

Qualys API Limits

The Qualys API enforces limits on the API calls a customer can make based on their subscription settings. The limits apply to the use of all Qualys APIs except “session” V2 API (session login/logout). Default API control settings are provided by the service. Note these settings may be customized per subscription by Qualys Support.

API limits currently apply to the Qualys API for Vulnerability Management and Policy Compliance, not APIs for Qualys apps like CA, WAS, WAF, MD, CM, Asset Management and Tagging API.

This document describes the API limits, how they are implemented by Qualys, and how you can track API usage and view recent API calls, including blocked calls.

Tell me about API Controls

API controls are applied per subscription based on your subscription’s service level. Default settings are provided and these may be customized per subscription by Qualys Support.

Concurrency Limit per Subscription (per API): This is the maximum number of concurrent API calls allowed within the subscription for each API (as per service level).

Rate Limit per Subscription (per API): Individual rate and count settings are applied (as per service level).

Rate Limit Count per Subscription (per API): The maximum number of API calls allowed within the subscription during the configured rate limit period.

Rate Limit Period per Subscription (in seconds, per API): The period of time that defines a window when API calls are counted within the subscription for each API. The window starts from the moment each API call is received by the service and extends backwards 1 hour or 1 day (Express/Consultant API Service).

SERVICE LEVEL	API CONTROLS
Express/Consultant API Service	Concurrency Limit per Subscription (per API): 1 call Rate Limit per Subscription (per API): 50 calls per Day <i>Not available for Enterprise Account</i>
Standard API Service	Concurrency Limit per Subscription (per API): 2 calls Rate Limit per Subscription (per API): 300 calls per Hour
Enterprise API Service	Concurrency Limit per Subscription (per API): 5 calls Rate Limit per Subscription (per API): 750 calls per Hour
Premium API Service	Concurrency Limit per Subscription (per API): 10 calls Rate Limit per Subscription (per API): 2000 calls per Hour

How it works

When an API call is received, Qualys first checks the concurrency limit; and if the concurrency limit has been exceeded the API call is blocked and an error is returned. In the case where the concurrency limit has not been exceeded, the service checks the rate limit; and if the rate limit has been exceeded the API call is blocked and an error is returned.

Concurrency Limit

The API concurrency is calculated each time an API call is received and checked against the concurrency limit for the subscription (2 by default for Standard API Service).

Example: A subscription for Standard API Service has the default API control settings and there are multiple users. A user makes 2 `asset_group.php` API calls and both API call instances are running. The `asset_group.php` API concurrency limit has been reached, so it's not possible for any user to make another successful `asset_group.php` API call until at least 1 `asset_group.php` API call instance completes. There must be a maximum of 1 `asset_group.php` API call instance running at the time the user makes a new `asset_group.php` API call.

When a user makes an API call for an API that has 2 concurrent API call instances already running, then the new API call is blocked, a Concurrency Limit Exceeded error is reported in the XML output, and an entry is added to the Qualys Activity like this: "API blocked (concurrency): `asset_group.php`"

Rate Limit

The rate count and period are calculated dynamically each time an API call is received. The rate period represents a rolling window when API calls are counted.

A user may distribute the quota of API calls arbitrarily within the time window. Using a subscription for Standard API Service this quota is 300 API calls per hour.


Example: A subscription for Standard API Service has the default API control settings. If 300 API calls are received in a 5 minute period and none are blocked by any API limiting rules, then you need to wait 55 minutes before making the next call to the API. During the wait period API calls will be blocked by the rate limiting rule.

When a user makes an API call for an API that is blocked due to exceeding the rate limit, a Rate Limit Exceeded appears in the XML output, and an entry is added to the Qualys Activity Log like this: "API blocked (rate): `asset_group.php`"

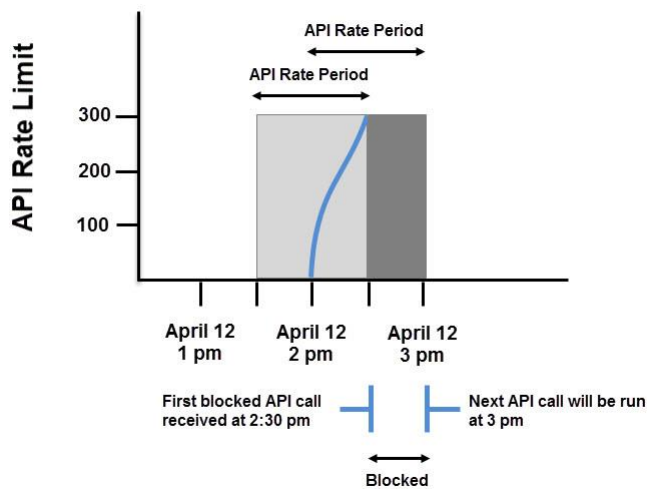
Let's take a look

Let's review API call history for a subscription for the Standard API Service with the default API limits

