sending the data. This ensures Zabbix has fully processed your new host and item.

2. Import the `Sender` class from zabbix_utils. This class works like Zabbix sender and lets your script send data to Zabbix.

3. Create a `Sender` instance that connects to the local Zabbix server at `127.0.0.1` on port `10051`.

4. Call the `send_value()` method on the `Sender` instance to send a value to your host's item and provide these details:

   - `hostname` - set to the `hostname` variable containing the hostname you retrieved from Zabbix agent.
   - `item_key` - set to the `item_key` variable you just defined earlier in your script.
   - `1` - the value to send.

4. Print a message to confirm the value was sent.

```
import time
time.sleep(10)

from zabbix_utils import Sender

sender = Sender(server='127.0.0.1', port=10051)

response = sender.send_value(hostname, item_key, 1)

print(f"Sender response: {response}")
```

After this code runs successfully, you should see a success message indicating Zabbix received your value:

```
Sender response: {"processed": 1, "failed": 0, "total": 1, "time": "0.000151", "chunk": 1}
```

Now you can also check your Zabbix web interface (*Monitoring* > *Latest data*) to see the value.

Complete script

Below is the complete script that combines all the steps: retrieving the hostname, creating the host and item using Zabbix API, and sending data to Zabbix.

```
import time
from zabbix_utils import Getter, ZabbixAPI, Sender

### Retrieve host name from Zabbix agent
agent = Getter(host='127.0.0.1', port=10050)
hostname = agent.get('system.hostname').value

### Connect and log in to Zabbix API
api = ZabbixAPI(url='127.0.0.1/zabbix')
api.login(user='Admin', password='zabbix')

### Create host in Zabbix
host = api.host.create(
    host=hostname,
    interfaces=[{
        'type': 1,
        'main': 1,
        'useip': 1,
        'ip': '127.0.0.1',
        'dns': '',
        'port': '10050',
    }],
    groups=[{'groupid': '2'}]
)

host_id = host['hostids'][0]

print(f"Host '{hostname}' created with ID {host_id}")

### Create item in Zabbix
item_key = 'app.myservice.heartbeat'

item = api.item.create(
    hostid=host_id,
    name='App heartbeat',
    key_=item_key,
    type=2,
    value_type=3
)

item_id = item['itemids'][0]

print(f"Item '{item_key}' created with ID {item_id}")

### Log out from API
api.logout()

### Wait for Zabbix to process new host
time.sleep(10)

### Send value to host
sender = Sender(server='127.0.0.1', port=10051)
response = sender.send_value(hostname, item_key, 1)

print(f"Sender response: {response}")
```

Use Zabbix API

zabbix_utils lets you use Zabbix API to manage Zabbix objects, including creating hosts, updating items, retrieving events, and more.

API requests can be made in synchronous or asynchronous mode:

- In synchronous mode, your Python script sends a request and waits for a response before continuing, which is suitable for simple, sequential, and predictable operations.
- In asynchronous mode, the script sends requests without waiting for each response, allowing other operations to proceed in parallel; this is more efficient for slow requests or large batches of data.

The examples on this page focus on synchronous mode, though asynchronous mode follows similar patterns. Additional examples are available in the zabbix_utils GitHub repository.

Import

To use zabbix_utils for working with Zabbix API, import the `ZabbixAPI` class in your script:

```
from zabbix_utils import ZabbixAPI
```

> **Note:**
> If you already use a community Python library, you can usually replace that import with `ZabbixAPI` from zabbix_utils.

Log in

Before making API requests, you must create a `ZabbixAPI` instance and log in to Zabbix API.

Choose the method that best fits how you manage credentials. You can use your username and password or API tokens.

During initialization

When creating a `ZabbixAPI` instance, you can provide the Zabbix web interface URL and your credentials all at once:

```
### Username and password:
api = ZabbixAPI(url="127.0.0.1/zabbix", user="Admin", password="zabbix")

### API token:
api = ZabbixAPI(url="127.0.0.1/zabbix", token="your_api_token")
```

If you prefer to group connection parameters, you can pass them as a Python dictionary and unpack it using Python's keyword argument syntax:

```
ZABBIX_AUTH = {
    "url": "127.0.0.1/zabbix",
    "user": "Admin",
    "password": "zabbix"
}

api = ZabbixAPI(**ZABBIX_AUTH)
```

Using `login()`

You can first create a `ZabbixAPI` instance with just the Zabbix web interface URL, and later use the `login()` method with your credentials:

```
### Username and password:
api = ZabbixAPI(url="127.0.0.1/zabbix")
api.login(user="Admin", password="zabbix")

### API token:
api = ZabbixAPI(url="127.0.0.1/zabbix")
api.login(token="your_api_token")
```

Using environment variables

You can first store your credentials as environment variables, for example, by using a terminal:

```
### Username and password:
export ZABBIX_USER="Admin"
export ZABBIX_PASSWORD="zabbix"

### Token:
export ZABBIX_TOKEN="your_api_token"
```

Then, if you start your script from the same terminal session, the `ZabbixAPI` instance in the script requires only the Zabbix web interface URL:

```
api = ZabbixAPI(url="127.0.0.1/zabbix")
```

You can also store the URL in an environment variable:

```
### Username and password:
export ZABBIX_USER="Admin"
export ZABBIX_PASSWORD="zabbix"
export ZABBIX_URL="https://127.0.0.1/zabbix"

### Token:
export ZABBIX_TOKEN="your_api_token"
export ZABBIX_URL="https://127.0.0.1/zabbix"
```

When the URL and credentials are all set as environment variables, the `ZabbixAPI` instance in the script does not require any arguments at all:

```
api = ZabbixAPI()
```

Log out

If you logged in using a username and password, call the `logout()` method after completing your API operations:

```
from zabbix_utils import ZabbixAPI

api = ZabbixAPI(url="127.0.0.1/zabbix")
api.login(user="Admin", password="zabbix")

### API actions go here

api.logout()
```

Note:
Calling the `logout()` method is not required when using API tokens.

API requests

After you've logged in, you can make any API requests by calling methods described in the Zabbix API method reference.

API methods are called using the following format:

```
api_instance.zabbix_object.method(parameters)
```

Note:
Some Zabbix API method or object names use words that are reserved keywords in Python (e.g., `import`). To avoid Python errors, add an underscore to the method or object name when using it in Python (e.g., `api.configuration.import_`).

For example, to see a list of hosts with `host.get`:

```
### 1. Import ZabbixAPI from zabbix_utils:
from zabbix_utils import ZabbixAPI

### 2. Create the ZabbixAPI instance and log in:
api = ZabbixAPI(url="127.0.0.1/zabbix")
api.login(user="Admin", password="zabbix")

### 3. Retrieve hosts matching a filter:
hosts = api.host.get(
```

```
    filter={
        "host": [
            "Zabbix server",
            "Linux server"
        ]
    }
)

### 4. Iterate over returned hosts and print each host's details:
for host in hosts:
    print(host)

### 5. Log out from the API to close the session:
api.logout()
```

Asynchronous mode

Asynchronous mode lets your script send API requests without waiting for each one to finish. This can make your script more efficient when it needs to perform many API requests or when some requests take a long time.

When using asynchronous mode, there are important differences compared to synchronous mode:

- Import Python's `asyncio` module (you must first install the required dependencies).
- Import `AsyncZabbixAPI` instead of `ZabbixAPI`.
- Write your code inside an `async` function.
- Use `await` when calling API methods, including `login()` and `logout()`.

> **Attention:**
> When creating an `AsyncZabbixAPI` instance, you cannot log in during initialization.

For example, to see a list of hosts with `host.get` in asynchronous mode:

```
### 1. Import asyncio for asynchronous mode, and AsyncZabbixAPI from zabbix_utils:
import asyncio
from zabbix_utils import AsyncZabbixAPI

### 2. Define the main async function where all API operations will run:
async def main():

    # 3. Create an AsyncZabbixAPI instance and log in (must await):
    api = AsyncZabbixAPI(url="127.0.0.1")
    await api.login(user="User", password="zabbix")

    # 4. Retrieve hosts matching a filter (must await):
    hosts = await api.host.get(
        filter={
            "host": [
                "Zabbix server",
                "Linux server"
            ]
        }
    )

    # 5. Iterate over returned hosts and print each host's details:
    for host in hosts:
        print(host)

    # 6. Log out from the API to close the session (must await):
    await api.logout()

### 7. Run the async main() function using asyncio's event loop:
asyncio.run(main())
```

Examples

The following examples show common Zabbix API tasks using the zabbix_utils library.

Additional examples are available in the `examples` directory of the zabbix_utils GitHub repository.

Collect data from Zabbix agent

Overview

zabbix_utils lets you collect data from Zabbix agent (similarly to Zabbix get).

Data can be collected in synchronous or asynchronous mode:

- In synchronous mode, your Python script requests and waits for data before continuing, which is suitable for simple, sequential, and predictable operations.
- In asynchronous mode, the script requests data without waiting for each response, allowing other operations to proceed in parallel; this is more efficient for slow requests or large batches of data.

The examples on this page focus on synchronous mode, though asynchronous mode follows similar patterns. Additional examples are available in the zabbix_utils GitHub repository.

Import

To use zabbix_utils for collecting item values, import the `Getter` class in your Python script:

```
from zabbix_utils import Getter
```

Request data

To request an item value:

1. Create a `Getter` instance, specifying the IP address and port of your Zabbix agent.
2. Call the `get()` method on the `Getter` instance, specifying the key of the item you want to retrieve.

For example, to request data for the `system.uname` item:

```
agent = Getter(host='192.0.2.1', port=10050)
response = agent.get('system.uname')
```

Using non-default IP

If the server running your script has multiple IP addresses, you can specify a `source_ip` for the `Getter` to use when connecting to Zabbix agent:

```
agent = Getter(
    host='192.0.2.1',
    port=10050,
    source_ip='10.10.7.1'
)
```

Using timeout

You can set a response `timeout` for the `Getter` to control how long your script should wait for a response from Zabbix agent before giving up:

```
agent = Getter(
    host='192.0.2.1',
    port=10050,
    timeout=30
)
```

Using encryption

The `Getter` does not include built-in encryption support, but you can provide it by creating a wrapper using third-party libraries:

```
def psk_wrapper(sock, tls):
    # ...
    # Implementation of TLS PSK wrapper for the socket
    # ...

agent = Getter(
    host='192.0.2.1',
```

```
    port=10050,
    socket_wrapper=psk_wrapper
)
```

Response

The response from Zabbix agent is processed by the library and returned as an `AgentResponse` object:

```
print(response)
### {
###     "error": null,
###     "raw": "Linux zabbix_server 5.15.0-3.60.5.1.el9uek.x86_64",
###     "value": "Linux zabbix_server 5.15.0-3.60.5.1.el9uek.x86_64"
### }

print(response.value)
### Linux zabbix_server 5.15.0-3.60.5.1.el9uek.x86_64

print(response.error)
### None
```

Asynchronous mode

Asynchronous mode lets your script collect values without waiting for each one to arrive. This can make your script more efficient when it needs to collect many values or when some values take a long time to be collected.

When using asynchronous mode, there are important differences compared to synchronous mode:

- Import Python's `asyncio` module (you must first install the required dependencies).
- Import `AsyncGetter` instead of `Getter`.
- Write your code inside an `async` function.
- Use `await` when calling the `get()` method.

For example, to collect a single value using asynchronous mode:

```
### 1. Import asyncio for asynchronous mode, and AsyncGetter from zabbix_utils:
import asyncio
from zabbix_utils import AsyncGetter

### 2. Define the main async function where all data requests will be executed:
async def main():
    agent = AsyncGetter(host='192.0.2.1', port=10050)

    # 3. Fetch the system.uname value from Zabbix agent (must await):
    response = await agent.get('system.uname')

    # 4. Print the value returned by Zabbix agent:
    print(response.value)

### 5. Run the async main() function using asyncio's event loop:
asyncio.run(main())
```

Send data to Zabbix server or proxy

zabbix_utils lets you send item values to a trapper item on Zabbix server or proxy (similarly to Zabbix sender).

You can send a single value, multiple values, or even target multiple Zabbix clusters.

Data can be sent in synchronous or asynchronous mode:

- In synchronous mode, your Python script sends values and waits for a response before continuing; this is suitable for simple, sequential, and predictable operations.
- In asynchronous mode, the script sends values without waiting for each response, allowing other operations to proceed in parallel; this is more efficient for slow requests or large batches of data.

The examples on this page focus on synchronous mode, though asynchronous mode follows similar patterns. Additional examples are available in the zabbix_utils GitHub repository.

Import

To use zabbix_utils for sending item values, import the Sender class in your script:

```
from zabbix_utils import Sender
```

For sending multiple values, you may also import the ItemValue class:

```
from zabbix_utils import Sender, ItemValue
```

Send single value

To send an item value:

1. Create a Sender instance, specifying the IP address and port of your Zabbix server or proxy.
2. Call the send_value() method on the Sender instance using the following format:

```
sender_instance.send_value('host', 'item.key', 'value', optional_timestamp, optional_nanoseconds)
```

For example, to send 1 to the service.status trapper item on the Linux server host:

```
sender = Sender(server='127.0.0.1', port=10051)
response = sender.send_value('Linux server', 'service.status', 1)
```

Using non-default IP

If the server running your script has multiple IP addresses, you can specify a source_ip for the Sender to use when sending values to Zabbix server or proxy:

```
sender = Sender(
    server='127.0.0.1',
    port=10051,
    source_ip='10.10.7.1'
)
```

Using timeout

You can set a response timeout for the Sender to control how long your script should wait for a response from Zabbix server or proxy before giving up:

```
sender = Sender(
    server='127.0.0.1',
    port=10051,
    timeout=30
)
```

Using agent configuration file

You can let zabbix_utils read the Server or ServerActive parameters from a local Zabbix agent or agent 2 configuration file. In such cases, you don't need to specify connection parameters when creating a Sender instance:

```
sender = Sender(
    use_config=True,
    config_path='/etc/zabbix/zabbix_agent2.conf'
)
```

> **Attention:**
> If ServerActive contains one or more Zabbix clusters with multiple server instances, Sender sends data to the first available server in each cluster. If ServerActive is not set, the address from Server with the default port (10051) is used.

Using encryption

The Sender does not include built-in encryption support, but you can provide it by creating a wrapper using third-party libraries:

```
def psk_wrapper(sock, tls):
    # ...
    # Implementation of TLS PSK wrapper for the socket
    # ...
```

```
sender = Sender(
    server='127.0.0.1',
    port=10051,
    socket_wrapper=psk_wrapper
)
```

Response for single value

The response returned by Zabbix server or proxy is processed by the library and returned as a `TrapperResponse` object:

```
print(response)
### {"processed": 1, "failed": 0, "total": 1, "time": "0.000123", "chunk": 1}

print(response.processed)
### 1

print(response.failed)
### 0

print(response.total)
### 1
```

Asynchronous mode

Asynchronous mode lets your Python script send values without waiting for a response from Zabbix server or proxy. This can make your script more efficient when it needs to send many values or when some values take a long time send.

When using asynchronous mode, there are a few important differences compared to synchronous mode:

- Import Python's `asyncio` module (you must first install the required dependencies).
- Import `AsyncSender` instead of `Sender`.
- Write your code inside an `async` function.
- Use `await` when calling the `send_value()` method.

For example, to send a single value using asynchronous mode:

```
### 1. Import asyncio for asynchronous mode, and AsyncSender from zabbix_utils:
import asyncio
from zabbix_utils import AsyncSender

### 2. Define the main async function where all data sending operations (must await) will be executed:
async def main():
    sender = AsyncSender(server='127.0.0.1', port=10051)
    response = await sender.send_value('Linux server', 'service.status', 1)

    # 3. Print the response returned by Zabbix server or proxy:
    print(response)

### 4. Run the async main() function using asyncio's event loop:
asyncio.run(main())
```

Send multiple values

To send multiple values:

1. Prepare an array of `ItemValue` objects, each using the same format as the `send_value()` method.
2. Create a `Sender` instance, specifying the IP address and port of your Zabbix server or proxy.
3. Call the `send()` method (instead of `send_value()`) on the `Sender` instance, specifying the array of objects with values to send.

For example, to send five values to different hosts:

```
items = [
    ItemValue('server-de', 'service.status', 'up', 1770887205, 100),
    ItemValue('server-fr', 'service.status', 'up', 1770887205, 100),
    ItemValue('server-uk', 'service.status', 'up', 1770887205, 100),
    ItemValue('server-nl', 'service.status', 'up', 1770887205, 100),
    ItemValue('server-pl', 'service.status', 'up', 1770887205, 100),
]
```

```
sender = Sender(server='127.0.0.1', port=10051)
response = sender.send(items)
```

Using custom chunk size

If you need to send more values than a trapper item can accept in a single request, you can split them into chunks.

By default, the chunk size is 250 values. You can change it by setting the `chunk_size` parameter when creating a `Sender` instance.

For example, to send five values in three chunks (2-2-1), set the `chunk_size` parameter to 2:

```
items = [
    ItemValue('server-de', 'service.status', 'up'),
    ItemValue('server-fr', 'service.status', 'up'),
    ItemValue('server-uk', 'service.status', 'up'),
    ItemValue('server-nl', 'service.status', 'up'),
    ItemValue('server-pl', 'service.status', 'up'),
]

sender = Sender(server='127.0.0.1', port=10051, chunk_size=2)
response = sender.send(items)
```

Send values to multiple Zabbix clusters

To send values to multiple Zabbix clusters:

1. Prepare an array of Zabbix clusters. If a cluster has multiple nodes, the value will be sent to the first **available** node of each cluster.
2. Create a `Sender`, specifying your array of Zabbix clusters.
3. Call the `send_value()` method on the `Sender` instance using the same format as the `send_value()` method.

For example, to send a value to the first available node in each cluster:

```
zabbix_clusters = [
    ['zabbix.cluster1.node1', 'zabbix.cluster1.node2:10051'],
    ['zabbix.cluster2.node1:10051', 'zabbix.cluster2.node2', 'zabbix.cluster2.node3']
]

sender = Sender(clusters=zabbix_clusters)
response = sender.send_value('Linux server', 'service.status', 1)
```

Response for multiple values

By default, `Sender` returns an aggregated result of sending values across all hosts or clusters:

```
print(response)
### {"processed": 2, "failed": 0, "total": 2, "time": "0.000108", "chunk": 2}
```

If you need more detailed information, you can inspect the results for each cluster and each chunk using the `response.details` attribute:

```
print(response)
### {"processed": 2, "failed": 0, "total": 2, "time": "0.000108", "chunk": 2}

if response.failed == 0:
    print(f"Value sent successfully in {response.time}")
else:
    print(response.details)
    # {
    #     127.0.0.1:10051: [
    #         {
    #             "processed": 1,
    #             "failed": 0,
    #             "total": 1,
    #             "time": "0.000051",
    #             "chunk": 1
    #         }
```

```
    #     ],
    #     zabbix.example.local:10051: [
    #         {
    #             "processed": 1,
    #             "failed": 0,
    #             "total": 1,
    #             "time": "0.000057",
    #             "chunk": 1
    #         }
    #     ]
    # }

    for node, chunks in response.details.items():
        for resp in chunks:
            print(f"Processed {resp.processed} of {resp.total} at {node.address}:{node.port}")
            # Processed 1 of 1 at 127.0.0.1:10051
            # Processed 1 of 1 at zabbix.example.local:10051
```

Debug logging

To troubleshoot issues when working with Zabbix API, collecting data from Zabbix agent, or sending data to Zabbix server or proxy, you can enable debug logging.

The zabbix_utils library uses Python's standard `logging` module.

To see detailed debug messages from the library:

1. Import Python's `logging` module in your script.
2. Configure the logging system to output messages in a readable format and set the logging level to DEBUG.

```
import logging
from zabbix_utils import Getter

logging.basicConfig(
    format=u'[%(asctime)s] %(levelname)s %(message)s',
    level=logging.DEBUG
)

agent = Getter(host='127.0.0.1', port=10050)
resp = agent.get('system.uname')

print(resp.value)
```

With debug logging enabled, detailed information about requests and responses is printed, for example:

```
[2023-10-01 12:00:01,587] DEBUG Content of the packet: b'ZBXD\x01\x0c\x00\x00\x00\x00\x00\x00\x00system.un
[2023-10-01 12:00:01,722] DEBUG Zabbix response header: b'ZBXD\x01C\x00\x00\x00C\x00\x00\x00'
[2023-10-01 12:00:01,723] DEBUG Zabbix response body: Linux test_server 5.15.0-3.60.5.1.el9uek.x86_64
[2023-10-01 12:00:01,724] DEBUG Response from [127.0.0.1:10050]: Linux test_server 5.15.0-3.60.5.1.el9uek.
Linux test_server 5.15.0-3.60.5.1.el9uek.x86_64
```

Modules

What is a PHP frontend module?

- A module is an entity with a unique ID, name, description, author, and other fields defined in it's manifest file, along with PHP, Javascript and other files located in a single directory inside the *modules* directory of your Zabbix frontend installation (for example, *zabbix/ui/modules*).
- A module should conform to simple rules to guarantee correct operation.
- A module must be installed (unpacked) and enabled in the frontend by the administrator.

What a module can be used for

- Adding new functionality via custom frontend sections;
- Creating custom dashboard widget types (see widget modules);
- Overriding or extending the existing functionality.

What a module cannot be used for

- Registering a new API method or modifying an existing one.

How modules work

- An enabled module is launched on each HTTP request, before executing the action code.
- The module will register new actions or redefine the existing ones.
- The module will add new frontend sections and remove or redefine the existing ones.
- The module will hook to frontend events like onBeforeAction and onTerminate, if needed.
- The requested action is finally executed by running the action code - either the default one, or module-defined.

