

17. API

Overview

Zabbix API allows you to programmatically retrieve and modify the configuration of Zabbix and provides access to historical data. It is widely used to:

- Create new applications to work with Zabbix;
- Integrate Zabbix with third party software;
- Automate routine tasks.

The Zabbix API is a web based API and is shipped as 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.

More info about the protocol and JSON can be found in the [JSON-RPC 2.0 specification](#) and the [JSON format homepage](#).

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 respectfully, but some of the APIs may provide a totally different set of methods.

Performing requests

Once you've 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 `http://company.com/zabbix`, the HTTP request to call the `apiinfo.version` method may look like this:

```
POST http://company.com/zabbix/api_jsonrpc.php HTTP/1.1
Content-Type: application/json-rpc

{"jsonrpc":"2.0","method":"apiinfo.version","id":1,"auth":null,"params":{}}
```

The request must have the `Content-Type` header set to one of these values: `application/json-rpc`, `application/json` or `application/jsonrequest`.

You can use any HTTP client or a JSON-RPC testing tool to perform API requests manually, but for developing applications we suggest you use one of the [community maintained libraries](#).

Example workflow

The following section will walk you through some usage examples in more detail.

Authentication

Before you can access any data inside of Zabbix you'll need to log in and obtain an authentication token. This can be done using the `user.login` method. Let us suppose that you want to log in as a standard Zabbix Admin user. Then your JSON request will look like this:

```
{
  "jsonrpc": "2.0",
  "method": "user.login",
  "params": {
    "user": "Admin",
    "password": "zabbix"
  },
  "id": 1,
  "auth": null
}
```

Let's take a closer look at the request object. It has the following properties:

- `jsonrpc` - the version of the JSON-RPC protocol used by the API; the Zabbix API implements JSON-RPC version 2.0;
- `method` - the API method being called;
- `params` - parameters that will be passed to the API method;
- `id` - an arbitrary identifier of the request;
- `auth` - a user authentication token; since we don't have one yet, it's set to `null`.

If you provided the credentials correctly, the response returned by the API will contain the user authentication token:

```
{
  "jsonrpc": "2.0",
  "result": "0424bd59b807674191e7d77572075f33",
  "id": 1
}
```

The response object in turn contains the following properties:

- `jsonrpc` - again, the version of the JSON-RPC protocol;
- `result` - the data returned by the method;
- `id` - identifier of the corresponding request.

Retrieving hosts

We now have a valid user authentication token that can be used to access the data in Zabbix. For example, let's use the `host.get` method to retrieve the IDs, host names and interfaces of all configured hosts:

```
{
  "jsonrpc": "2.0",
  "method": "host.get",
  "params": {
    "output": [
      "hostid",
      "host"
    ],
    "selectInterfaces": [
      "interfaceid",
      "ip"
    ]
  },
  "id": 2,
  "auth": "0424bd59b807674191e7d77572075f33"
}
```

Note that the `auth` property is now set to the authentication token we've obtained by calling `user.login`.

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
}
```

For performance reasons we recommend to always list the object properties you want to retrieve and avoid retrieving everything.

Creating a new item

Let's create a new [item](#) on "Zabbix server" using the data we've obtained from the previous `host.get` request. This can be done by using the `item.create` method:

```
{
  "jsonrpc": "2.0",
  "method": "item.create",
  "params": {
    "name": "Free disk space on $1",
    "key_": "vfs.fs.size[/home/joe/,free]",
    "hostid": "10084",
    "type": 0,
    "value_type": 3,
    "interfaceid": "1",
    "delay": 30
  },
  "auth": "0424bd59b807674191e7d77572075f33",
  "id": 3
}
```

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
}
```

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

So if create methods accept arrays, we can add multiple [triggers](#) like so:

```
{
  "jsonrpc": "2.0",
  "method": "trigger.create",
  "params": [
    {
```

```
        "description": "Processor load is too high on {HOST.NAME}",
        "expression": "{Linux
server:system.cpu.load[percpu,avg1].last()}>5",
    },
    {
        "description": "Too many processes on {HOST.NAME}",
        "expression": "{Linux server:proc.num[].avg(5m)}>300",
    }
],
"auth": "0424bd59b807674191e7d77572075f33",
"id": 4
}
```

A 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, that is, set its status to "0":

```
{
  "jsonrpc": "2.0",
  "method": "item.update",
  "params": {
    "itemid": "10092",
    "status": 0
  },
  "auth": "0424bd59b807674191e7d77572075f33",
  "id": 5
}
```

A successful response will contain the ID of the updated item:

```
{
  "jsonrpc": "2.0",
  "result": {
    "itemids": [
```

```
        "10092"  
    ],  
    },  
    "id": 5  
}
```

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, that is, set their status to 0:

```
{  
  "jsonrpc": "2.0",  
  "method": "trigger.update",  
  "params": [  
    {  
      "triggerid": "13938",  
      "status": 0  
    },  
    {  
      "triggerid": "13939",  
      "status": 0  
    }  
  ],  
  "auth": "0424bd59b807674191e7d77572075f33",  
  "id": 6  
}
```

A successful response will contain the IDs of the updated triggers:

```
{  
  "jsonrpc": "2.0",  
  "result": {  
    "triggerids": [  
      "13938",  
      "13939"  
    ]  
  },  
  "id": 6  
}
```

This is the preferred method of updating. Some API methods like `host.massupdate` allow to write more simple code, but it's not recommended to use those methods, since they will be removed in the future releases.

Error handling

Up to that point everything we've tried has worked fine. But what happens if we try to make an incorrect call to the API? Let's try to create another host by calling `host.create` but omitting the mandatory `groups` parameter.

```
{
  "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"
      }
    ]
  },
  "id": 7,
  "auth": "0424bd59b807674191e7d77572075f33"
}
```

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 occurred, instead of the `result` property, the response object will contain an `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 different cases, such as, using incorrect input values, a session timeout or trying to access unexisting objects. Your application should be able to gracefully handle these kinds of errors.

API versions

To simplify API versioning, starting from 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're working with. This can be useful for adjusting your application to use version-specific features.

We guarantee feature backward compatibility inside of a major version. When making backward incompatible changes between major releases, we usually leave the old features as deprecated in the next release, and only remove them in the release after that. Occasionally, we 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.

You can follow all of the changes made to the API in the [API changelog](#).

Further reading

You now know enough to start working with the Zabbix API, but don't stop here. For further reading we suggest you have a look at the [list of available APIs](#).

From:

<https://www.zabbix.com/documentation/2.2/> - **Zabbix Documentation 2.2**

Permanent link:

<https://www.zabbix.com/documentation/2.2/manual/api>

Last update: **2019/01/28 14:49**