Where to go next Whether you prefer to learn by doing or read the guidelines first, these pages contain the information and steps required to build your own modules:

- Step-by-step tutorials for writing your first module
- Module file structure
- Widget module specifics
- Module examples to reuse

Module file structure

All code related to a module is stored in a single directory inside the **modules** directory of your Zabbix frontend installation (for example, *zabbix/ui/modules*).

```
example_module_directory/          (required)
   manifest.json                   (required)  Metadata and action definition.
   Module.php                                  Module initialization and event handling.
   actions/                                    Action controller files.
       SomethingView.php
       SomethingCreate.php
       SomethingDelete.php
       data_export/
           ExportAsXml.php
           ExportAsExcel.php
   views/                                      View files.
       example.something.view.php
       example.something.delete.php
   assets/                                     Any additional files to be used in views. Must be specified
       js/                                     JavaScript files used in views.
           example.something.view.js.php
       css/                                    CSS files used in views.
           example.something.css
        image.png                              Images used in views.
        example.something.file                 Any file for using in views.
```

Module file tree

Writing a module A sample module writing process consists of the following steps (where available, click the file or folder name to view additional details about the step):

1. Create a new directory for the module inside **zabbix/ui/modules/**.
2. Add manifest.json file with module metadata.
3. Create *views* folder and define a module view(s).
4. Create *actions* folder and define a module action(s).
5. Create Module.php (or Widget.php for dashboard widgets) file and define initialization and event handling rules.

6. Create *assets* folder for JavaScript files (place into *assets/js*), CSS styles (place into *assets/css*), or any other additional files.
7. Make sure to specify required views, actions and asset files in the manifest.json.
8. Register the module in Zabbix frontend and start using it.

Note:
You can register and enable a module as soon as you create manifest.json file. Once the module is enabled, you can preview all changes made to module files immediately by refreshing Zabbix frontend.

manifest.json

Any module needs the manifest.json file. The file must be located in the module's primary directory (for example, *zabbix/ui/modules/module_name/manifest.json*).

As a bare minimum, manifest.json should specify these fields:

```
{
    "manifest_version": 2.0,
    "id": "my_ip_address",
    "name": "My IP Address",
    "namespace": "MyIPAddress",
    "version": "1.0"
}
```

Parameters supported in **manifest.json** (click the parameter name for a detailed description):

Parameter	Description	Required
manifest_version	Manifest version of the module.	Yes
id	Unique module ID.	
name	Module name that will be displayed in the Administration section.	
namespace	PHP namespace for module classes.	
version	Module version.	
type	Type of the module. For widget must be set to *widget*	Yes for widgets, otherwise no
widget	Widget configuration. Used for widgets only.	
actions	Actions to register with the module.	
assets	CSS styles and JavaScript files to include.	No
author	Module author.	
config	Default values for custom module options.	
description	Module description.	
url	A link to the module description.	

manifest_version

Manifest version of the module. Currently, supported version is **2.0**.

Type: Double

Example:

```
"manifest_version": 2.0
```

id

Module ID. Must be unique. To avoid future naming conflicts, it is recommended to use prefix for modules (author or company name, or any other). For example, if a module is an example for lessons and the module name is "My module", then the ID will be "example_my_module".

Type: String

Example:

```
"id": "example_my_module"
```

name

Module name that will be displayed in the Administration section.

Type: String

Example:

```
"name": "My module"
```

namespace

PHP namespace for module classes.

Type: String

Example:

```
"namespace": "ClockWidget"
```

version

Module version. The version will be displayed in the Administration section.

Type: String

Example:

```
"version": "1.0"
```

type

Type of the module. Required for widgets and must equal "widget".

Type: String

Default: "module"

Example:

```
"type": "widget"
```

actions

Actions to register with the module. Defining *class* object key for each action is required, other action keys are optional.

Type: Object

Supported object keys if *type* is *module*:

- **write.your.action.name** (object) - action name, should be written in lowercase [a-z], separating words with dots. Supports the keys:
  - **class** (string; required) - action class name.
  - **layout** (string) - action layout. Supported values: *layout.json*, *layout.htmlpage* (default), *null*.
  - **view** (string) - action view.

Example:

```
"actions": {
    "module.example.list": {
        "class": "ExampleList",
        "view": "example.list",
        "layout": "layout.htmlpage"
        }
    }
```

Supported object keys if *type* is *widget*:

- **widget.{id}.view** (object) - file and class name for widget view. Replace **{id}** with the widget's id value (for example, *widget.example_clock.view*). Supports the keys:
  - **class** (string; required) - action class name for widget view mode to extend the default CControllerDashboardWidgetView class. The class source file must be located in the *actions* directory.
  - **view** (string) - widget view. Must be located in the *views* directory. If the view file is *widget.view.php*, which is expected by default, this parameter maybe omitted. If using a different name, specify it here.
- **widget.{id}.edit** (object) - file name for widget configuration view. Replace **{id}** with the widget's id value (for example, *widget.example_clock.edit*). Supports the keys:
  - **class** (string; required) - action class name for widget configuration view mode. The class source file must be located in the *actions* directory.
  - **view** (string) - widget configuration view. Must be located in the *views* directory. If the view file is *widget.edit.php*, which is expected by default, this parameter maybe omitted. If using a different name, specify it here.

Example:

```
"actions": {
    "widget.tophosts.view": {
        "class": "WidgetView"
    },
    "widget.tophosts.column.edit": {
        "class": "ColumnEdit",
        "view": "column.edit",
        "layout": "layout.json"
    }
}
```

assets

CSS styles and JavaScript files to include.

Type: Object

Supported object keys:

- **css** (array) - CSS files to include. The files must be located in the *assets/css*.
- **js** (array) - JavaScript files to include. The files must be located in the *assets/js*.

Example:

```
"assets": {
    "css": ["widget.css"],
    "js": ["class.widget.js"]
}
```

author

Module author. The author will be displayed in the Administration section.

Type: String

Example:

```
"author": "John Smith"
```

config

Default values for the module options. The object may contain any custom keys. If specified, these values will be written into the database during module registration. New variables added later will be written upon the first call. Afterwards, the variable values can only be changed directly in the database.

Type: Object

Example:

```
"config": {
    "username": "Admin",
    "password": "",
    "auth_url": "https://example.com/auth"
}
```

description

Module description.

Type: String

Example:

```
"description": "This is a clock widget."
```

widget

Widget configuration. Used, if *type* is set to *widget*.

Type: Object

Supported object keys:

- **name** (string) - used in the widget list and as default header. If empty, ”name” parameter from the module will be used.

- **size** (object) - default widget dimensions. Supports keys:
  - *width* (integer) - default widget width.
  - *height* (integer) - default widget height.
- **form_class** (string) - class with widget fields form. Must be located in the *includes* directory. If the class is *WidgetForm.php*, which is expected by default, this parameter maybe omitted. If using a different name, specify it here.
- **js_class** (string) - name of a JavaScript class for widget view mode to extend the default CWidget class. The class will be loaded with the dashboard. The class source file must be located in the assets/js directory. See also: assets.
- **use_time_selector** (boolean) - determines whether the widget requires dashboard time selector. Supported values: *true*, *false* (default).
- **refresh_rate** (integer) - widget refresh rate in seconds (default: 60).

Example:

```
"widget": {
    "name": "",
    "size": {
        "width": 12,
        "height": 5
    },
    "form_class": "WidgetForm",
    "js_class": "CWidget",
    "use_time_selector": false,
    "refresh_rate": 60
}
```

url

A link to the module description. For widgets, this link will be opened when clicking the help icon in the *Add widget* or *Edit widget* window. If **url** is not specified, clicking the help icon will open the general Dashboard widgets page.

Type: String

Example:

```
"url": "http://example.com"
```

Actions

Actions are responsible for 'business logic' of the module. An action usually consists of a controller and an action view.

A module can:

- Call actions that are already defined in Zabbix frontend.
- Override default actions with custom actions.
- Define completely new actions.

To override a default action behavior with some custom behavior, define an action with the same name in the module configuration. When the action is called, the module action will be executed instead of the default Zabbix action.

Action files should be stored in the *actions* folder. The actions need to be specified in the manifest.json.

Controller

Action controller workflow:

1. Check that all parameters passed in an HTTP request are valid:
   - Call the controller's *checkInput()* method;
   - Use validation rules defined in CNewValidator.php;
   - Call *validateInput()* method.
2. Check user permissions.
3. Prepare the data according to passed parameters: if *checkInput()* returns true, Zabbix calls the controller's *doAction()* method.
4. Prepare the **$data** array for the view. Use *CControllerResponseData* and *setResponse()* method to store response in the **$data** array.

Example:

```
/**
 * Validate input parameters.
 *
 * @return bool
 */
protected function checkInput(): bool {
    $ret = $this->validateInput([
        'status' => 'in '.implode(',', [HOST_STATUS_MONITORED, HOST_STATUS_NOT_MONITORED])
    ]);

    if (!$ret) {
        $this->setResponse(new CControllerResponseFatal());
    }

    return $ret;
}

/**
 * Check user permissions.
 *
 * @return bool
 */
protected function checkPermissions() {
    return $this->getUserType() >= USER_TYPE_ZABBIX_ADMIN;
}

/**
 * Execute action and generate response object.
 */
protected function do Action(): void {
    $data = [
        'hosts_count' => API::Host()->get([
            'countOutput' => true,
            'filter' => [
                'status' => $this->getInput('status')
            ]
        ])
    ];

    $this->setResponse(new CControllerResponseData($data));
}
```

Note:
You can view the full list of available controller classes in Zabbix source code.

Views

View file receives the data from a controller and then prepares the HTML look of it.

Note:
Defining view(s) for a frontend module is optional, unless the module is a widget.
 Dashboard widgets need at least two views: one for edit mode and one for view mode (should be stored in the *views* directory).

It is possible to use pre-defined Zabbix HTML classes (from the */zabbix/ui/include/classes/html*) in the view as well as add new HTML and CSS classes. New classes should be stored in the module's *assets* folder.

Example:

```
...
(new CColHeader(_('Name')))
```

This will add a new column name *Name* and style the top table row as on other Zabbix pages.

Action view

This is a reference file for defining an action view.

```
<?php declare(strict_types = 1);

/**
* @var CView $this
*/

$this->includeJsFile('example.something.view.js.php');

(new CWidget())
    ->setTitle(_('Something view'))
    ->addItem(new CDiv($data['name']))
    ->addItem(new CPartial('module.example.something.reusable', [
        'contacts' => $data['contacts']
    ])
    ->show();
```

Assets

The folder *assets* may contain any files and subfolders that do not belong to other directories. You can use it for:

- JavaScript styles (must be inside *assets/js*);
- CSS styles (must be inside *assets/css*);
- Images;
- Fonts;
- Anything else you need to include.

assets/js

assets/js directory is reserved and should only contain JavaScript files. To be used by the widget, specify these files in the *manifest.json*.

For example:

```
"assets": {
    "js": ["class.widget.js"]
}
```

assets/css

assets/css is reserved and should only contain CSS style files. To be used by the widget, specify these files in the *manifest.json*.

For example:

```
"assets": {
    "css": ["mywidget.css"]
}
```

CSS styles

CSS files may contain a custom attribute `theme` to define different style for a specific frontend themes.

Available themes and their attribute values:

- **Blue** - [theme='blue-theme']
- **Blue (classic)** - [theme='blue-classic-theme']
- **Dark** - [theme='dark-theme']
- **Dark blue** - [theme='dark-blue-theme']
- **Dark (classic)** - [theme='dark-classic-theme']
- **High-contrast light** - [theme='hc-light']
- **High-contrast dark** - [theme='hc-dark']

Example:

```
.widget {
    background-color: red;
}

[theme='dark-theme'] .widget {
    background-color: green;
}
```

Register a new module

This section explains how to add a new module to Zabbix frontend.

Pre-requisites Before proceeding make sure, that:

- The module is located inside the *modules* directory of your Zabbix frontend installation (for example, *zabbix/ui/modules*).
- The module has at least a basic version of manifest.json file.
- You have access to the Administration menu section in Zabbix (requires Super admin user role type).

> **Note:**
> The frontend will not install or even recognize incompatible modules.

Adding a module Open *Administration* > *General* > *Modules* page and click *Scan directory*.



Locate your module in the list and activate it.

To activate a module, click the *Disabled* hyperlink - the module's state will change to *Enabled*.

Click the module name to view additional information about the module, such as author, version, or short description (if defined in the manifest).

Widget preview Widget modules, once added, become immediately visible in the dashboard widget list.

You can open a dashboard, switch to edit mode and add the widget to the dashboard as usual.

When you make some changes to the widget, refresh the dashboard to view how the widget looks with the most recent updates.

Widgets

Widgets are Zabbix frontend modules used for the dashboards. Unless otherwise noted, all module guidelines are also applicable to widgets.

However, a widget is notably different from a module. To build a widget:

- specify the type "widget" in the manifest.json file ("type": "widget");
- include at least two views: one for the widget presentation mode and one for the widget configuration mode (example.widget.view.php and example.widget.edit.php);
- and a controller for widget presentation (WidgetView.php);
- use and extend default widget classes.

Configuration

This page describes classes that can be used to create a widget configuration view with custom configuration fields. The widget configuration view is the part of the widget that allows the user to configure widget parameters for presentation.

Widget

Primary widget class, extends the base class of all dashboard widgets - *CWidget*. Required for overriding the default widget behavior.

The *Widget* class should be located in the root directory of the widget (for example, *zabbix/ui/modules/my_custom_widget*).

Widget.php example

```
<?php

namespace Modules\MyCustomWidget;

use Zabbix\Core\CWidget;

class Widget extends CWidget {

    public const MY_CONSTANT = 0;

    public function getTranslationStrings(): array {
        return [
            'class.widget.js' => [
                'No data' => _('No data')
            ]
        ];
    }
}
```

WidgetForm

The *WidgetForm* class extends the default class *CWidgetForm* and contains a set of *CWidgetField* fields that are required for defining widget configuration storage structure in the database and handling input validation.

The *WidgetForm* class should be located in the *includes* directory. If the class has a different name, the name should be specified in the *widget/form_class* parameter in the *manifest.json* file.

includes/WidgetForm.php example

```
<?php

namespace Modules\MyCustomWidget\Includes;

use Modules\MyCustomWidget\Widget;

use Zabbix\Widgets\{
    CWidgetField,
    CWidgetForm
};

use Zabbix\Widgets\Fields\{
    CWidgetFieldMultiSelectItem,
    CWidgetFieldTextBox,
    CWidgetFieldColor
};


class WidgetForm extends CWidgetForm {

    public const DEFAULT_COLOR_PALETTE = [
        'FF465C', 'B0AF07', '0EC9AC', '524BBC', 'ED1248', 'D1E754', '2AB5FF', '385CC7', 'EC1594', 'BAE37D'
        '6AC8FF', 'EE2B29', '3CA20D', '6F4BBC', '00A1FF', 'F3601B', '1CAE59', '45CFDB', '894BBC', '6D6D6D'
    ];

    public function addFields(): self {
        return $this
            ->addField(
                (new CWidgetFieldMultiSelectItem('itemid', _('Item')))
                    ->setFlags(CWidgetField::FLAG_NOT_EMPTY | CWidgetField::FLAG_LABEL_ASTERISK)
                    ->setMultiple(false)
            )
            ->addField(
                new CWidgetFieldTextBox('description', _('Description'))
            )
```

```
            ->addField(
                (new CWidgetFieldColor('chart_color', _('Color')))->setDefault('FF0000')
            );
    }
}
```

CWidgetFormView

The *CWidgetFormView* class is required for specifying the presentation logic of the fields defined in the *WidgetForm* class, determining their appearance and behavior when rendered in the configuration view.

The *CWidgetFormView* class supports the following methods:

- *addField()* - receives an instance of the *CWidgetFieldView* class as a parameter; each *CWidgetField* class, has a respective *CWidgetFieldView* class for using in the widget configuration view.
- *addFieldset()* - receives an instance of the *CWidgetFieldsGroupView* class that combines fields into a collapsible container.
- *addFieldsGroup()* - receives an instance of *CWidgetFormFieldsetCollapsibleView* that visually (with a border) combines fields into a group.
- *includeJsFile()* - allows adding a JavaScript file to the widget configuration view.
- *addJavaScript()* - allows adding inline JavaScript that will be executed as soon as the widget configuration view is loaded.

The *CWidgetFormView* class should be located in the *views* directory.

views/widget.edit.php example

```
<?php

/**
 * My custom widget form view.
 *
 * @var CView $this
 * @var array $data
 */

use Modules\MyCustomWidget\Includes\WidgetForm;

(new CWidgetFormView($data))
    ->addField(
        (new CWidgetFieldMultiSelectItemView($data['fields']['itemid']))->setPopupParameter('numeric', tru
    )
    ->addFieldset(
        (new CWidgetFormFieldsetCollapsibleView(_('Advanced configuration')))
            ->addField(
                new CWidgetFieldTextBoxView($data['fields']['description'])
            )
            ->addField(
                new CWidgetFieldColorView($data['fields']['chart_color'])
            )
    )
    ->includeJsFile('widget.edit.js.php')
    ->addJavaScript('my_custom_widget_form.init('.json_encode([
        'color_palette' => WidgetForm::DEFAULT_COLOR_PALETTE
    ]).');')
    ->show();
```

JavaScript

A JavaScript class can be used to add dynamic behavior and interactivity to the widget configuration view. For example, you can initialize a color picker, defined in the *CWidgetFormView* class.

The JavaScript class should be loaded with the form, therefore it should be referenced in the *CWidgetFormView* class by using the methods *includeJsFile()* and *addJavaScript()*.

In the example below, a singleton class instance is immediately created and stored under the *window.my_custom_widget_form* name. Thus, opening the form for the second time will re-create the instance.

The JavaScript class should be located in the *views* directory.

views/widget.edit.js.php example

```
<?php

use Modules\MyCustomWidget\Widget;

?>

window.my_custom_widget_form = new class {

    init({color_palette}) {
        colorPalette.setThemeColors(color_palette);

        for (const colorpicker of jQuery('.<?= ZBX_STYLE_COLOR_PICKER ?> input')) {
            jQuery(colorpicker).colorpicker();
        }

        const overlay = overlays_stack.getById('widget_properties');

        for (const event of ['overlay.reload', 'overlay.close']) {
            overlay.$dialogue[0].addEventListener(event, () => { jQuery.colorpicker('hide'); });
        }
    }
};
```

CWidgetField

The *CWidgetField* class is a base class from which all form field classes (*CWidgetFieldCheckBox*, *CWidgetFieldTextArea*, *CWidgetFieldRadioButtonList*, etc.) are inherited. Classes extending *CWidgetField* are responsible for receiving, saving, and validating widget configuration values.

The following *CWidgetField* classes are available.

CWidgetField class	Database field type	Description
CWidgetFieldCheckBox	int32	Single checkbox.
CWidgetFieldCheckBoxList	array of int32	Multiple checkboxes under a single configuration field.
CWidgetFieldColor	string	Color selection field.
CWidgetFieldDatePicker	string	Date selection field.
CWidgetFieldIntegerBox	int32	Field to enter an integer. Can be used to configure minimum and maximum values.
CWidgetFieldLatLng	string	Text box for entering comma-separated latitude, longitude, and map zoom level.
CWidgetFieldMultiSelect	-	Base class for multiselect fields, extended by all `CWidgetFieldMultiSelect*` classes.
CWidgetFieldMultiSelectAction	ID	Multiselect field for selecting actions from the list of actions defined in *Alerts* > *Actions*.
CWidgetFieldMultiSelectGraph	ID	Multiselect field for selecting custom graphs.
CWidgetFieldMultiSelectGraphPrototype	ID	Multiselect field for selecting custom graph prototypes.
CWidgetFieldMultiSelectGroup	ID	Multiselect field for selecting host groups.
CWidgetFieldMultiSelectHost	ID	Multiselect field for selecting hosts.
CWidgetFieldMultiSelectHostInventory	array of int32	Multiselect field for selecting host inventory fields.
CWidgetFieldMultiSelectItem	ID	Multiselect field for selecting items.
CWidgetFieldMultiSelectItemPrototype	ID	Multiselect field for selecting item prototypes.
CWidgetFieldMultiSelectMap	ID	Multiselect field for selecting maps.
CWidgetFieldMultiSelectMediaType	ID	Multiselect field for selecting media types.
CWidgetFieldMultiSelectOverrideHost	ID	Multiselect field for selecting a data source for hosts from a dashboard or another widget.
CWidgetFieldMultiSelectService	ID	Multiselect field for selecting services.
CWidgetFieldMultiSelectSla	ID	Multiselect field for selecting SLAs.
CWidgetFieldMultiSelectUser	ID	Multiselect field for selecting users.
CWidgetFieldNavTree	array of (multiple mixed)	Stores a navigation tree structure of maps for use in the *Map navigation tree* widget.
CWidgetFieldNumericBox	string	Field to enter a float number.
CWidgetFieldPatternSelect	-	Base class for pattern select fields, extended by all `CWidgetPatternSelect*` classes.

CWidgetField class	Database field type	Description
CWidgetFieldPatternSelectHost	string	Multiselect field for selecting hosts. Supports host name patterns.
CWidgetFieldPatternSelectItem	string	Multiselect field for selecting items. Supports item name patterns.
CWidgetFieldRadioButtonList	int32	Radio box group that consists of one or more radio boxes.
CWidgetFieldRangeControl	int32	Slider to select an integer type value.
CWidgetFieldReference	string	Stores a unique identifier for this widget on the dashboard; used for widget communication.
CWidgetFieldSelect	int32	Dropdown select box.
CWidgetFieldSeverities	array of int32	*CWidgetFieldCheckBoxList* preset with trigger severities.
CWidgetFieldSparkline	array of (multiple mixed)	Allows configuring a sparkline chart.
CWidgetFieldTags	array of (string, int32, string)	Allows configuring one or more tag filter rows.
CWidgetFieldTextArea	string	Text area for entering multi-line text.
CWidgetFieldTextBox	string	Text box for entering single-line text.
CWidgetFieldThresholds	array of (string, string)	Allows configuring color and number pairs.
CWidgetFieldTimePeriod	array of string	Time period selecting field.
CWidgetFieldTimeZone	string	Dropdown with timezones.
CWidgetFieldUrl	string	Text box for entering URLs.

Note:
The maximum length for `int32` fields is 10 characters, `ID` fields 20 characters, and `string` fields 65535 characters. For composite types (e.g., `array of (string, int32, string)`), the limit applies to each element individually. Field-specific length restrictions can be set using methods such as `setMaxLength()`.

Presentation

This page describes the components that can be used to create a widget presentation view. The widget presentation view is the part of the widget that receives the data according to its configuration and displays it on the dashboard in a container.

The presentation view consists of three parts:

- Widget action
- Widget view
- JavaScript

Widget action

The widget action class (*WidgetView*) contains methods for operations with widgets in the presentation view mode. The majority of widget actions use and/or extend the default controller class *CControllerDashboardWidgetView*.

The widget action class should be located in the *actions* directory and specified in the *actions* parameter (*actions/widget.{id}.view/class*) in the *manifest.json* file.

actions/WidgetView.php example (implemented in the Zabbix-native *System information* widget)

```
class WidgetView extends CControllerDashboardWidgetView {

    protected function doAction(): void {
        $this->setResponse(new CControllerResponseData([
            'name' => $this->getInput('name', $this->widget->getDefaultName()),
            'system_info' => CSystemInfoHelper::getData(),
            'info_type' => $this->fields_values['info_type'],
            'user_type' => CWebUser::getType(),
            'user' => [
                'debug_mode' => $this->getDebugMode()
            ]
        ]));
    }
}
```

Widget view

The widget view class (*CWidgetView*) is responsible for building the widget presentation view.

The widget view class should be located in the *views* directory. If the file containing the widget view class has a different name than the default (*widget.view.php*), then it must be specified in the *manifest.json* file *actions* parameter (*actions/widget.{id}.view/view*).

views/widget.view.php example

```
<?php

/**
 * My custom widget view.
 *
 * @var CView $this
 * @var array $data
 */

(new CWidgetView($data))
    ->addItem(
        new CTag('h1', true, $data['name'])
    )
    ->show();
```

JavaScript

The JavaScript class is responsible for determining widget behavior, such as updating widget data, resizing the widget, displaying widget elements, etc.

All JavaScript operations use and/or extend the base JavaScript class of all dashboard widgets - *CWidget*. The *CWidget* class contains a set of methods with the default implementation for widget behavior. Depending on widget complexity, these methods can be utilized as is or extended.

The *CWidget* class contains the following methods:

- Methods that define widget lifecycle: *onInitialize()*, *onStart()*, *onActivate()*, *onDeactivate()*, *onDestroy()*, *onEdit()*.
- Methods that handle updating and displaying widget data: *promiseUpdate()*, *getUpdateRequestData()*, *processUpdateResponse(response)*, *processUpdateErrorResponse(error)*, *setContents(response)*.
- Methods that modify widget appearance: *onResize()*, *hasPadding()*.

The JavaScript class should be located in the *assets/js* directory and specified in the *assets* (*assets/js*) parameter in the *manifest.json* file.

Lifecycle methods

The widget lifecycle methods are invoked by the dashboard, and at different stages of the widget's lifecycle during its existence within the dashboard.

The ***onInitialize()*** method defines the initial state and/or values of the widget, without performing any HTML or data manipulation. This method is invoked when a widget is created (a widget object is instantiated), typically by adding the widget to a dashboard page or loading the dashboard page.

Example:

```
onInitialize() {
    this._time_offset = 0;
    this._interval_id = null;
    this._clock_type = CWidgetClock.TYPE_ANALOG;
    this._time_zone = null;
    this._show_seconds = true;
    this._time_format = 0;
    this._tzone_format = 0;
    this._show = [];
    this._has_contents = false;
    this._is_enabled = true;
}
```

The ***onStart()*** method defines the HTML structure of the widget, without performing any data manipulation. This method is invoked before the first activation of the dashboard page, that is, before the dashboard and its widgets are fully displayed to the user.

Example:

```
onStart() {
    this._events.resize = () => {
        const padding = 25;
```

```
        const header_height = this._view_mode === ZBX_WIDGET_VIEW_MODE_HIDDEN_HEADER
            ? 0
            : this._header.offsetHeight;

        this._target.style.setProperty(
            '--content-height',
            `${this._cell_height * this._pos.height - padding * 2 - header_height}px`
        );
    }
}
```

The ***onActivate()*** method makes the widget active and interactive by enabling custom event listeners (for responding to user actions) and initiating the widget update cycle (for keeping its content up-to-date). This method is invoked when the dashboard page is activated, that is, when it becomes fully displayed in the user interface.

Note that before the *onActivate()* method is invoked, the widget is in the inactive state (`WIDGET_STATE_INACTIVE`). After successful invocation, the widget transitions to the active state (`WIDGET_STATE_ACTIVE`). In the active state, the widget is responsive, listens to events, updates its content periodically, and can interact with other widgets.

Example:

```
onActivate() {
    this._startClock();

    this._resize_observer = new ResizeObserver(this._events.resize);
    this._resize_observer.observe(this._target);
}
```

The ***onDeactivate()*** method stops any activity and interactivity of the widget by deactivating custom event listeners and stopping the widget update cycle. This method is invoked when the dashboard page is deactivated, that is, switched away or deleted, or when the widget is deleted from the dashboard page.

Note that before the *onDeactivate()* method is invoked, the widget is in the active state (`WIDGET_STATE_ACTIVE`). After successful invocation, the widget transitions to the inactive state (`WIDGET_STATE_INACTIVE`).

Example:

```
onDeactivate() {
    this._stopClock();
    this._resize_observer.disconnect();
}
```

The ***onDestroy()*** method performs cleanup tasks before the widget is deleted from the dashboard, which can include closing a database connection that was established during widget initialization, cleaning up temporary data to free up system memory and avoid resource leaks, unregistering event listeners related to resize events or button clicks to prevent unnecessary event handling and memory leaks, etc. This method is invoked when the widget or the dashboard page that contains it is deleted.

Note that before the *onDestroy()* method is invoked, a widget in an active state (`WIDGET_STATE_ACTIVE`) is always deactivated with the invocation of the *onDeactivate()* method.

Example:

```
onDestroy() {
    if (this._filter_widget) {
        this._filter_widget.off(CWidgetMap.WIDGET_NAVTREE_EVENT_MARK, this._events.mark);
        this._filter_widget.off(CWidgetMap.WIDGET_NAVTREE_EVENT_SELECT, this._events.select);
    }
}
```

The ***onEdit()*** method defines the appearance and behavior of the widget when the dashboard transitions into edit mode. This method is invoked when the dashboard transitions into edit mode, typically when a user interacts with the widget's *Edit* button or the dashboard's *Edit dashboard* button.

Example:

```
onEdit() {
    this._deactivateGraph();
}
```

Update process methods

The widget update process methods are responsible for retrieving updated data from Zabbix server or any other data source and displaying it in the widget.

The ***promiseUpdate()*** method initiates the data update process by retrieving data, typically using web requests or API calls. This method is invoked when a dashboard page is displayed and periodically after, until the dashboard page is switched to another dashboard page.

The following is an example of the default implementation of the *promiseUpdate()* method used by most Zabbix-native widgets. In the default implementation, the *promiseUpdate()* method follows a general pattern for retrieving data from the server. It creates a new `Curl` object with the appropriate URL and request parameters, sends a `POST` request using the *fetch()* method with the data object constructed by the *getUpdateRequestData()* method, and processes the response (or an error response) with the *processUpdateResponse(response)* or *processUpdateErrorResponse(error)* accordingly. This implementation is suitable for most widgets as they typically retrieve data in a JSON format and handle it in a consistent manner.

```
promiseUpdate() {
    const curl = new Curl('zabbix.php');

    curl.setArgument('action', `widget.${this._type}.view`);

    return fetch(curl.getUrl(), {
        method: 'POST',
        headers: {'Content-Type': 'application/json'},
        body: JSON.stringify(this.getUpdateRequestData()),
        signal: this._update_abort_controller.signal
    })
        .then((response) => response.json())
        .then((response) => {
            if ('error' in response) {
                this.processUpdateErrorResponse(response.error);

                return;
            }

            this.processUpdateResponse(response);
        });
    }
```

The ***getUpdateRequestData()*** method prepares the server request data for updating the widget by gathering various properties and their corresponding values (widget identifiers, filter settings, time ranges, etc.) from the widget's state and configuration, and constructing a data object that represents the necessary information to be sent to the server in the update request. This method is invoked only as part of the default *promiseUpdate()* method, that is, during the widget update process.

Default implementation:

```
getUpdateRequestData() {
    return {
        templateid: this._dashboard.templateid ?? undefined,
        dashboardid: this._dashboard.dashboardid ?? undefined,
        widgetid: this._widgetid ?? undefined,
        name: this._name !== '' ? this._name : undefined,
        fields: Object.keys(this._fields).length > 0 ? this._fields : undefined,
        view_mode: this._view_mode,
        edit_mode: this._is_edit_mode ? 1 : 0,
        dynamic_hostid: this._dashboard.templateid !== null || this.supportsDynamicHosts()
            ? (this._dynamic_hostid ?? undefined)
            : undefined,
        ...this._contents_size
    };
}
```

The ***processUpdateResponse(response)*** method handles the response received from the server after the update request, and, if the update process has been successful and without errors, clears widget data and displays new contents with the *setContents()* method. This method is invoked only as part of the default *promiseUpdate()* method, that is, during the widget update process.

Default implementation:

```
processUpdateResponse(response) {
    this._setHeaderName(response.name);

    this._updateMessages(response.messages);
    this._updateInfo(response.info);
    this._updateDebug(response.debug);

    this.setContents(response);
}
```

The ***processUpdateErrorResponse(error)*** method handles the response received from the server after the update request if the response is an error, and displays the error message/-s. This method is invoked only as part of the default *promiseUpdate()* method, that is, during the widget update process.

Default implementation:

```
processUpdateErrorResponse(error) {
    this._updateMessages(error.messages, error.title);
}
```

The ***setContents(response)*** method displays widget contents if the widget update process has been successful and without errors, which can include manipulating DOM elements, updating UI components, applying styles or formatting, etc. This method is invoked only as part of the default *processUpdateResponse(response)* method, that is, during the process of handling the response received from the server after the update request.

Default implementation:

```
setContents(response) {
    this._body.innerHTML = response.body ?? '';
}
```

Presentation modification methods

The widget presentation modification methods are responsible for modifying widget appearance.

The ***onResize()*** method is responsible for adjusting widget's visual elements to accommodate the new widget size, which can include rearranging elements, adjusting element dimensions, text truncation, implementing lazy loading to improve responsiveness during resizing, etc. This method is invoked when the widget is resized, for example, when the user manually resizes the widget or when the browser window is resized.

Example:

```
onResize() {
    if (this.getState() === WIDGET_STATE_ACTIVE) {
        this._startUpdating();
    }
}
```

The ***hasPadding()*** method is responsible for applying an 8px vertical padding at the bottom of the widget when it is configured to show its header. This method is invoked when the dashboard page is activated, that is, when it becomes the displayed page in the user interface.

Default implementation:

```
hasPadding() {
    return this.getViewMode() !== ZBX_WIDGET_VIEW_MODE_HIDDEN_HEADER;
}
```

For some widgets it is necessary to use all of the available widget space to configure, for example, a custom background color. The following is an example of the implementation of the *hasPadding()* method used in the Zabbix-native *Item value* widget.

```
hasPadding() {
    return false;
}
```

Tutorials

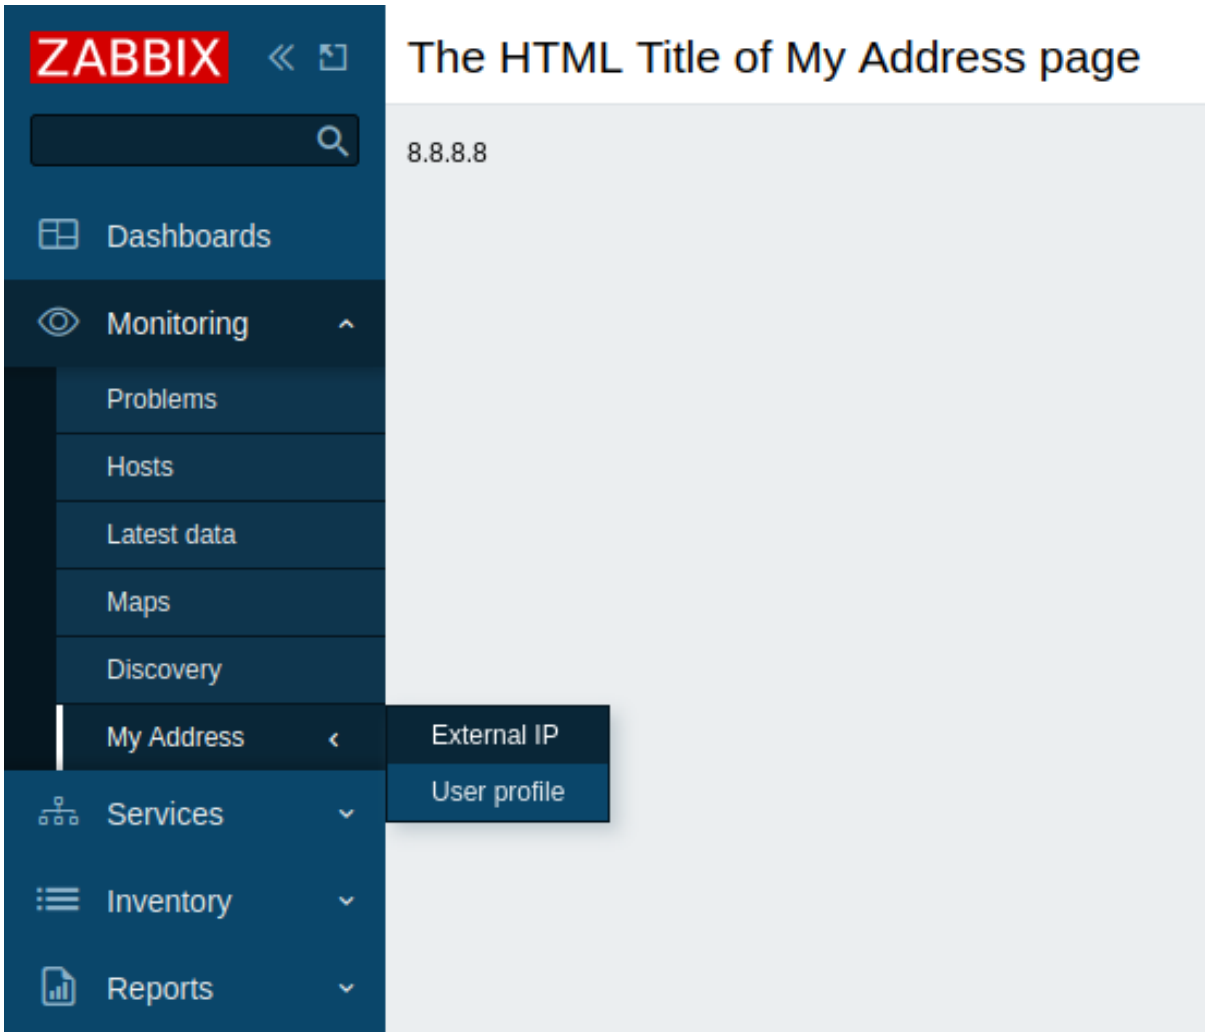
This section contains practical step-by-step tutorials to illustrate how to build a custom module and a widget in Zabbix.

Create a module (tutorial)

This is a step-by-step tutorial that shows how to create a simple Zabbix frontend module. You can download all files of this module as a ZIP archive: MyAddress.zip.

What you'll build

During this tutorial, you will first build a frontend module that adds a new *My Address* menu section and then convert it into a more advanced frontend module that makes an HTTP request to *https://api.seeip.org* and displays the response - the IP address of your computer - on a new page in the newly created *My Address* menu section. Here's how the finished module will look like:



Part I - New menu section

Add a blank module to Zabbix frontend

1. Create a directory *MyAddress* in the *modules* directory of your Zabbix frontend installation (for example, *zabbix/ui/modules*).
2. Create a *manifest.json* file with basic module metadata (see the description of supported parameters).

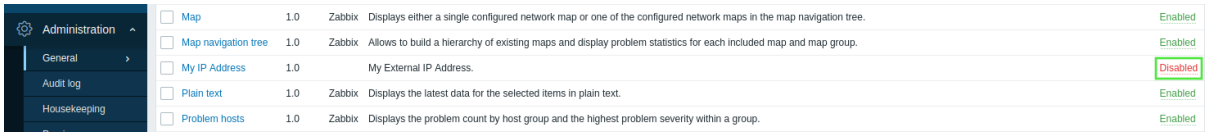
ui/modules/MyAddress/manifest.json

```
{
    "manifest_version": 2.0,
    "id": "my-address",
    "name": "My IP Address",
    "version": "1.0",
    "namespace": "MyAddress",
    "description": "My External IP Address."
}
```

3. In Zabbix frontend, go to *Administration* > *General* > *Modules* section and click the *Scan directory* button.



4. Locate the new module *My IP Address* in the list and click the "Disabled" hyperlink to change the module's status from "Disabled" to "Enabled" (if the module is not listed, see troubleshooting section).



The module is now registered in the frontend. However, it is not visible yet, because you still need to define the module functionality. Once you add content to the module directory, you will immediately see the changes in Zabbix frontend upon refreshing the page.

Create a menu section

1. Create a *Module.php* file in the *MyAddress* directory.

This file implements a new *Module* class that extends the default *CModule* class. The *Module* class will insert a new *My Address* menu section into the main menu.

The *setAction()* method specifies an action to be executed upon clicking the menu section. To start with, you can use the predefined action *userprofile.edit*, which will open the *User profile* page. In part III of this tutorial, you will learn how to create a custom action.

ui/modules/MyAddress/Module.php

```
<?php

namespace Modules\MyAddress;

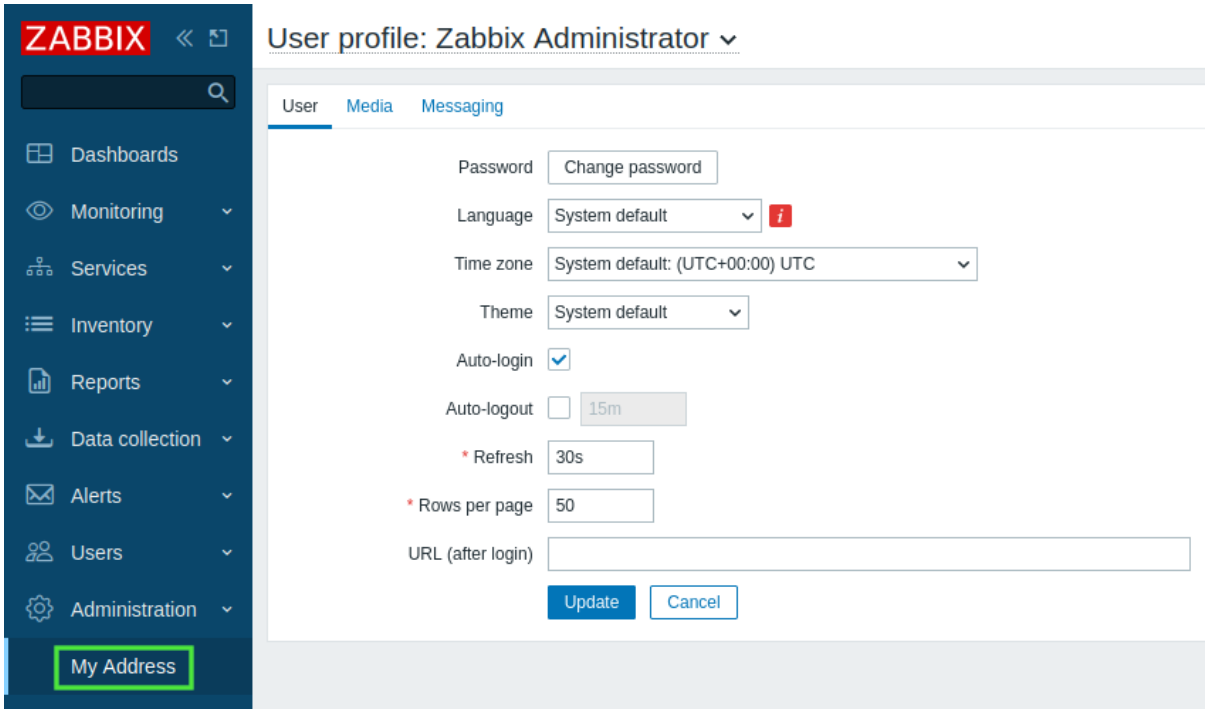
use Zabbix\Core\CModule,
    APP,
    CMenuItem;

class Module extends CModule {

    public function init(): void {
        APP::Component()->get('menu.main')
            ->add((new CMenuItem(_('My Address')))
            ->setAction('userprofile.edit'));
    }
}
```

Note:
You can replace `'userprofile.edit'` with other actions, for example, `'charts.view'` (opens custom graphs), `'problems.view'` (opens *Monitoring* > *Problems*), or `'report.status'` (opens *System information* report).

3. Refresh Zabbix frontend. There is now a new *My Address* section at the bottom of the Zabbix main menu. Click *My Address* to open the *User profile* page.



Part II - Menu section location change

In this part, you will move the *My Address* menu section to the *Monitoring* section and then add a nested menu to it. As a result, users will be able to access two sub-menu pages from the *Monitoring* > *My Address* menu section.

1. Open and edit the *Module.php* file.

ui/modules/MyAddress/Module.php

```
<?php

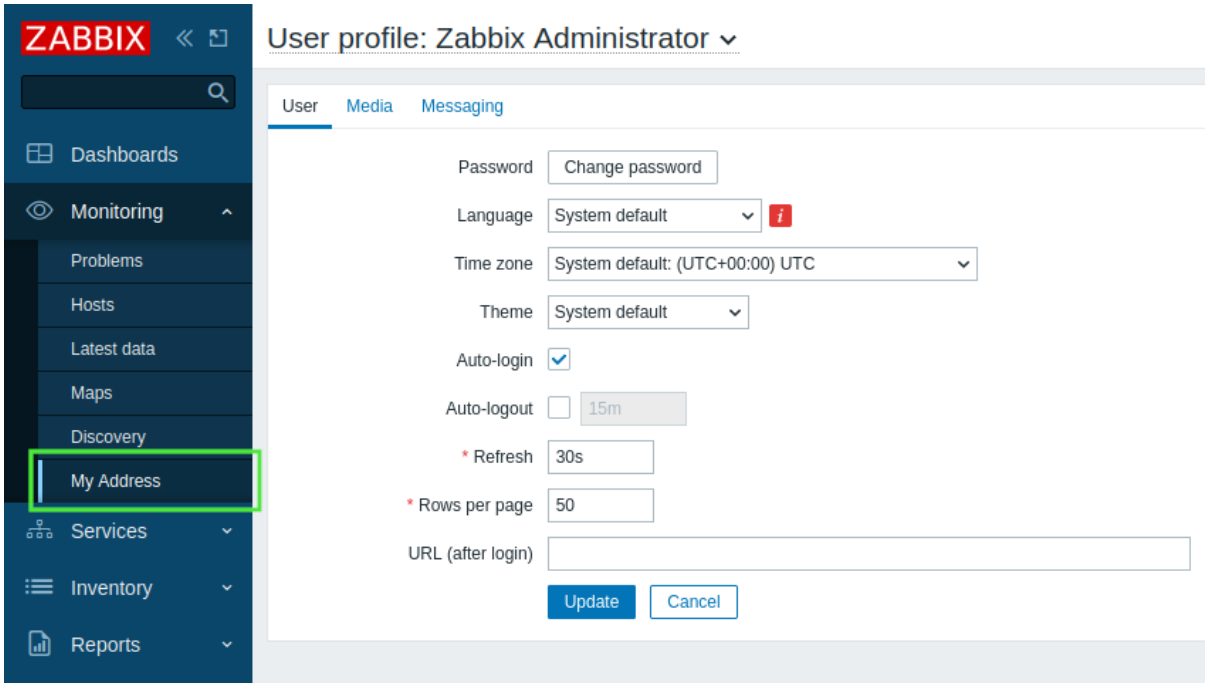
namespace Modules\MyAddress;

use Zabbix\Core\CModule,
    APP,
    CMenuItem;

class Module extends CModule {

    public function init(): void {
        APP::Component()->get('menu.main')
            ->findOrAdd(_('Monitoring'))
            ->getSubmenu()
            ->insertAfter(_('Discovery'),
                (new CMenuItem(_('My Address')))->setAction('userprofile.edit')
            );
    }
}
```

2. Refresh Zabbix frontend. Expand the *Monitoring* menu section and observe that the *My address* section is now located below the *Discovery* section.



3. To add nested pages to the *My Address* menu section, open and edit the *Module.php* file again.

This step creates two subsections:

- *External IP* that executes a new 'my.address' action that will be defined in the next steps;
- *User profile* that executes the predefined 'userprofile.edit' action to open the *User profile* page.

Note that for the nested menu, you need to use the *CMenu* class in addition to the classes used in the previous steps.

ui/modules/MyAddress/Module.php

```
<?php

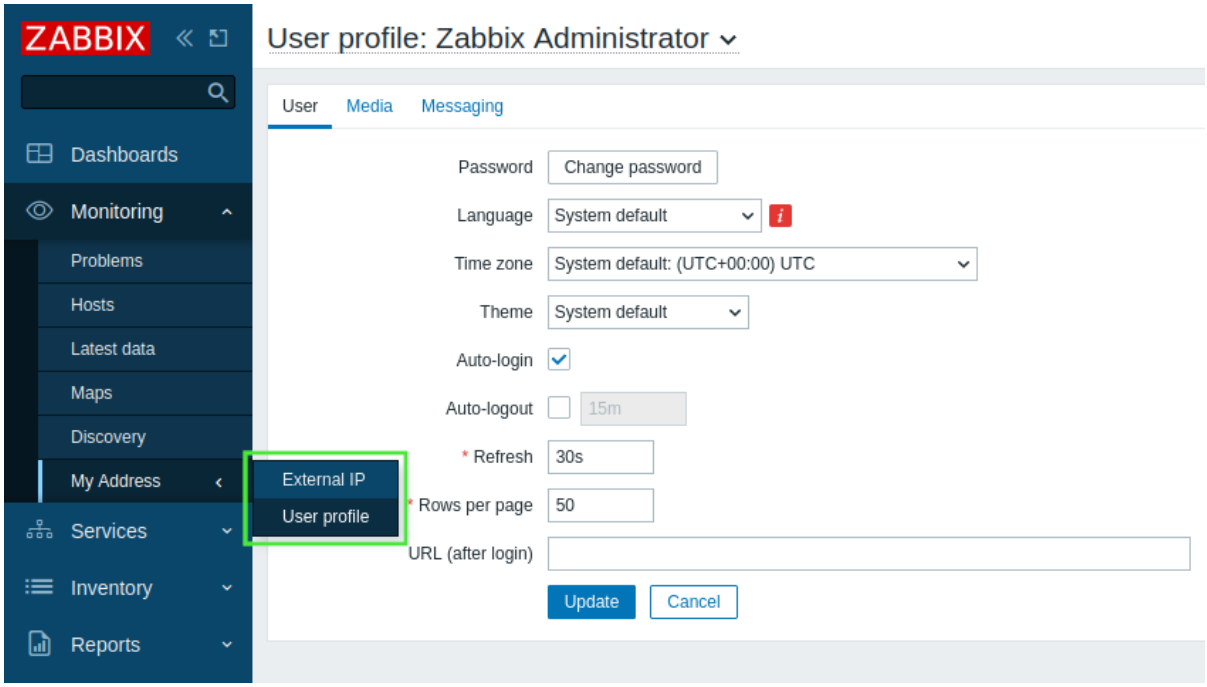
namespace Modules\MyAddress;

use Zabbix\Core\CModule,
    APP,
    CMenu,
    CMenuItem;

class Module extends CModule {

    public function init(): void {
        APP::Component()->get('menu.main')
            ->findOrAdd(_('Monitoring'))
            ->getSubmenu()
            ->insertAfter(_('Discovery'),
                (new CMenuItem(_('My Address')))->setSubMenu(
                    new CMenu([
                        (new CMenuItem(_('External IP')))->setAction('my.address'),
                        (new CMenuItem(_('User profile')))->setAction('userprofile.edit')
                    ])
                )
            );
    }
}
```

4. Refresh Zabbix frontend. Observe that the *My address* menu section now contains a third-level menu with two pages - *External IP* and *User profile*.



Part III - Module action

An action is implemented in two files - *actions/MyAddress.php* and *views/my.address.php*. The ***actions/MyAddress.php*** file takes care of the business logic implementation, while the ***views/my.address.php*** file is responsible for the view.

1. Create a directory *actions* in the *MyAddress* directory.
2. Create a *MyAddress.php* file in the *actions* directory.

The action logic will be defined in the *MyAddress* class. This action class will implement four functions: *init()*, *checkInput()*, *checkPermissions()*, and *doAction()*. Zabbix frontend calls the *doAction()* function when the action is requested. This function is responsible for the business logic of the module.

Attention:
The data must be organized as an associative array. The array can be multidimensional and may contain any data expected by the view.

ui/modules/MyAddress/actions/MyAddress.php

```
<?php

namespace Modules\MyAddress\Actions;

use CController,
    CControllerResponseData;

class MyAddress extends CController {

    public function init(): void {
        $this->disableCsrfValidation();
    }

    protected function checkInput(): bool {
        return true;
    }

    protected function checkPermissions(): bool {
        return true;
    }

    protected function doAction(): void {
        $data = ['my-ip' => file_get_contents("https://api.seeip.org")];
        $response = new CControllerResponseData($data);
```

```
        $this->setResponse($response);
    }
}
```

3. Create a new directory *views* in the *MyAddress* directory.
4. Create a *my.address.php* file in the *views* directory and define the module view.

Note that the variable `$data` is available in the view without specifically defining it. The framework automatically passes the associative array to the view.

ui/modules/MyAddress/views/my.address.php

```
<?php

(new CHtmlPage())
    ->setTitle(_('The HTML Title of My Address Page'))
    ->addItem(new CDiv($data['my-ip']))
    ->show();
```

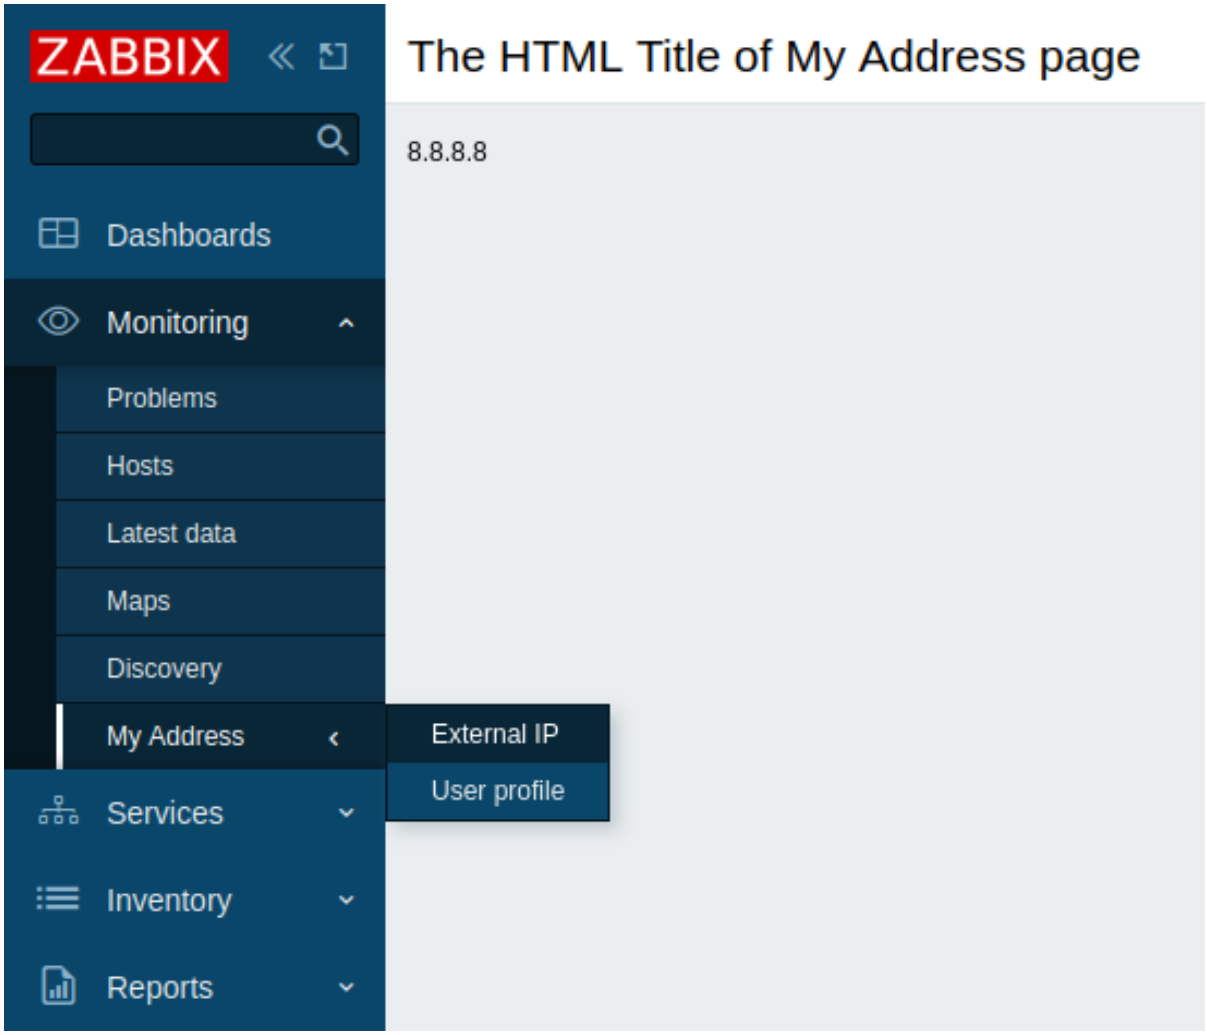
5. The module action has to be registered in the *manifest.json* file. Open *manifest.json* and add a new object `actions` that contains:

- the action key with the action name written in lowercase (a-z) and with words separated by dots (for example, `my.address`);
- the action class name (`MyAddress`) as a value for the `class` key of the `my.address` object;
- the action view name (`my.address`) as a value for the `view` key of the `my.address` object.

ui/modules/MyAddress/manifest.json

```
{
    "manifest_version": 2.0,
    "id": "my-address",
    "name": "My IP Address",
    "version": "1.0",
    "namespace": "MyAddress",
    "description": "My External IP Address.",
    "actions": {
        "my.address": {
            "class": "MyAddress",
            "view": "my.address"
        }
    }
}
```

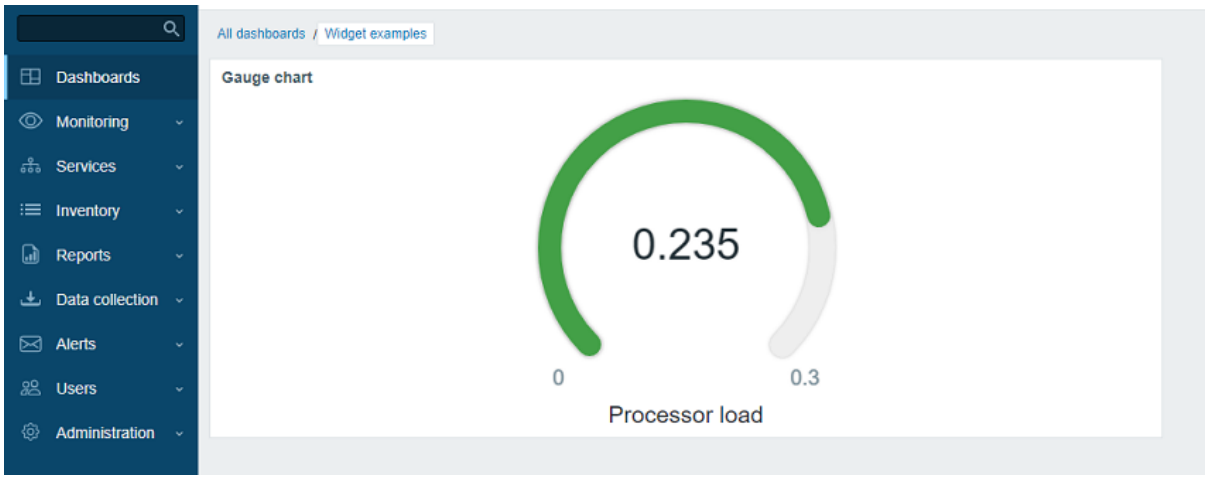
6. Refresh Zabbix frontend. Click *My address* > *External IP* to see the IP address of your computer.



Create a widget (tutorial)

This is a step-by-step tutorial that shows how to create a simple dashboard widget. You can download all files of this widget as a ZIP archive: lesson_gauge_chart.zip.

What you'll build During this tutorial, you will first build a basic "Hello, world!" widget and then convert it into a more advanced widget that displays an item value as a gauge chart. Here's how the finished widget will look like:



Part I - "Hello, world!" In this section you will learn how to create the minimum required widget elements and add a new widget to Zabbix frontend.

Add a blank widget to Zabbix frontend

1. Create a directory `lesson_gauge_chart` in the `modules` directory of your Zabbix frontend installation (for example, *zabbix/ui/modules*).

> **Note:**
> All custom widgets are treated as external modules and must be added to the `modules` directory of your Zabbix frontend installation (for example, *zabbix/ui/modules*). The directory `zabbix/ui/widgets` is reserved for Zabbix built-in widgets and gets updated along with Zabbix UI.

2. Create a `manifest.json` file with basic widget metadata (see the description of supported parameters).

ui/modules/lesson_gauge_chart/manifest.json

```
{
    "manifest_version": 2.0,
    "id": "lesson_gauge_chart",
    "type": "widget",
    "name": "Gauge chart",
    "namespace": "LessonGaugeChart",
    "version": "1.1",
    "author": "Zabbix"
}
```

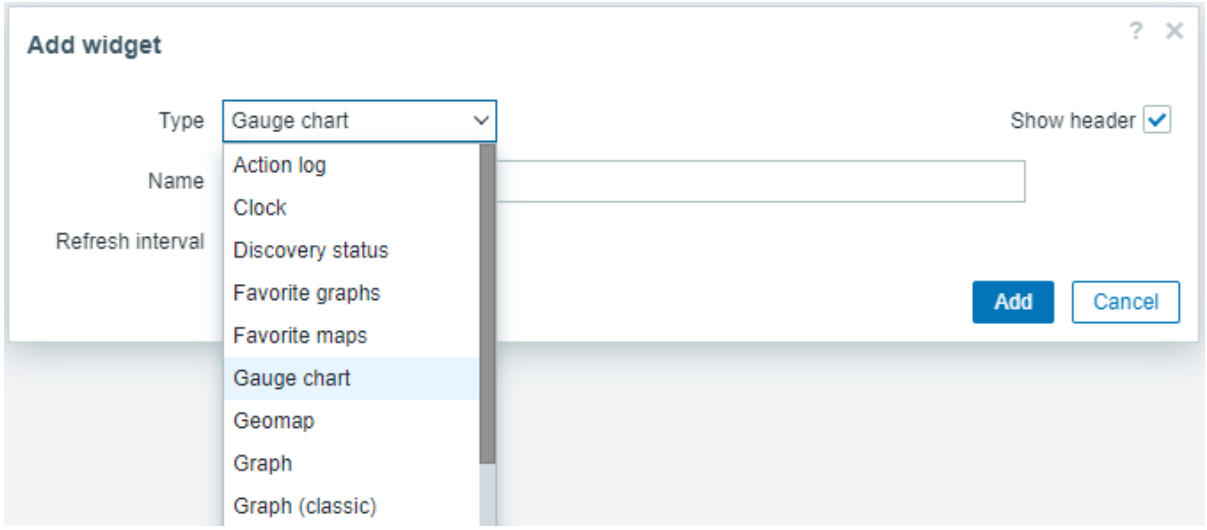
3. In Zabbix frontend, go to *Administration* > *General* > *Modules* section and click the *Scan directory* button.



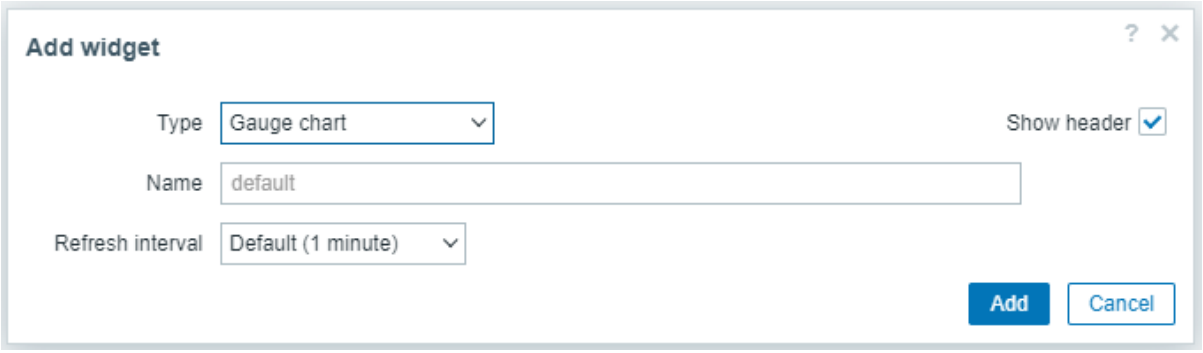
4. Locate the new module *Gauge chart* in the list and click the "Disabled" hyperlink to change the module's status from "Disabled" to "Enabled" (if the module is not listed, see troubleshooting section).



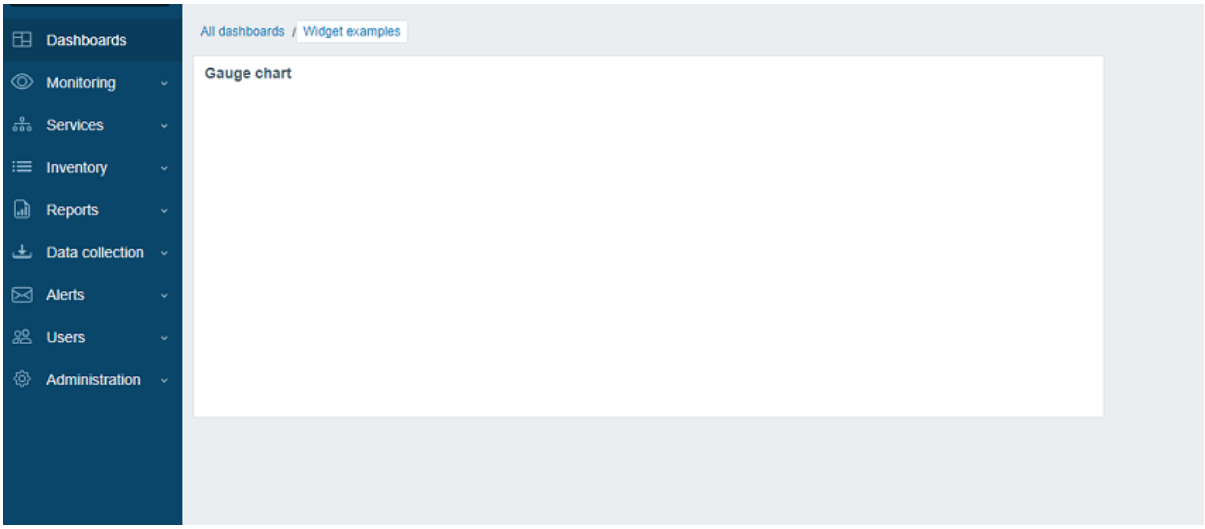
5. Open a dashboard, switch it to edit mode and add a new widget. In the *Type* field, select "Gauge chart".



6. At this point, the *Gauge chart* widget configuration contains only common widget fields *Name* and *Refresh interval*. Click *Add* to add the widget to the dashboard.



7. A blank widget should appear on the dashboard. Click *Save changes* in the upper-right corner to save the dashboard.



Add a widget view

Note:
The widget's **view** file should be located in the `views` directory (for this tutorial, `ui/modules/lesson_gauge_chart/views/`). If the file has the default name `widget.view.php`, you do not need to register it in the `manifest.json` file. If the file has a different name, specify it in the `actions/widget.lesson_gauge_chart.view` section of the `manifest.json` file.

1. Create a directory `views` in the `lesson_gauge_chart` directory.
2. Create a `widget.view.php` file in the `views` directory.

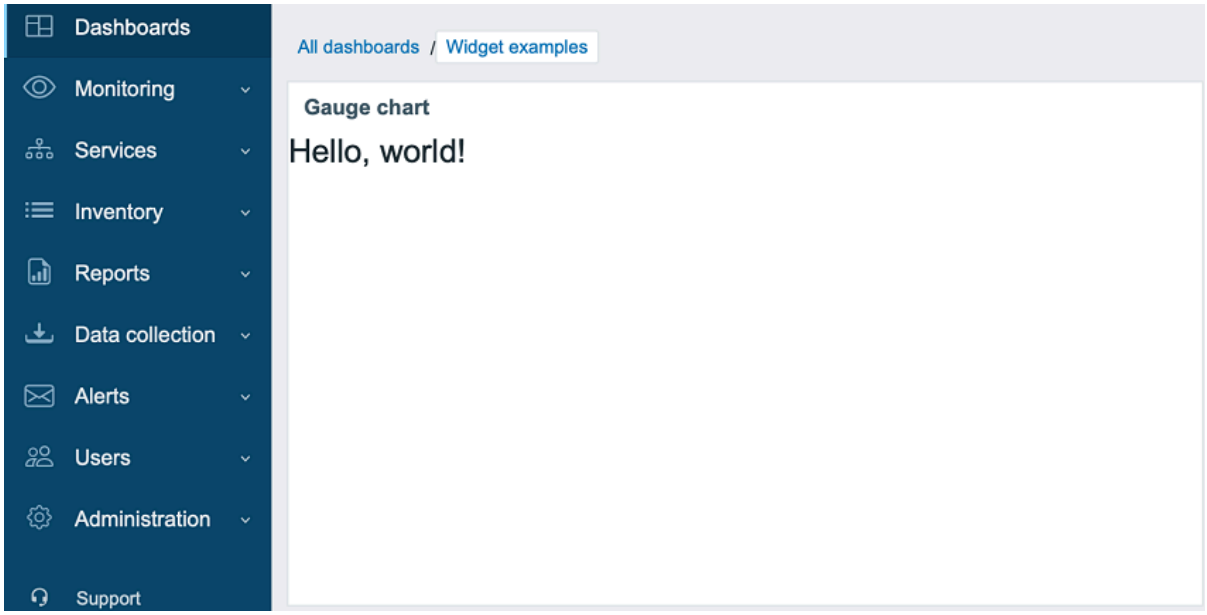
`ui/modules/lesson_gauge_chart/views/widget.view.php`

```
<?php

/**
 * Gauge chart widget view.
 *
 * @var CView $this
 * @var array $data
 */

(new CWidgetView($data))
    ->addItem(
        new CTag('h1', true, 'Hello, world!')
    )
    ->show();
```

3. Refresh the dashboard. The *Gauge chart* widget now displays "Hello, world!".



Part II - Gauge chart Add settings to a configuration view and use them in a widget view

In this section, you will learn how to add a widget configuration field and show the entered value in the widget view as text.

The widget configuration consists of a form (`Zabbix\Widgets\CWidgetForm`) and a widget form view (`widget.edit.php`). To add fields (`Zabbix\Widgets\CWidgetField`), you need to create a `WidgetForm` class, which will extend `Zabbix\Widgets\CWidgetForm`.

The form contains the set of fields (`Zabbix\Widgets\CWidgetField`) of various types, which are used to validate user-entered values. The form field (`Zabbix\Widgets\CWidgetField`) for each input element type converts the value into a single format to store it in the database.

> **Note:**
> The widget's **form** file should be located in the `includes` directory (for this tutorial, `ui/modules/lesson_gauge_chart/includes/`). If the file has the default name `WidgetForm.php`, you do not need to register it in the `manifest.json` file. If the file has a different name, specify it in the `widget/form_class` section of the `manifest.json` file.

1. Create a new directory `includes` in the `lesson_gauge_chart` directory.
2. Create a `WidgetForm.php` file in the `includes` directory.

`ui/modules/lesson_gauge_chart/includes/WidgetForm.php`

```
<?php

namespace Modules\LessonGaugeChart\Includes;

use Zabbix\Widgets\CWidgetForm;

class WidgetForm extends CWidgetForm {
}
```

3. Add a *Description* field to widget configuration form. This is a regular text field, where a user can enter any character set. You can use the `CWidgetFieldTextBox` class for it.

`ui/modules/lesson_gauge_chart/includes/WidgetForm.php`

```
<?php

namespace Modules\LessonGaugeChart\Includes;

use Zabbix\Widgets\CWidgetForm;

use Zabbix\Widgets\Fields\CWidgetFieldTextBox;

class WidgetForm extends CWidgetForm {
```

```
    public function addFields(): self {
        return $this
            ->addField(
               new CWidgetFieldTextBox('description', _('Description'))
            );
    }
}
```

4. In the `views` directory, create a widget configuration view file `widget.edit.php` and add a view for the new *Description* field. For the `CWidgetFieldTextBox` field class, the view is `CWidgetFieldTextBoxView`.

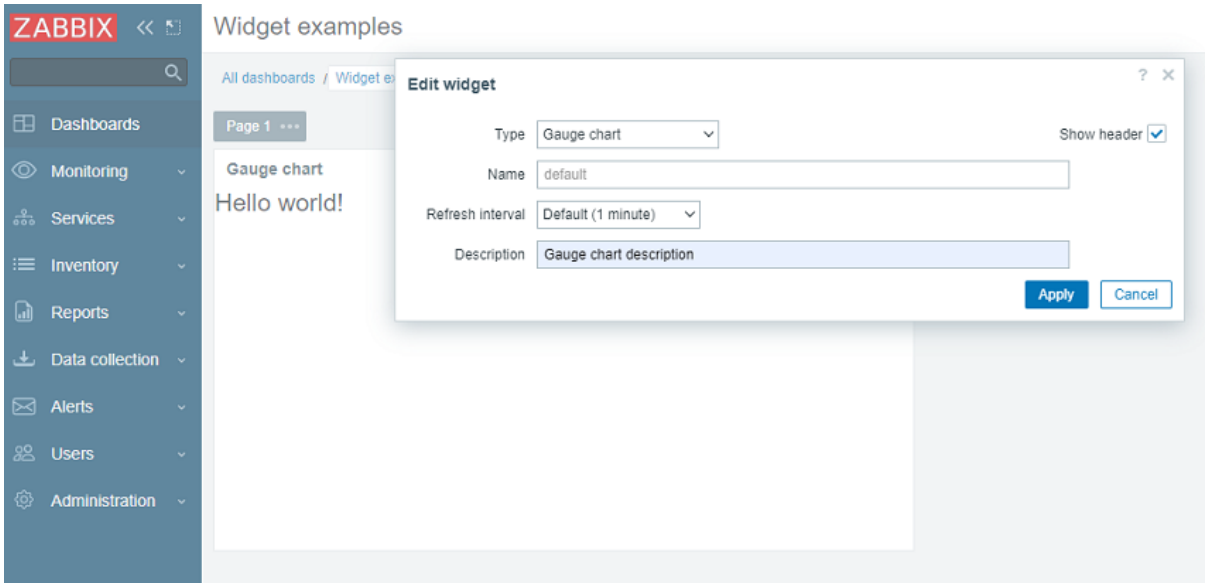
`ui/modules/lesson_gauge_chart/views/widget.edit.php`

```
<?php

/**
 * Gauge chart widget form view.
 *
 * @var CView $this
 * @var array $data
 */

(new CWidgetFormView($data))
    ->addField(
        new CWidgetFieldTextBoxView($data['fields']['description'])
    )
    ->show();
```

5. Go to the dashboard and click the gear icon in the widget to open the widget configuration form.
6. The widget configuration form now contains a new *Description* text field. Enter any value, for example, *Gauge chart description*.



7. Click *Apply* in the widget configuration form. Then click *Save changes* in the upper-right corner to save the dashboard. Note that the new description is not visible anywhere, and the widget still displays "Hello, world!".

For the new description to appear in the widget, the *Description* field value needs to be retrieved from the database and passed to the widget view. For that, you need to create an action class.

8. Create a new directory `actions` in the `lesson_gauge_chart` directory.
9. Create a `WidgetView.php` file in the `actions` directory. The `WidgetView` action class will extend the `CControllerDashboardWidge` class.

Values of the widget configuration fields are stored in the `$fields_values` property of the action class.

`ui/modules/lesson_gauge_chart/actions/WidgetView.php`

```php
<?php

namespace Modules\LessonGaugeChart\Actions;

use CControllerDashboardWidgetView,
    CControllerResponseData;

class WidgetView extends CControllerDashboardWidgetView {

    protected function doAction(): void {
        $this->setResponse(new CControllerResponseData([
            'name' => $this->getInput('name', $this->widget->getName()),
            'description' => $this->fields_values['description'],
            'user' => [
                'debug_mode' => $this->getDebugMode()
            ]
        ]));
    }
}
```

10. Open `manifest.json` and register `WidgetView` as an action class in the `actions/widget.lesson_gauge_chart.view` section.

`ui/modules/lesson_gauge_chart/manifest.json`

```
{
    "manifest_version": 2.0,
    "id": "lesson_gauge_chart",
    "type": "widget",
    "name": "Gauge chart",
    "namespace": "LessonGaugeChart",
    "version": "1.0",
    "author": "Zabbix",
    "actions": {
        "widget.lesson_gauge_chart.view": {
            "class": "WidgetView"
        }
    }
}
```

11. Now you can use the value of the description field, contained in `$data['description']`, in the widget view. Open `views/widget.view.php` and replace the static text "Hello, world!" with `$data['description']`.

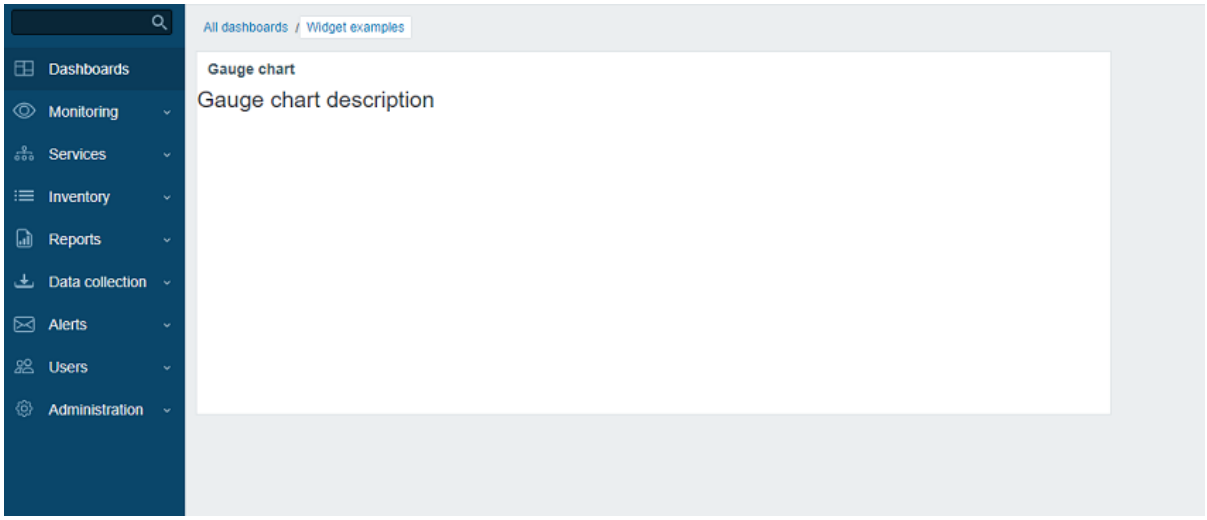
`ui/modules/lesson_gauge_chart/views/widget.view.php`

```php
<?php

/**
 * Gauge chart widget view.
 *
 * @var CView $this
 * @var array $data
 */

(new CWidgetView($data))
    ->addItem(
        new CTag('h1', true, $data['description'])
    )
    ->show();
```

12. Refresh the dashboard page. You should now see the widget description text instead of "Hello, world!".



Retrieve an item value via API

The widget should show the last value of an item of user's choice. For that, you need to add the ability to select items in the widget configuration.

In this section, you will learn how to add an item selection field to the widget form and how to add the visual part of this field to the configuration view. Then, the widget controller will be able to retrieve item data and its value via an API request. Once received, the value can be displayed in the widget view.

1. Open `includes/WidgetForm.php` and add the `CWidgetFieldMultiSelectItem` field. This will allow selecting an item in the configuration form.

`ui/modules/lesson_gauge_chart/includes/WidgetForm.php`

```
<?php

namespace Modules\LessonGaugeChart\Includes;

use Zabbix\Widgets\{
    CWidgetField,
    CWidgetForm
};

use Zabbix\Widgets\Fields\{
    CWidgetFieldMultiSelectItem,
    CWidgetFieldTextBox
};

/**
 * Gauge chart widget form.
 */
class WidgetForm extends CWidgetForm {

    public function addFields(): self {
        return $this
            ->addField(
                (new CWidgetFieldMultiSelectItem('itemid', _('Item')))
                    ->setFlags(CWidgetField::FLAG_NOT_EMPTY | CWidgetField::FLAG_LABEL_ASTERISK)
                    ->setMultiple(false)
            )
            ->addField(
                new CWidgetFieldTextBox('description', _('Description'))
            );
    }
}
```

2. Open `views/widget.edit.php` and add the field visual component to the configuration view.

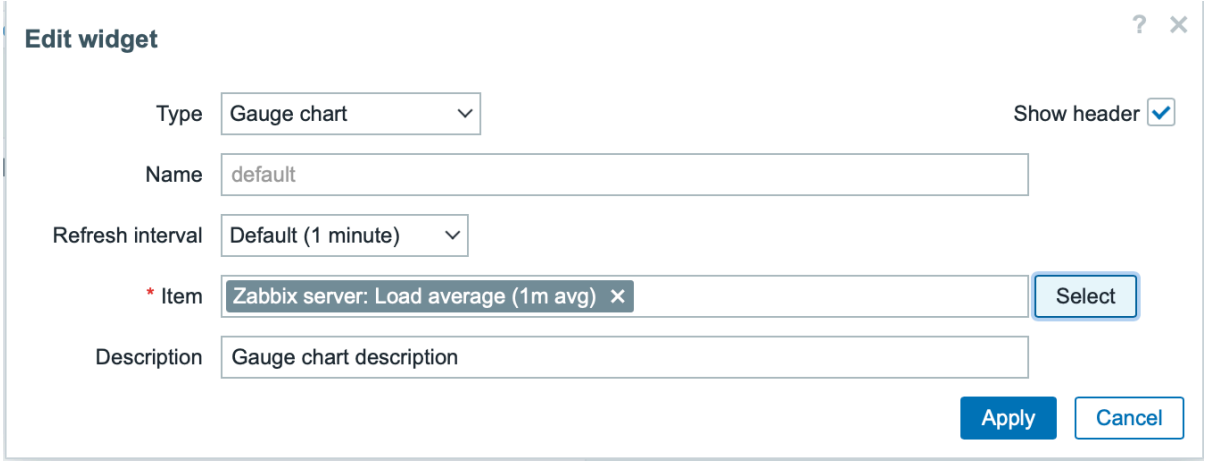
`ui/modules/lesson_gauge_chart/views/widget.edit.php`

```php
<?php

/**
 * Gauge chart widget form view.
 *
 * @var CView $this
 * @var array $data
 */

(new CWidgetFormView($data))
    ->addField(
        new CWidgetFieldMultiSelectItemView($data['fields']['itemid'])
    )
    ->addField(
        new CWidgetFieldTextBoxView($data['fields']['description'])
    )
    ->show();
```

3. Return to the dashboard and click the gear icon in the widget to open the widget configuration form.
4. The widget configuration form now contains a new input field *Item*. Select the host "Zabbix server" and the item "Load average (1m avg)".



5. Click *Apply* in the widget configuration form. Then click *Save changes* in the upper-right corner to save the dashboard.
6. Open and modify `actions/WidgetView.php`.

From now on, the item ID will be available in the widget controller in `$this->fields_values['itemid']`. The `doAction()` controller method collects the item data (name, value type, units) using the API method `item.get` and the item last value using the API method `history.get`.

ui/modules/lesson_gauge_chart/actions/WidgetView.php

```php
<?php

namespace Modules\LessonGaugeChart\Actions;

use API,
    CControllerDashboardWidgetView,
    CControllerResponseData;

class WidgetView extends CControllerDashboardWidgetView {

    protected function doAction(): void {
        $db_items = API::Item()->get([
            'output' => ['itemid', 'value_type', 'name', 'units'],
            'itemids' => $this->fields_values['itemid'],
            'webitems' => true,
            'filter' => [
                'value_type' => [ITEM_VALUE_TYPE_UINT64, ITEM_VALUE_TYPE_FLOAT]
            ]
```

```
        ]);

        $value = null;

        if ($db_items) {
            $item = $db_items[0];

            $history = API::History()->get([
                'output' => API_OUTPUT_EXTEND,
                'itemids' => $item['itemid'],
                'history' => $item['value_type'],
                'sortfield' => 'clock',
                'sortorder' => ZBX_SORT_DOWN,
                'limit' => 1
            ]);

            if ($history) {
                $value = convertUnitsRaw([
                    'value' => $history[0]['value'],
                    'units' => $item['units']
                ]);
            }
        }

        $this->setResponse(new CControllerResponseData([
            'name' => $this->getInput('name', $this->widget->getName()),
            'value' => $value,
            'description' => $this->fields_values['description'],
            'user' => [
                'debug_mode' => $this->getDebugMode()
            ]
        ]));
    }
}
```

7. Open `views/widget.view.php` and add the item value to the widget view.

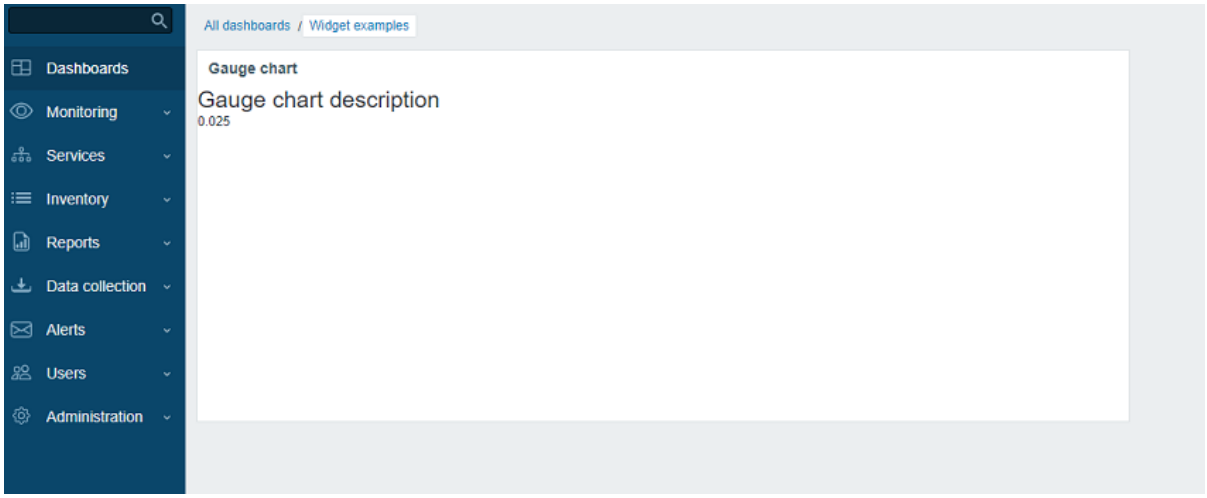
`ui/modules/lesson_gauge_chart/views/widget.view.php`

```
<?php

/**
 * Gauge chart widget view.
 *
 * @var CView $this
 * @var array $data
 */

(new CWidgetView($data))
    ->addItem([
        new CTag('h1', true, $data['description']),
        new CDiv($data['value'] !== null ? $data['value']['value'] : _('No data'))
    ])
    ->show();
```

8. Refresh the dashboard page. The widget will display the latest item value.



Add advanced configuration settings to a configuration view

In this section, you will learn how to add an expandable/collapsible *Advanced configuration* section with optional parameters, such as color, minimum and maximum values, units and the *Description* field created earlier.

1. Create a `Widget.php` file in the main widget directory `lesson_gauge_chart` to create a new class `Widget`.

The `Widget` class will extend the `CWidget` base class to add/override the default widget settings (in this case - translations). JavaScript, provided below, displays the string "No data" in case of missing data. The "No data" string is present in the Zabbix UI translation files.

If there are any widget constants, it is recommended to also specify them in the `Widget` class.

`ui/modules/lesson_gauge_chart/Widget.php`

```
<?php

namespace Modules\LessonGaugeChart;

use Zabbix\Core\CWidget;

class Widget extends CWidget {

    public const UNIT_AUTO = 0;
    public const UNIT_STATIC = 1;

    public function getTranslationStrings(): array {
        return [
            'class.widget.js' => [
                'No data' => _('No data')
            ]
        ];
    }
}
```

2. Open `includes/WidgetForm.php` and add the new fields *Color* (color picker), *Min* (numeric field), *Max* (numeric field), and *Units* (select), and define the default color palette for the color picker, so that it can be used in the next steps.

`ui/modules/lesson_gauge_chart/includes/WidgetForm.php`

```
<?php

namespace Modules\LessonGaugeChart\Includes;

use Modules\LessonGaugeChart\Widget;

use Zabbix\Widgets\{
    CWidgetField,
    CWidgetForm
};
```

```
use Zabbix\Widgets\Fields\{
    CWidgetFieldColor,
    CWidgetFieldMultiSelectItem,
    CWidgetFieldNumericBox,
    CWidgetFieldSelect,
    CWidgetFieldTextBox
};

/**
 * Gauge chart widget form.
 */
class WidgetForm extends CWidgetForm {

    public const DEFAULT_COLOR_PALETTE = [
        'FF465C', 'B0AF07', '0EC9AC', '524BBC', 'ED1248', 'D1E754', '2AB5FF', '385CC7', 'EC1594', 'BAE37D',
        '6AC8FF', 'EE2B29', '3CA20D', '6F4BBC', '00A1FF', 'F3601B', '1CAE59', '45CFDB', '894BBC', '6D6D6D'
    ];

    public function addFields(): self {
        return $this
            ->addField(
                (new CWidgetFieldMultiSelectItem('itemid', _('Item')))
                    ->setFlags(CWidgetField::FLAG_NOT_EMPTY | CWidgetField::FLAG_LABEL_ASTERISK)
                    ->setMultiple(false)
            )
            ->addField(
                (new CWidgetFieldColor('chart_color', _('Color')))->setDefault('FF0000')
            )
            ->addField(
                (new CWidgetFieldNumericBox('value_min', _('Min')))
                    ->setDefault(0)
                    ->setFlags(CWidgetField::FLAG_NOT_EMPTY | CWidgetField::FLAG_LABEL_ASTERISK)
            )
            ->addField(
                (new CWidgetFieldNumericBox('value_max', _('Max')))
                    ->setDefault(100)
                    ->setFlags(CWidgetField::FLAG_NOT_EMPTY | CWidgetField::FLAG_LABEL_ASTERISK)
            )
            ->addField(
                (new CWidgetFieldSelect('value_units', _('Units'), [
                    Widget::UNIT_AUTO => _x('Auto', 'history source selection method'),
                    Widget::UNIT_STATIC => _x('Static', 'history source selection method')
                ]))->setDefault(Widget::UNIT_AUTO)
            )
            ->addField(
                (new CWidgetFieldTextBox('value_static_units'))
            )
            ->addField(
                new CWidgetFieldTextBox('description', _('Description'))
            );
    }
}
```

3. Open `views/widget.edit.php` and add the field visual components to the configuration view.

`ui/modules/lesson_gauge_chart/views/widget.edit.php`

```
<?php

/**
 * Gauge chart widget form view.
 *
 * @var CView $this
 * @var array $data
```

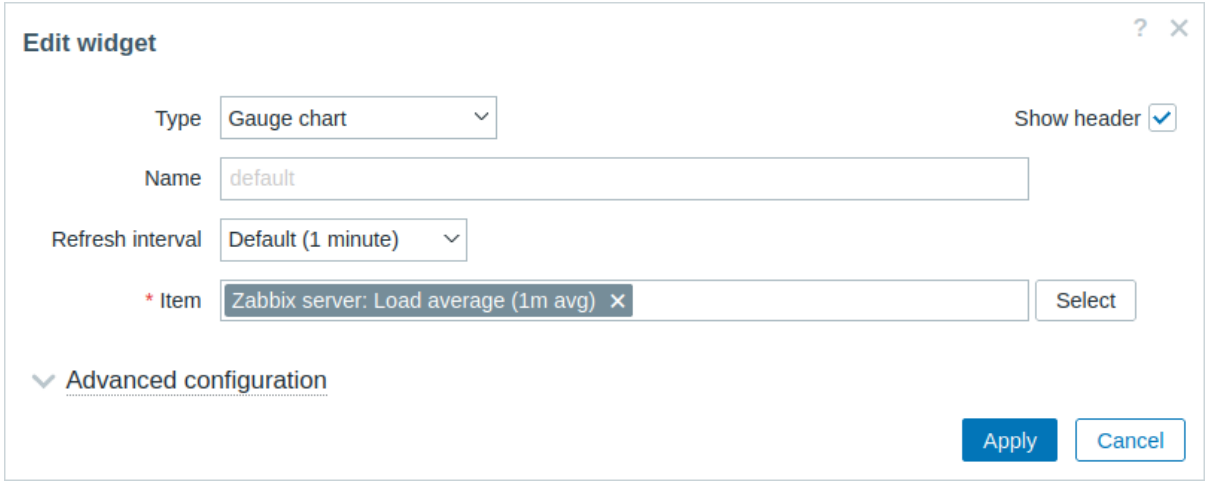
```
 */

$lefty_units = new CWidgetFieldSelectView($data['fields']['value_units']);
$lefty_static_units = (new CWidgetFieldTextBoxView($data['fields']['value_static_units']))
    ->setPlaceholder(_('value'))
    ->setWidth(ZBX_TEXTAREA_TINY_WIDTH);

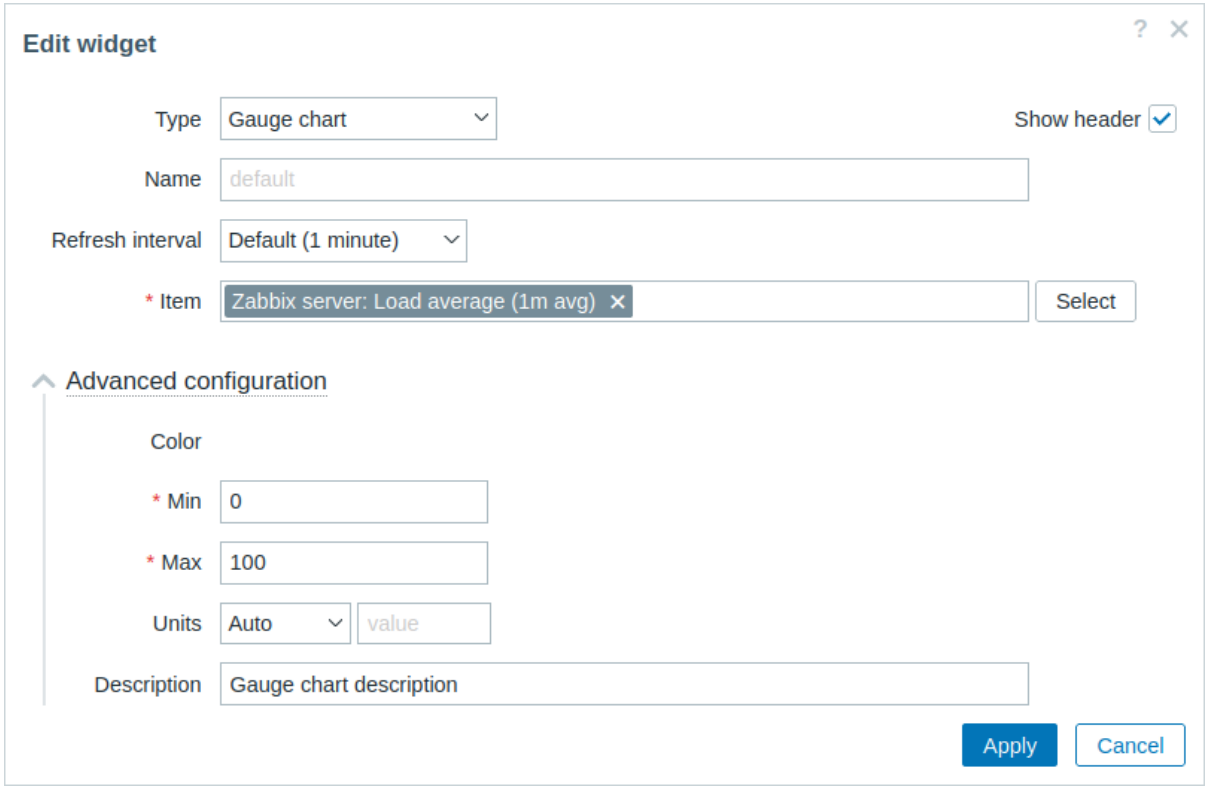
(new CWidgetFormView($data))
    ->addField(
        (new CWidgetFieldMultiSelectItemView($data['fields']['itemid']))
            ->setPopupParameter('numeric', true)
    )
    ->addFieldset(
        (new CWidgetFormFieldsetCollapsibleView(_('Advanced configuration')))
            ->addField(
                new CWidgetFieldColorView($data['fields']['chart_color'])
            )
            ->addField(
                new CWidgetFieldNumericBoxView($data['fields']['value_min'])
            )
            ->addField(
                new CWidgetFieldNumericBoxView($data['fields']['value_max'])
            )
            ->addItem([
                $lefty_units->getLabel(),
                (new CFormField([
                    $lefty_units->getView()->addClass(ZBX_STYLE_FORM_INPUT_MARGIN),
                    $lefty_static_units->getView()
                ]))
            ])
            ->addField(
                new CWidgetFieldTextBoxView($data['fields']['description'])
            )
    )
    ->show();
```

Note:
The `addField()` method of the `CWidgetFormView` class takes a CSS class string as the second parameter.

4. Return to the dashboard, switch to edit mode and click the gear icon in the widget to open the widget configuration form. The widget configuration form now contains a new expandable/collapsible section *Advanced configuration*.



5. Expand the *Advanced configuration* section to see additional widget configuration fields. Note that the field *Color* has no color picker yet. This is because the color picker must be initialized with JavaScript, which will be added in the next section - *Add JavaScript to the widget*.



Add JavaScript to the widget

In this section, you will learn how to add a gauge chart - made using JavaScript - that shows if the latest value is normal or too high/too low.

1. Create a `widget.edit.js.php` file in the `views` directory.

JavaScript will be responsible for initializing the color picker in the configuration view.

`ui/modules/lesson_gauge_chart/views/widget.edit.js.php`

```
<?php

use Modules\LessonGaugeChart\Widget;

?>

window.widget_lesson_gauge_chart_form = new class {

    init({color_palette}) {
        this._unit_select = document.getElementById('value_units');
        this._unit_value = document.getElementById('value_static_units');

        this._unit_select.addEventListener('change', () => this.updateForm());

        colorPalette.setThemeColors(color_palette);

        for (const colorpicker of jQuery('.<?= ZBX_STYLE_COLOR_PICKER ?> input')) {
            jQuery(colorpicker).colorpicker();
        }

        const overlay = overlays_stack.getById('widget_properties');

        for (const event of ['overlay.reload', 'overlay.close']) {
            overlay.$dialogue[0].addEventListener(event, () => { jQuery.colorpicker('hide'); });
        }

        this.updateForm();
    }
```

```
    updateForm() {
        this._unit_value.disabled = this._unit_select.value == <?= Widget::UNIT_AUTO ?>;
    }
};
```

2. Open `views/widget.edit.php` and add the `widget.edit.js.php` file with the JavaScript to the configuration view. To do this, use the `includeJsFile()` method. To add inline JavaScript, use the `addJavaScript()` method.

ui/modules/lesson_gauge_chart/views/widget.edit.php

```
<?php

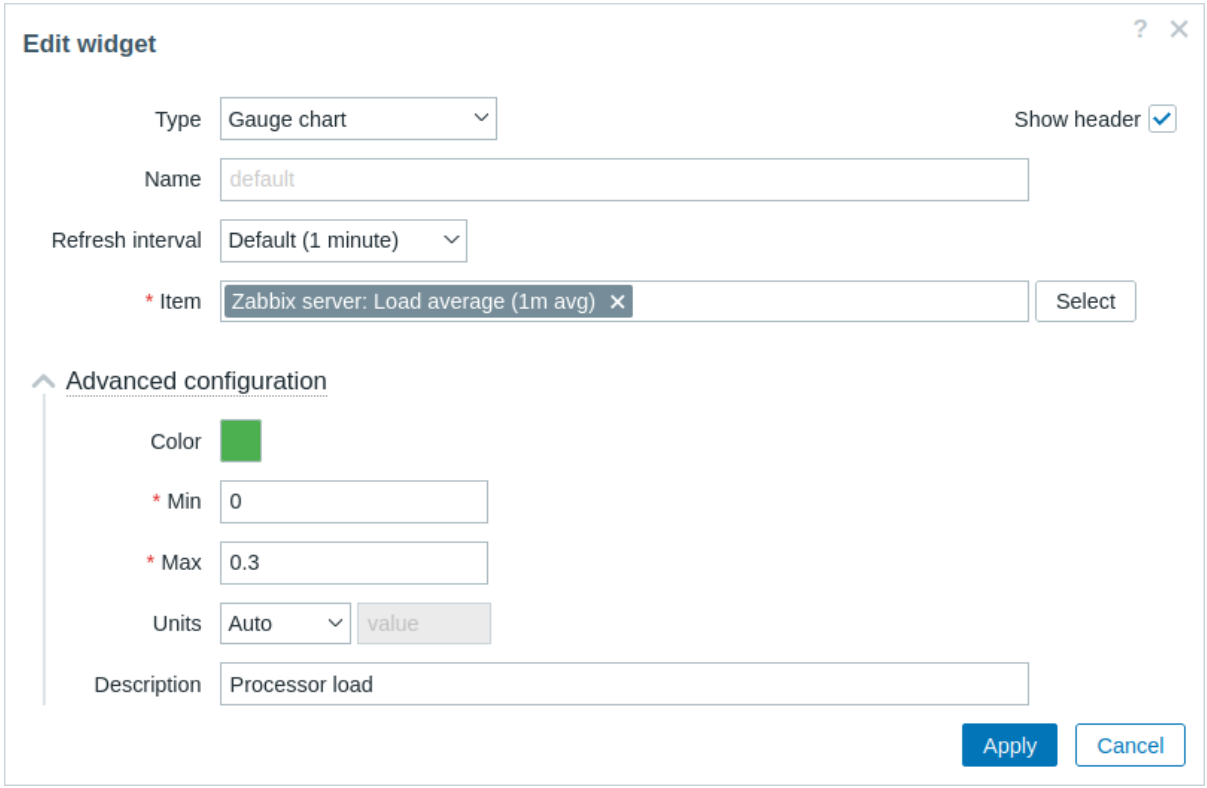
/**
 * Gauge chart widget form view.
 *
 * @var CView $this
 * @var array $data
 */

use Modules\LessonGaugeChart\Includes\WidgetForm;

$lefty_units = new CWidgetFieldSelectView($data['fields']['value_units']);
$lefty_static_units = (new CWidgetFieldTextBoxView($data['fields']['value_static_units']))
    ->setPlaceholder(_('value'))
    ->setWidth(ZBX_TEXTAREA_TINY_WIDTH);

(new CWidgetFormView($data))
    ->addField(
        (new CWidgetFieldMultiSelectItemView($data['fields']['itemid']))
            ->setPopupParameter('numeric', true)
    )
    ->addFieldset(
        (new CWidgetFormFieldsetCollapsibleView(_('Advanced configuration')))
            ->addField(
                new CWidgetFieldColorView($data['fields']['chart_color'])
            )
            ->addField(
                new CWidgetFieldNumericBoxView($data['fields']['value_min'])
            )
            ->addField(
                new CWidgetFieldNumericBoxView($data['fields']['value_max'])
            )
            ->addItem([
                $lefty_units->getLabel(),
                (new CFormField([
                    $lefty_units->getView()->addClass(ZBX_STYLE_FORM_INPUT_MARGIN),
                    $lefty_static_units->getView()
                ]))
            ])
            ->addField(
                new CWidgetFieldTextBoxView($data['fields']['description'])
            )
    )
    ->includeJsFile('widget.edit.js.php')
    ->addJavaScript('widget_lesson_gauge_chart_form.init('.json_encode([
        'color_palette' => WidgetForm::DEFAULT_COLOR_PALETTE
    ], JSON_THROW_ON_ERROR).');')
    ->show();
```

3. Return to the dashboard, click the gear icon in the widget to open the widget configuration form. Now, expand the *Advanced configuration* section to see the initialized color picker. Fill in the fields with values and select a color for the gauge chart.



4. Click *Apply* in the widget configuration form. Then click *Save changes* in the upper-right corner to save the dashboard.
5. Open `actions/WidgetView.php` and update the controller.

The `$this->fields_values` property now contains the values of all the *Advanced configuration* fields. Finalize the controller to enable passing the configuration and selected item value to the widget view.

`ui/modules/lesson_gauge_chart/actions/WidgetView.php`

```
<?php

namespace Modules\LessonGaugeChart\Actions;

use API,
    CControllerDashboardWidgetView,
    CControllerResponseData;

class WidgetView extends CControllerDashboardWidgetView {

    protected function doAction(): void {
        $db_items = API::Item()->get([
            'output' => ['itemid', 'value_type', 'name', 'units'],
            'itemids' => $this->fields_values['itemid'],
            'webitems' => true,
            'filter' => [
                'value_type' => [ITEM_VALUE_TYPE_UINT64, ITEM_VALUE_TYPE_FLOAT]
            ]
        ]);

        $history_value = null;

        if ($db_items) {
            $item = $db_items[0];

            $history = API::History()->get([
                'output' => API_OUTPUT_EXTEND,
                'itemids' => $item['itemid'],
                'history' => $item['value_type'],
                'sortfield' => 'clock',
                'sortorder' => ZBX_SORT_DOWN,
```

```
                'limit' => 1
            ]);

            if ($history) {
                $history_value = convertUnitsRaw([
                    'value' => $history[0]['value'],
                    'units' => $item['units']
                ]);
            }
        }

        $this->setResponse(new CControllerResponseData([
            'name' => $this->getInput('name', $this->widget->getName()),
            'history' => $history_value,
            'fields_values' => $this->fields_values,
            'user' => [
                'debug_mode' => $this->getDebugMode()
            ]
        ]));
    }
}
```

6. Open and modify `views/widget.view.php`.

You need to create a container for the gauge chart, which you will draw in the next steps, and a container for the description.

To pass values to JavaScript as a JSON object, use the `setVar()` method.

`ui/modules/lesson_gauge_chart/views/widget.view.php`

```
<?php

/**
 * Gauge chart widget view.
 *
 * @var CView $this
 * @var array $data
 */

(new CWidgetView($data))
    ->addItem([
        (new CDiv())->addClass('chart'),
        $data['fields_values']['description']
            ? (new CDiv($data['fields_values']['description']))->addClass('description')
            : null
    ])
    ->setVar('history', $data['history'])
    ->setVar('fields_values', $data['fields_values'])
    ->show();
```

7. Create a new directory `assets` in the `lesson_gauge_chart` directory. This directory will be used for storing JavaScript, CSS, and potentially any other assets, such as fonts or images.
8. For widget view JavaScript, create a directory `js` in the `assets` directory.
9. Create a `class.widget.js` file in the `assets/js` directory.

This JavaScript widget class will extend the base JavaScript class of all dashboard widgets - `CWidget`.

The dashboard relies on a correct implementation of a widget and communicates any relevant information to the widget through calling the respective JavaScript methods. The dashboard also expects the widget to generate events when some interaction occurs. Thus, the `CWidget` class contains a set of methods with the default implementation of widget behavior, which can be customized by extending the class.

In this case, some customization is necessary, therefore custom logic will be implemented for the following widget behavior:

- widget initialization that is responsible for defining the initial state of the widget (see the `onInitialize()` method);

- displaying widget contents (that is, drawing the gauge chart) if the widget update process has been successful and without errors (see the `processUpdateResponse(response)` method and the related `_resizeChart()` and `_updatedChart()` methods)
- resizing the widget (see the `onResize()` method and the related `_resizeChart()` method)

For other aspects of the gauge chart widget, the default implementation for widget behavior will be used. To learn more about the JavaScript methods of the `CWidget` class, see: JavaScript.

Since this JavaScript is required for the widget view, it should be loaded with the dashboard page. To enable JavaScript loading, you will need to update the `assets/js` and `js_class` parameters in the `manifest.json` file as shown in step 10.

`ui/modules/lesson_gauge_chart/assets/js/class.widget.js`

```
class WidgetLessonGaugeChart extends CWidget {

    static UNIT_AUTO = 0;
    static UNIT_STATIC = 1;

    onInitialize() {
        super.onInitialize();

        this._refresh_frame = null;
        this._chart_container = null;
        this._canvas = null;
        this._chart_color = null;
        this._min = null;
        this._max = null;
        this._value = null;
        this._last_value = null;
        this._units = '';
    }

    processUpdateResponse(response) {
        if (response.history === null) {
            this._value = null;
            this._units = '';
        }
        else {
            this._value = Number(response.history.value);
            this._units = response.fields_values.value_units == WidgetLessonGaugeChart.UNIT_AUTO
                ? response.history.units
                : response.fields_values.value_static_units;
        }

        this._chart_color = response.fields_values.chart_color;
        this._min = Number(response.fields_values.value_min);
        this._max = Number(response.fields_values.value_max);

        super.processUpdateResponse(response);
    }

    setContents(response) {
        if (this._canvas === null) {
            super.setContents(response);

            this._chart_container = this._body.querySelector('.chart');
            this._chart_container.style.height =
                `${this._getContentsSize().height - this._body.querySelector('.description').clientHeight}
            this._canvas = document.createElement('canvas');

            this._chart_container.appendChild(this._canvas);

            this._resizeChart();
        }
```

```
        this._updatedChart();
    }

    onResize() {
        super.onResize();

        if (this._state === WIDGET_STATE_ACTIVE) {
            this._resizeChart();
        }
    }

    _resizeChart() {
        const ctx = this._canvas.getContext('2d');
        const dpr = window.devicePixelRatio;

        this._canvas.style.display = 'none';
        const size = Math.min(this._chart_container.offsetWidth, this._chart_container.offsetHeight);
        this._canvas.style.display = '';

        this._canvas.width = size * dpr;
        this._canvas.height = size * dpr;

        ctx.scale(dpr, dpr);

        this._canvas.style.width = `${size}px`;
        this._canvas.style.height = `${size}px`;

        this._refresh_frame = null;

        this._updatedChart();
    }

    _updatedChart() {
        if (this._last_value === null) {
            this._last_value = this._min;
        }

        const start_time = Date.now();
        const end_time = start_time + 400;

        const animate = () => {
            const time = Date.now();

            if (time <= end_time) {
                const progress = (time - start_time) / (end_time - start_time);
                const smooth_progress = 0.5 + Math.sin(Math.PI * (progress - 0.5)) / 2;
                let value = this._value !== null ? this._value : this._min;
                value = (this._last_value + (value - this._last_value) * smooth_progress - this._min) / (t

                const ctx = this._canvas.getContext('2d');
                const size = this._canvas.width;
                const char_weight = size / 12;
                const char_shadow = 3;
                const char_x = size / 2;
                const char_y = size / 2;
                const char_radius = (size - char_weight) / 2 - char_shadow;

                const font_ratio = 32 / 100;

                ctx.clearRect(0, 0, size, size);

                ctx.beginPath();
```

```
                ctx.shadowBlur = char_shadow;
                ctx.shadowColor = '#bbb';
                ctx.strokeStyle = '#eee';
                ctx.lineWidth = char_weight;
                ctx.lineCap = 'round';
                ctx.arc(char_x, char_y, char_radius, Math.PI * 0.749, Math.PI * 2.251, false);
                ctx.stroke();

                ctx.beginPath();
                ctx.strokeStyle = `#${this._chart_color}`;
                ctx.lineWidth = char_weight - 2;
                ctx.lineCap = 'round';
                ctx.arc(char_x, char_y, char_radius, Math.PI * 0.75,
                    Math.PI * (0.75 + (1.5 * Math.min(1, Math.max(0, value)))), false
                    );
                ctx.stroke();

                ctx.shadowBlur = 2;
                ctx.fillStyle = '#1f2c33';
                ctx.font = `${(char_radius * font_ratio)|0}px Arial`;
                ctx.textAlign = 'center';
                ctx.textBaseline = 'middle';
                ctx.fillText(`${this._value !== null ? this._value : t('No data')}${this._units}`,
                    char_x, char_y, size - char_shadow * 4 - char_weight * 2
                );

                ctx.fillStyle = '#768d99';
                ctx.font = `${(char_radius * font_ratio * .5)|0}px Arial`;
                ctx.textBaseline = 'top';

                ctx.textAlign = 'left';
                ctx.fillText(`${this._min}${this._min != '' ? this._units : ''}`,
                    char_weight * .75, size - char_weight * 1.25, size / 2 - char_weight
                );

                ctx.textAlign = 'right';
                ctx.fillText(`${this._max}${this._max != '' ? this._units : ''}`,
                    size - char_weight * .75, size - char_weight * 1.25, size / 2 - char_weight
                );

                requestAnimationFrame(animate);
            }
            else {
                this._last_value = this._value;
            }
        };

        requestAnimationFrame(animate);
    }
}
```

10. Open `manifest.json` and add:

- file name (`class.widget.js`) to the array in the `assets/js` section;
- class name (`WidgetLessonGaugeChart`) to the `js_class` parameter in the `widget` section.

The `WidgetLessonGaugeChart` class will now be automatically loaded with the dashboard.

`ui/modules/lesson_gauge_chart/manifest.json`

```
{
    "manifest_version": 2.0,
    "id": "lesson_gauge_chart",
    "type": "widget",
    "name": "Gauge chart",
```

```
    "namespace": "LessonGaugeChart",
    "version": "1.0",
    "author": "Zabbix",
    "actions": {
        "widget.lesson_gauge_chart.view": {
            "class": "WidgetView"
        }
    },
    "widget": {
        "js_class": "WidgetLessonGaugeChart"
    },
    "assets": {
        "js": ["class.widget.js"]
    }
}
```

Add CSS styles to the widget

In this section you will learn how to add custom CSS styles to make the widget look more appealing.

1. For widget styles, create a new directory `css` in the `assets` directory.
2. Create a `widget.css` file in the `assets/css` directory. To style widget elements, use the selector `div.dashboard-widget-{widget id}`. To configure CSS for the whole widget, use the selector `form.dashboard-widget-{widget id}`

`ui/modules/lesson_gauge_chart/assets/css/widget.css`

```
div.dashboard-widget-lesson_gauge_chart {
    display: grid;
    grid-template-rows: 1fr;
    padding: 0;
}

div.dashboard-widget-lesson_gauge_chart .chart {
    display: grid;
    align-items: center;
    justify-items: center;
}

div.dashboard-widget-lesson_gauge_chart .chart canvas {
    background: white;
}

div.dashboard-widget-lesson_gauge_chart .description {
    padding-bottom: 8px;
    font-size: 1.750em;
    line-height: 1.2;
    text-align: center;
}

.dashboard-grid-widget-hidden-header div.dashboard-widget-lesson_gauge_chart .chart {
    margin-top: 8px;
}
```

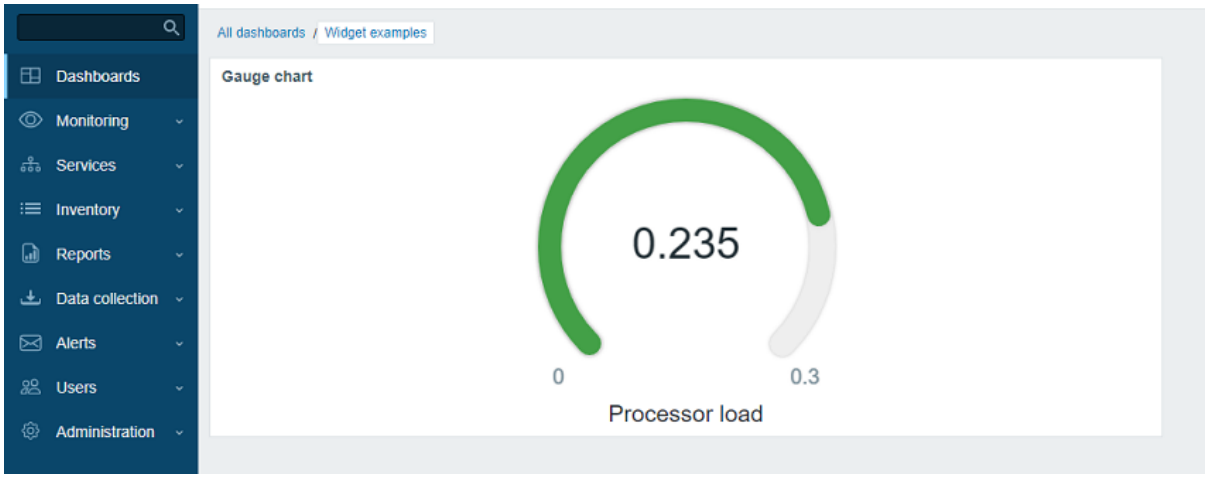
3. Open `manifest.json` and add the CSS file name (`widget.css`) to the array in the `assets/css` section. This will allow the CSS styles defined in `widget.css` to load with the dashboard page.

`ui/modules/lesson_gauge_chart/manifest.json`

```
{
    "manifest_version": 2.0,
    "id": "lesson_gauge_chart",
    "type": "widget",
    "name": "Gauge chart",
    "namespace": "LessonGaugeChart",
    "version": "1.0",
    "author": "Zabbix",
```

```
    "actions": {
        "widget.lesson_gauge_chart.view": {
            "class": "WidgetView"
        }
    },
    "widget": {
        "js_class": "WidgetLessonGaugeChart"
    },
    "assets": {
        "css": ["widget.css"],
        "js": ["class.widget.js"]
    }
}
```

4. Refresh the dashboard page to see the finished version of the widget.



Examples

This section provides files of sample modules and widgets, which you can use as a base for your custom modules.

To use a module:

1. Download the ZIP archive.
2. Unpack the content into a separate directory inside the *modules* directory of your Zabbix frontend installation (for example, *zabbix/ui/modules*).
3. Register the module in Zabbix frontend.

Module example

- When creating a host group, grant read permissions to configured user groups - hg_auto_perm.zip

Widget examples

- Minimal widget - widget_min.zip
- "Hello, world" widget using CSS only - hello_world_css.zip
- "Hello, world" widget using JavaScript only - hello_world_js.zip
- "Hello, world" widget using PHP - hello_world_php.zip

Note:
You can also use Zabbix native widgets as examples.

Changes to extension development

Changes from 8.0 to 8.2

Zabbix manpages

Reference documentation for Zabbix command-line tools and services, listing available commands, options, and usage examples.

zabbix_agent2

```
ZABBIX_AGENT2(8)            System Manager's Manual           ZABBIX_AGENT2(8)

NAME
       zabbix_agent2 - Zabbix agent 2

SYNOPSIS
       zabbix_agent2 [-c config-file]
       zabbix_agent2 [-c config-file] [-v] -p
       zabbix_agent2 [-c config-file] [-v] -t item-key
       zabbix_agent2 [-c config-file] -R runtime-option
       zabbix_agent2 [-c config-file] -T
       zabbix_agent2 -h
       zabbix_agent2 -V

DESCRIPTION
       zabbix_agent2  is  an  application for monitoring parameters of various
       services.

OPTIONS
       -c, --config config-file
              Use the alternate config-file instead of the default one.

       -R, --runtime-control runtime-option
              Perform administrative functions according to runtime-option.

           Runtime control options
               userparameter_reload
                   Reload user parameters from the configuration file

               log_level_increase
                   Increase log level

               log_level_decrease
                   Decrease log level

               help
                   List available runtime control options

               metrics
                   List available metrics

               version
                   Display version

       -p, --print
              Print known items and exit.  For each item  either  generic  de-
              faults  are used, or specific defaults for testing are supplied.
              These defaults are listed in square brackets as item key parame-
              ters.  Returned values are enclosed in square brackets and  pre-
              fixed  with  the type of the returned value, separated by a pipe
              character.  For user parameters type is always t, as  the  agent
              can  not determine all possible return values.  Items, displayed
```

```
              as working, are not guaranteed to work from the Zabbix server or
              zabbix_get when querying a running agent daemon  as  permissions
              or environment may be different.  Returned value types are:

              d   Number with a decimal part.

              m   Not  supported.   This could be caused by querying an item that
                  only works in the active mode like a log monitoring item or  an
                  item  that  requires multiple collected values.  Permission is-
                  sues or incorrect user parameters could also result in the  not
                  supported state.

              s   Text.  Maximum length not limited.

              t   Text.  Same as s.

              u   Unsigned integer.

       -t, --test item-key
              Test single item and exit.  See --print for output description.

       -v, --verbose
              Enable verbose output when test item or print known items.

       -T, --test-config
              Validate configuration file and exit.

       -h, --help
              Display this help and exit.

       -V, --version
              Output version information and exit.

FILES
       /usr/local/etc/zabbix_agent2.conf
              Default  location  of  Zabbix agent 2 configuration file (if not
              modified during compile time).

SEE ALSO
       Documentation

       zabbix_agentd(8),  zabbix_get(1),  zabbix_proxy(8),   zabbix_sender(1),
       zabbix_server(8), zabbix_js(1), zabbix_web_service(8)

Zabbix                            2022-01-06                  ZABBIX_AGENT2(8)
```

zabbix_agentd

```
ZABBIX_AGENTD(8)            System Manager's Manual           ZABBIX_AGENTD(8)

NAME
       zabbix_agentd - Zabbix agent daemon

SYNOPSIS
       zabbix_agentd [-c config-file]
       zabbix_agentd [-c config-file] -p
       zabbix_agentd [-c config-file] -t item-key
       zabbix_agentd [-c config-file] -R runtime-option
       zabbix_agentd [-c config-file] -T
       zabbix_agentd -h
       zabbix_agentd -V
```

DESCRIPTION

zabbix_agentd is a daemon for monitoring various server parameters.

OPTIONS

-c, --config config-file

Use the alternate config-file instead of the default one.

-f, --foreground

Run Zabbix agent in foreground.

-R, --runtime-control runtime-option

Perform administrative functions according to runtime-option.

Runtime control options

userparameter_reload

Reload user parameters from the configuration file

log_level_increase[=target]

Increase log level, affects all processes if target is not specified

log_level_decrease[=target]

Decrease log level, affects all processes if target is not specified

Log level control targets

process-type

All processes of specified type (active checks, collector, listener)

process-type,N

Process type and number (e.g., listener,3)

pid Process identifier, up to 65535. For larger values specify target as "process-type,N"

-p, --print

Print known items and exit. For each item either generic defaults are used, or specific defaults for testing are supplied. These defaults are listed in square brackets as item key parameters. Returned values are enclosed in square brackets and prefixed with the type of the returned value, separated by a pipe character. For user parameters type is always t, as the agent can not determine all possible return values. Items, displayed as working, are not guaranteed to work from the Zabbix server or zabbix_get when querying a running agent daemon as permissions or environment may be different. Returned value types are:

d Number with a decimal part.

m Not supported. This could be caused by querying an item that only works in the active mode like a log monitoring item or an item that requires multiple collected values. Permission issues or incorrect user parameters could also result in the not supported state.

s Text. Maximum length not limited.

t Text. Same as s.

u Unsigned integer.

```
       -t, --test item-key
              Test single item and exit.  See --print for output description.

       -T, --test-config
              Validate configuration file and exit.

       -h, --help
              Display this help and exit.

       -V, --version
              Output version information and exit.

FILES
       /usr/local/etc/zabbix_agentd.conf
              Default location of Zabbix agent configuration file (if not mod-
              ified during compile time).

SEE ALSO
       Documentation

       zabbix_get(1),  zabbix_proxy(8),  zabbix_sender(1),   zabbix_server(8),
       zabbix_js(1), zabbix_agent2(8), zabbix_web_service(8)

Zabbix                            2022-01-06                  ZABBIX_AGENTD(8)
```

zabbix_get

```
ZABBIX_GET(1)               General Commands Manual              ZABBIX_GET(1)

NAME
       zabbix_get - Zabbix get utility

SYNOPSIS
       zabbix_get  -s  host-name-or-IP  [-p  port-number]  [-I IP-address] [-t
       timeout] -k item-key
       zabbix_get -s host-name-or-IP  [-p  port-number]  [-I  IP-address]  [-t
       timeout]   --tls-connect  cert  --tls-ca-file  CA-file  [--tls-crl-file
       CRL-file] [--tls-agent-cert-issuer cert-issuer]  [--tls-agent-cert-sub-
       ject  cert-subject]  --tls-cert-file  cert-file --tls-key-file key-file
       [--tls-cipher13 cipher-string] [--tls-cipher cipher-string] -k item-key
       zabbix_get -s host-name-or-IP  [-p  port-number]  [-I  IP-address]  [-t
       timeout]     --tls-connect    psk    --tls-psk-identity    PSK-identity
       --tls-psk-file PSK-file  [--tls-cipher13  cipher-string]  [--tls-cipher
       cipher-string] -k item-key
       zabbix_get -h
       zabbix_get -V

DESCRIPTION
       zabbix_get  is  a  command  line  utility  for getting data from Zabbix
       agent.

OPTIONS
       -s, --host host-name-or-IP
              Specify host name or IP address of a host.

       -p, --port port-number
              Specify port number of agent running on the  host.   Default  is
              10050.

       -I, --source-address IP-address
              Specify source IP address.
```

-t, --timeout seconds
       Specify timeout. Valid range: 1-600 seconds (default: 30)

-k, --key item-key
       Specify key of item to retrieve value for.

-P, --protocol value
       Protocol used to communicate with agent. Values:

            auto   connect using JSON protocol, fallback and retry
                        with plaintext protocol (default)

            json   connect using JSON protocol

            plaintext   connect using plaintext protocol where just
                        item key is sent (6.4.x and older releases)

--tls-connect value
       How to connect to agent. Values:

            unencrypted
                        connect without encryption (default)

            psk   connect using TLS and a pre-shared key

            cert   connect using TLS and a certificate

--tls-ca-file CA-file
       Full pathname of a file containing the top-level CA(s) certifi-
       cates for peer certificate verification.

--tls-crl-file CRL-file
       Full pathname of a file containing revoked certificates.

--tls-agent-cert-issuer cert-issuer
       Allowed agent certificate issuer.

--tls-agent-cert-subject cert-subject
       Allowed agent certificate subject.

--tls-cert-file cert-file
       Full pathname of a file containing the certificate or certifi-
       cate chain.

--tls-key-file key-file
       Full pathname of a file containing the private key.

--tls-psk-identity PSK-identity
       PSK-identity string.

--tls-psk-file PSK-file
       Full pathname of a file containing the pre-shared key.

--tls-cipher13 cipher-string
       Cipher string for OpenSSL 1.1.1 or newer for TLS 1.3. Override
       the default ciphersuite selection criteria. This option is not
       available if OpenSSL version is less than 1.1.1.

--tls-cipher cipher-string
       GnuTLS priority string (for TLS 1.2 and up) or OpenSSL cipher
       string (only for TLS 1.2). Override the default ciphersuite se-
       lection criteria.

```
       -h, --help
              Display this help and exit.

       -V, --version
              Output version information and exit.

EXIT STATUS
       The exit status is 0 if zabbix_get successfully retrieves a value  from
       the Zabbix agent. If value retrieval failed, the exit status is 1.

EXAMPLES
       zabbix_get -s 127.0.0.1 -p 10050 -k "system.cpu.load[all,avg1]"
       zabbix_get   -s   127.0.0.1  -p  10050  -k  "system.cpu.load[all,avg1]"
       --tls-connect    cert     --tls-ca-file     /home/zabbix/zabbix_ca_file
       --tls-agent-cert-issuer   "CN=Signing   CA,OU=IT   operations,O=Example
       Corp,DC=example,DC=com" --tls-agent-cert-subject "CN=server1,OU=IT  op-
       erations,O=Example  Corp,DC=example,DC=com"  --tls-cert-file /home/zab-
       bix/zabbix_get.crt --tls-key-file /home/zabbix/zabbix_get.key
       zabbix_get  -s  127.0.0.1  -p  10050   -k   "system.cpu.load[all,avg1]"
       --tls-connect   psk   --tls-psk-identity   "PSK   ID   Zabbix   agentd"
       --tls-psk-file /home/zabbix/zabbix_agentd.psk

SEE ALSO
       Documentation

       zabbix_agentd(8), zabbix_proxy(8), zabbix_sender(1),  zabbix_server(8),
       zabbix_js(1), zabbix_agent2(8), zabbix_web_service(8)

Zabbix                            2022-01-06                      ZABBIX_GET(1)
```

zabbix_js

```
ZABBIX_JS(1)                General Commands Manual                ZABBIX_JS(1)

NAME
       zabbix_js - Zabbix JS utility

SYNOPSIS
       zabbix_js -s script-file -p input-param [-l log-level] [-t timeout]
       zabbix_js -s script-file -i input-file [-l log-level] [-t timeout]
       zabbix_js -h
       zabbix_js -V

DESCRIPTION
       zabbix_js  is  a  command  line  utility  that can be used for embedded
       script testing.

OPTIONS
       -s, --script script-file
              Specify the file name of the script to execute. If '-' is speci-
              fied as file name, the script will be read from stdin.

       -p, --param input-param
              Specify the input parameter passed to the script.  Parameter  is
              passed  verbatim  as a string and is available in the script via
              the 'value' variable. Mutually exclusive with -i.

       -i, --input input-file
              Specify a file containing the  input  parameter  passed  to  the
              script.  The  file  contents are passed verbatim as a string and
```

```
              are available in the script via the 'value' variable. If '-'  is
              specified as file name, the input will be read from stdin. Mutu-
              ally exclusive with -p.

       -w, --webdriver url
              Specify the webdriver URL.

       -l, --loglevel log-level
              Specify log level (Range: 0-5).

       -t, --timeout timeout
              Specify  the timeout in seconds. Valid range: 1-600 seconds (de-
              fault: 10)

       -h, --help
              Display this help and exit.

       -V, --version
              Output version information and exit.

EXAMPLES
       zabbix_js -s script-file.js -p example

SEE ALSO
       Documentation

       zabbix_agentd(8), zabbix_proxy(8), zabbix_sender(1),  zabbix_server(8),
       zabbix_get(1), zabbix_agent2(8), zabbix_web_service(8)

Zabbix                            2022-01-06                        ZABBIX_JS(1)
```

zabbix_proxy

```
ZABBIX_PROXY(8)             System Manager's Manual            ZABBIX_PROXY(8)

NAME
       zabbix_proxy - Zabbix proxy daemon

SYNOPSIS
       zabbix_proxy [-c config-file]
       zabbix_proxy [-c config-file] -R runtime-option
       zabbix_proxy [-c config-file] -T
       zabbix_proxy -h
       zabbix_proxy -V

DESCRIPTION
       zabbix_proxy is a daemon that collects monitoring data from devices and
       sends it to Zabbix server.

OPTIONS
       -c, --config config-file
              Use the alternate config-file instead of the default one.

       -f, --foreground
              Run Zabbix proxy in foreground.

       -R, --runtime-control runtime-option
              Perform administrative functions according to runtime-option.

           Runtime control options
               config_cache_reload
```

Reload configuration cache. Ignored if cache is being currently loaded. Active Zabbix proxy will connect to the Zabbix server and request configuration data. Default configuration file (unless -c option is specified) will be used to find PID file and signal will be sent to process, listed in PID file.

snmp_cache_reload
Reload SNMP cache.

housekeeper_execute
Execute the housekeeper. Ignored if housekeeper is being currently executed.

diaginfo[=section]
Log internal diagnostic information of the specified section. Section can be historycache, preprocessing, locks. By default diagnostic information of all sections is logged.

prof_enable[=target]
Enable profiling, affects all processes if target is not specified

prof_disable[=target]
Disable profiling, affects all processes if target is not specified

log_level_increase[=target]
Increase log level, affects all processes if target is not specified.

log_level_decrease[=target]
Decrease log level, affects all processes if target is not specified.

Log level control targets
process-type
All processes of specified type (availability manager, browser poller, configuration syncer, data sender, discovery manager, history syncer, housekeeper, http poller, icmp pinger, internal poller, ipmi manager, ipmi poller, java poller, odbc poller, poller, agent poller, http agent poller, snmp poller, preprocessing manager, preprocessing worker, self-monitoring, snmp trapper, supervisor, task manager, trapper, unreachable poller, vmware collector)

process-type,N
Process type and number (e.g., poller,3)

pid Process identifier

Profiling control targets
process-type
All processes of specified type (availability manager, browser poller, configuration syncer, data sender, discovery manager, history syncer, housekeeper, http poller, icmp pinger, internal poller, ipmi manager, ipmi poller, java poller, odbc poller, poller, agent poller, http agent poller, snmp poller, preprocessing manager, self-monitoring, snmp trapper, task manager, trapper, unreachable poller, vmware collector)

```
               process-type,N
                   Process type and number (e.g., history syncer,1)

               pid Process identifier

               scope
                   (rwlock, mutex, processing) can be used  with  process-type
                   (e.g., history syncer,1,processing)

       -T, --test-config
              Validate configuration file and exit.

       -h, --help
              Display this help and exit.

       -V, --version
              Output version information and exit.

FILES
       /usr/local/etc/zabbix_proxy.conf
              Default location of Zabbix proxy configuration file (if not mod-
              ified during compile time).

SEE ALSO
       Documentation

       zabbix_agentd(8),  zabbix_get(1),  zabbix_sender(1),  zabbix_server(8),
       zabbix_js(1), zabbix_agent2(8), zabbix_web_service(8)

Zabbix                            2022-01-06                   ZABBIX_PROXY(8)
```

zabbix_sender

```
ZABBIX_SENDER(1)            General Commands Manual           ZABBIX_SENDER(1)

NAME
       zabbix_sender - Zabbix sender utility

SYNOPSIS
       zabbix_sender [-v] -z server [-p port] [-I IP-address] [-t timeout] -s
       host -k key -o value
       zabbix_sender [-v] -z server [-p port] [-I IP-address] [-t timeout] [-s
       host] [-T] [-N] [-r] -i input-file
       zabbix_sender [-v] -c config-file [-z server] [-p port] [-I IP-address]
       [-t timeout] [-s host] -k key -o value
       zabbix_sender [-v] -c config-file [-z server] [-p port] [-I IP-address]
       [-t timeout] [-s host] [-T] [-N] [-r] -i input-file
       zabbix_sender [-v] -z server [-p port] [-I IP-address] [-t timeout] -s
       host --tls-connect cert --tls-ca-file CA-file [--tls-crl-file CRL-file]
       [--tls-server-cert-issuer cert-issuer] [--tls-server-cert-subject
       cert-subject] --tls-cert-file cert-file --tls-key-file key-file
       [--tls-cipher13 cipher-string] [--tls-cipher cipher-string] -k key -o
       value
       zabbix_sender [-v] -z server [-p port] [-I IP-address] [-t timeout] [-s
       host] --tls-connect cert --tls-ca-file CA-file [--tls-crl-file
       CRL-file] [--tls-server-cert-issuer cert-issuer]
       [--tls-server-cert-subject cert-subject] --tls-cert-file cert-file
       --tls-key-file key-file [--tls-cipher13 cipher-string] [--tls-cipher
       cipher-string] [-T] [-N] [-r] -i input-file
       zabbix_sender [-v] -c config-file [-z server] [-p port] [-I IP-address]
```

```
[-t timeout] [-s host] --tls-connect cert --tls-ca-file CA-file
[--tls-crl-file CRL-file] [--tls-server-cert-issuer cert-issuer]
[--tls-server-cert-subject cert-subject] --tls-cert-file cert-file
--tls-key-file key-file [--tls-cipher13 cipher-string] [--tls-cipher
cipher-string] -k key -o value
zabbix_sender [-v] -c config-file [-z server] [-p port] [-I IP-address]
[-t timeout] [-s host] --tls-connect cert --tls-ca-file CA-file
[--tls-crl-file CRL-file] [--tls-server-cert-issuer cert-issuer]
[--tls-server-cert-subject cert-subject] --tls-cert-file cert-file
--tls-key-file key-file [--tls-cipher13 cipher-string] [--tls-cipher
cipher-string] [-T] [-N] [-r] -i input-file
zabbix_sender [-v] -z server [-p port] [-I IP-address] [-t timeout] -s
host --tls-connect psk --tls-psk-identity PSK-identity --tls-psk-file
PSK-file [--tls-cipher13 cipher-string] [--tls-cipher cipher-string] -k
key -o value
zabbix_sender [-v] -z server [-p port] [-I IP-address] [-t timeout] [-s
host] --tls-connect psk --tls-psk-identity PSK-identity --tls-psk-file
PSK-file [--tls-cipher13 cipher-string] [--tls-cipher cipher-string]
[-T] [-N] [-r] -i input-file
zabbix_sender [-v] -c config-file [-z server] [-p port] [-I IP-address]
[-t timeout] [-s host] --tls-connect psk --tls-psk-identity PSK-iden-
tity --tls-psk-file PSK-file [--tls-cipher13 cipher-string] [--tls-ci-
pher cipher-string] -k key -o value
zabbix_sender [-v] -c config-file [-z server] [-p port] [-I IP-address]
[-t timeout] [-s host] --tls-connect psk --tls-psk-identity PSK-iden-
tity --tls-psk-file PSK-file [--tls-cipher13 cipher-string] [--tls-ci-
pher cipher-string] [-T] [-N] [-r] -i input-file
zabbix_sender -h
zabbix_sender -V
```

DESCRIPTION

zabbix_sender is a command line utility for sending monitoring data to Zabbix server or proxy. On the Zabbix server an item of type Zabbix trapper should be created with corresponding key. Note that incoming values will only be accepted from hosts specified in Allowed hosts field for this item.

OPTIONS

-c, --config config-file

Use config-file. Zabbix sender reads server details from the agentd configuration file. By default, Zabbix sender does not read any configuration file. Only parameters Hostname, ServerActive, SourceIP, TLSConnect, TLSCAFile, TLSCRLFile, TLSServerCertIssuer, TLSServerCertSubject, TLSCertFile, TLSKeyFile, TLSPSKIdentity and TLSPSKFile are supported. Hostname defined through HostnameItem parameter will not be picked up, in this case the hostname should be specified via command line (see -s option). All addresses defined in the agent ServerActive configuration parameter are used for sending data. If sending of batch data fails to one address, the following batches are not sent to this address.

-z, --zabbix-server server

Hostname or IP address of Zabbix server. If a host is monitored by a proxy, proxy hostname or IP address should be used instead. When used together with --config, overrides the entries of ServerActive parameter specified in agentd configuration file.

-p, --port port

Specify port number of Zabbix server trapper running on the server. Default is 10051. When used together with --config, overrides the port entries of ServerActive parameter specified in agentd configuration file.

-I, --source-address IP-address
       Specify source IP address. When used together with --config, overrides SourceIP parameter specified in agentd configuration file.

-t, --timeout seconds
       Specify timeout. Valid range: 1-300 seconds (default: 60)

-s, --host host
       Specify host name the item belongs to (as registered in Zabbix frontend). Host IP address and DNS name will not work. When used together with --config, overrides Hostname parameter specified in agentd configuration file.

-k, --key key
       Specify item key to send value to.

-o, --value value
       Specify item value.

-i, --input-file input-file
       Load values from input file. Specify - as <input-file> to read values from standard input. Each line of file contains whitespace delimited: <hostname> <key> <value>. Each value must be specified on its own line. Each line must contain 3 whitespace delimited entries: <hostname> <key> <value>, where "hostname" is the name of monitored host as registered in Zabbix frontend, "key" is target item key and "value" - the value to send. Specify - as <hostname> to use hostname from agent configuration file or from --host argument.

       An example of a line of an input file:

        "Linux DB3" db.connections 43

       The value type must be correctly set in item configuration of Zabbix frontend. Zabbix sender will send up to 250 values in one connection. Contents of the input file must be in the UTF-8 encoding. All values from the input file are sent in a sequential order top-down. Entries must be formatted using the following rules:

              • Quoted and non-quoted entries are supported.

              • Double-quote is the quoting character.

              • Entries with whitespace must be quoted.

              • Double-quote and backslash characters inside quoted entry must be escaped with a backslash.

              • Escaping is not supported in non-quoted entries.

              • Linefeed escape sequences (\n) are supported in quoted strings.

              • Linefeed escape sequences are trimmed from the end of an entry.

-T, --with-timestamps
       This option can be only used with --input-file option.

Each line of the input file must contain 4 whitespace delimited entries: <hostname> <key> <timestamp> <value>. Timestamp should be specified in Unix timestamp format. If target item has triggers referencing it, all timestamps must be in an increasing order, otherwise event calculation will not be correct.

An example of a line of the input file:

"Linux DB3" db.connections 1429533600 43

For more details please see option --input-file.

If a timestamped value is sent for a host that is in a “no data” maintenance type then this value will be dropped; however, it is possible to send a timestamped value in for an expired maintenance period and it will be accepted.

-N, --with-ns
This option can be only used with --with-timestamps option.

Each line of the input file must contain 5 whitespace delimited entries: <hostname> <key> <timestamp> <ns> <value>.

An example of a line of the input file:

"Linux DB3" db.connections 1429533600 7402561 43

For more details please see option --input-file.

-r, --real-time
Send values one by one as soon as they are received. This can be used when reading from standard input.

--tls-connect value
How to connect to server or proxy. Values:

unencrypted
connect without encryption (default)

psk connect using TLS and a pre-shared key

cert connect using TLS and a certificate

--tls-ca-file CA-file
Full pathname of a file containing the top-level CA(s) certificates for peer certificate verification.

--tls-crl-file CRL-file
Full pathname of a file containing revoked certificates.

--tls-server-cert-issuer cert-issuer
Allowed server certificate issuer.

--tls-server-cert-subject cert-subject
Allowed server certificate subject.

--tls-cert-file cert-file
Full pathname of a file containing the certificate or certificate chain.

--tls-key-file key-file
Full pathname of a file containing the private key.

--tls-psk-identity PSK-identity
       PSK-identity string.

--tls-psk-file PSK-file
       Full pathname of a file containing the pre-shared key.

--tls-cipher13 cipher-string
       Cipher string for OpenSSL 1.1.1 or newer for TLS 1.3. Override the default ciphersuite selection criteria. This option is not available if OpenSSL version is less than 1.1.1.

--tls-cipher cipher-string
       GnuTLS priority string (for TLS 1.2 and up) or OpenSSL cipher string (only for TLS 1.2). Override the default ciphersuite selection criteria.

-v, --verbose
       Verbose mode, -vv for more details.

-h, --help
       Display this help and exit.

-V, --version
       Output version information and exit.

EXIT STATUS
       The exit status is 0 if the values were sent and all of them were successfully processed by server. If data was sent, but processing of at least one of the values failed, the exit status is 2. If data sending failed, the exit status is 1.

EXAMPLES
       zabbix_sender -c /etc/zabbix/zabbix_agentd.conf -k mysql.queries -o 342.45
       Send 342.45 as the value for mysql.queries item of monitored host. Use monitored host and Zabbix server defined in agent configuration file.

       zabbix_sender -c /etc/zabbix/zabbix_agentd.conf -s "Monitored Host" -k mysql.queries -o 342.45
       Send 342.45 as the value for mysql.queries item of Monitored Host host using Zabbix server defined in agent configuration file.

       zabbix_sender -z 192.168.1.113 -i data_values.txt
       Send values from file data_values.txt to Zabbix server with IP 192.168.1.113. Host names and keys are defined in the file.

       echo "- hw.serial.number 1287872261 SQ4321ASDF" | zabbix_sender -c /usr/local/etc/zabbix_agentd.conf -T -i -
       Send a timestamped value from the commandline to Zabbix server, specified in the agent configuration file. Dash in the input data indicates that hostname also should be used from the same configuration file.

       echo '"Zabbix server" trapper.item ""' | zabbix_sender -z 192.168.1.113 -p 10000 -i -
       Send empty value of an item to the Zabbix server with IP address 192.168.1.113 on port 10000 from the commandline. Empty values must be indicated by empty double quotes.

       zabbix_sender -z 192.168.1.113 -s "Monitored Host" -k mysql.queries -o 342.45 --tls-connect cert --tls-ca-file /home/zabbix/zabbix_ca_file

```
       --tls-cert-file      /home/zabbix/zabbix_agentd.crt      --tls-key-file
       /home/zabbix/zabbix_agentd.key
              Send  342.45  as  the  value for mysql.queries item in Monitored
              Host host to server with IP 192.168.1.113 using  TLS  with  cer-
              tificate.

       zabbix_sender -z 192.168.1.113  -s "Monitored Host" -k mysql.queries -o
       342.45  --tls-connect  psk  --tls-psk-identity  "PSK  ID Zabbix agentd"
       --tls-psk-file /home/zabbix/zabbix_agentd.psk
              Send 342.45 as the value for  mysql.queries  item  in  Monitored
              Host  host  to  server with IP 192.168.1.113 using TLS with pre-
              shared key (PSK).

SEE ALSO
       Documentation

       zabbix_agentd(8),  zabbix_get(1),  zabbix_proxy(8),   zabbix_server(8),
       zabbix_js(1), zabbix_agent2(8), zabbix_web_service(8)

Zabbix                            2022-01-06                  ZABBIX_SENDER(1)
```

zabbix_server

```
ZABBIX_SERVER(8)            System Manager's Manual           ZABBIX_SERVER(8)

NAME
       zabbix_server - Zabbix server daemon

SYNOPSIS
       zabbix_server [-c config-file]
       zabbix_server [-c config-file] -R runtime-option
       zabbix_server [-c config-file] -T
       zabbix_server -h
       zabbix_server -V

DESCRIPTION
       zabbix_server is the core daemon of Zabbix software.

OPTIONS
       -c, --config config-file
              Use the alternate config-file instead of the default one.

       -f, --foreground
              Run Zabbix server in foreground.

       -R, --runtime-control runtime-option
              Perform administrative functions according to runtime-option.

           Runtime control options
               config_cache_reload
                   Reload configuration cache.  Ignored if cache is being cur-
                   rently  loaded.   Default configuration file (unless -c op-
                   tion is specified) will be used to find PID file and signal
                   will be sent to process, listed in PID file.

               snmp_cache_reload
                   Reload SNMP cache.

               housekeeper_execute
                   Execute the housekeeper.  Ignored if housekeeper  is  being
                   currently executed.
```

```
    trigger_housekeeper_execute
        Execute  the  trigger  housekeeper  (remove  problems  for
        deleted triggers).  Ignored if trigger housekeeper is being
        currently executed.

    secrets_reload
        Reload secrets from Vault.  Ignored if configuration  cache
        or secrets are being currently loaded.

    diaginfo[=section]
        Log  internal  diagnostic information of the specified sec-
        tion. Section can be historycache, preprocessing, alerting,
        lld, valuecache, locks, and connector.  By default, diagnos-
        tic information of all sections is logged.

    prof_enable[=target]
        Enable profiling, affects all processes if  target  is  not
        specified

    prof_disable[=target]
        Disable  profiling,  affects all processes if target is not
        specified

    ha_status
        Display high availability cluster status.  Can be performed
        only on active node.

    ha_remove_node
        ha_remove_node[=name|ID] Remove high availability  node  by
        its  name or ID listed in ha_status.  Can be performed only
        on active node.

    ha_set_failover_delay[=delay]
        Set high availability failover delay.   The  delay  can  be
        specified  using  time  suffixes.  Can be performed only on
        active node.

    log_level_increase[=target]
        Increase log level, affects all processes if target is  not
        specified

    log_level_decrease[=target]
        Decrease  log level, affects all processes if target is not
        specified

Log level control targets
    process-type
        All processes of specified  type  (alerter,  alert manager,
        alert syncer, availability manager, browser poller, config-
        uration syncer, configuration syncer worker, connector man-
        ager, connector worker, discovery manager, escalator, ha m-
        anager,  history poller,  history syncer,  housekeeper,  h-
        ttp poller, http agent poller, icmp pinger,  internal poll-
        er,  ipmi manager,  ipmi poller, java poller,  lld manager,
        lld worker, odbc poller, poller, agent poller, preprocessi-
        ng manager, preprocessing worker, proxy poller, proxy group
        manager,   report manager,  report writer, self-monitoring,
        service manager,  snmp poller,  snmp  trapper,  supervisor,
        task manager, timer, trapper,  trigger housekeeper, unreac-
        hable poller, vmware collector)

    process-type,N
```

```
                    Process type and number (e.g., poller,3)

                pid Process identifier

            Profiling control targets
                process-type
                    All  processes  of  specified type (alerter, alert manager,
                    alert syncer, availability manager, browser poller, config-
                    uration syncer, configuration syncer worker, connector man-
                    ager,    connector worker,    discovery manager, escalator,
                    ha manager,  history poller,  history syncer,  housekeeper,
                    http poller,  http agent poller, icmp pinger,  internal po-
                    ller,  ipmi manager, ipmi poller, java poller, lld manager,
                    lld worker,  odbc poller,  poller,  agent poller,  preproc-
                    essing manager, proxy poller,  proxy group manager,   repo-
                    rt manager,   report writer, self-monitoring, service mana-
                    ger, snmp poller, snmp trapper,  task manager, timer, trap-
                    per, trigger housekeeper, unreachable poller, vmware colle-
                    ctor)

                process-type,N
                    Process type and number (e.g., history syncer,1)

                pid Process identifier

                scope
                    (rwlock, mutex, processing) can be used  with  process-type
                    (e.g., history syncer,1,processing)

       -T, --test-config
              Validate configuration file and exit.

       -h, --help
              Display this help and exit.

       -V, --version
              Output version information and exit.

FILES
       /usr/local/etc/zabbix_server.conf
              Default  location  of  Zabbix  server configuration file (if not
              modified during compile time).

SEE ALSO
       Documentation

       zabbix_agentd(8),  zabbix_get(1),  zabbix_proxy(8),   zabbix_sender(1),
       zabbix_js(1), zabbix_agent2(8), zabbix_web_service(8)

Zabbix                           2022-01-06                  ZABBIX_SERVER(8)
```

zabbix_web_service

```
ZABBIX_WEB_SERVICE(8)       System Manager's Manual      ZABBIX_WEB_SERVICE(8)

NAME
       zabbix_web_service - Zabbix web service

SYNOPSIS
       zabbix_web_service [-c config-file]
       zabbix_web_service [-c config-file] -T
```

```
       zabbix_web_service -h
       zabbix_web_service -V

DESCRIPTION
       zabbix_web_service is an application for providing web services to Zab-
       bix components.

OPTIONS
       -c, --config config-file
              Use the alternate config-file instead of the default one.

       -T, --test-config
              Validate configuration file and exit.

       -h, --help
              Display this help and exit.

       -V, --version
              Output version information and exit.

FILES
       /usr/local/etc/zabbix_web_service.conf
              Default  location  of  Zabbix web service configuration file (if
              not modified during compile time).

SEE ALSO
       Documentation

       zabbix_agentd(8),  zabbix_get(1),  zabbix_proxy(8),   zabbix_sender(1),
       zabbix_server(8), zabbix_js(1), zabbix_agent2(8)

Zabbix                            2022-01-06             ZABBIX_WEB_SERVICE(8)
```

Copyright notice

Zabbix documentation is NOT distributed under the AGPL-3.0 license. Use of Zabbix documentation is subject to the following terms:

You may create a printed copy of this documentation solely for your own personal use. Conversion to other formats is allowed as long as the actual content is not altered or edited in any way. You shall not publish or distribute this documentation in any form or on any media, except if you distribute the documentation in a manner similar to how Zabbix disseminates it (that is, electronically for download on a Zabbix web site) or on a USB or similar medium, provided however that the documentation is disseminated together with the software on the same medium. Any other use, such as any dissemination of printed copies or use of this documentation, in whole or in part, in another publication, requires the prior written consent from an authorized representative of Zabbix. Zabbix reserves any and all rights to this documentation not expressly granted above.