Example: An API call was received on April 3 at 10 AM. The service calculated the API rate period by creating a window that extends backwards 1 hour from the time the API call was received to April 3 at 9 AM. The total number of API calls received in the window is 200 so the API call instance received on April 3 at 10 AM runs successfully.



Example: 300 API calls were received starting April 12 at 2 PM. The first blocked API call was received on April 12 at 2:30 PM. Users could not run API calls for 30 minutes. The next time an API can be received and run is April 12 at 3 PM, assuming there is a maximum of 1 API call instance currently running at that time.



Errors Returned in XML output

Each API call returns an informational message in the XML output when the API call was blocked because the concurrency limit or rate limit has been exceeded for the API being called. Please note if an API call was blocked, only one error is returned. In the case where the concurrency limit has been exceeded, a Concurrency Limit Exceeded error will be reported (and a Rate Limit Exceeded error will not be reported).

Concurrency Limit Exceeded Error

An API call returns this error in the XML output in the case where a user makes an API call and the total number of concurrent API instances, which are currently running, exceeds the limit for the subscription.

For a V1 API, the error will appear like this:

```
<GENERIC_RETURN>
  <API name="asset_group_list.php" username="acme_es1" at="2017-04-12T14:52:39Z" />
  <RETURN status="FAILED" number="1999">
    This API cannot be run again until 1 currently running API instance has finished.
  </RETURN>
</GENERIC_RETURN>
```

For a V2 API, the error will appear like this:

```
<SIMPLE_RETURN>
  <RESPONSE>
    <DATETIME>2017-04-12T14:52:39Z </DATETIME>
    <CODE>1960</CODE>
    <TEXT> This API cannot be run again until 1 currently running API instance has finished.</TEXT>
    <ITEM_LIST>
      <ITEM>
        <KEY>CALLS_TO_FINISH</KEY>
        <VALUE>2</VALUE>
      </ITEM>
    </ITEM_LIST>
  </RESPONSE>
</SIMPLE_RETURN>
```

Rate Limit Exceeded Error

An API call returns this error in the XML output in the case where a user makes an API call and the rate limit for the API, as defined for the subscription, has already been reached. In other words, the rate limit count (maximum number of API call instances) has already been reached for the rate limit period.

For a V1 API, the error will appear like this:

```
<GENERIC_RETURN>
  <API name="asset_group_list.php" username="acme_es1" at="2017-04-12T14:52:39Z" />
  <RETURN status="FAILED" number="1999">
    This API cannot be run again for another 23 hours, 57 minutes and 54 seconds.
  </RETURN>
</GENERIC_RETURN>
```

For a V2 API, the error will appear like this:

```
<SIMPLE_RETURN>
<RESPONSE>
  <DATETIME>2017-04-12T14:52:39Z </DATETIME>
  <CODE>1965</CODE>
  <TEXT> This API cannot be run again for another 23 hours, 57 minutes and 54 seconds.</TEXT>
  <ITEM_LIST>
    <ITEM>
      <KEY>SECONDS_TO_WAIT</KEY>
      <VALUE>68928</VALUE>
    </ITEM>
  </ITEM_LIST>
</RESPONSE>
</SIMPLE_RETURN>
```

API Usage in HTTP response headers

Your subscription’s API usage and quota information is exposed in the HTTP response headers generated by Qualys APIs (all APIs except “session” V2 API).

The HTTP response headers generated by Qualys APIs are described below.

HEADER	DESCRIPTION
X-RateLimit-Limit	Maximum number of API calls allowed in any given time period of <number-sec> seconds, where <number-sec> is the value of X-RateLimit-Window-Sec.
X-RateLimit-Window-Sec	Time period (in seconds) during which up to <number-limit> API calls are allowed, where <number-limit> is the value of X-RateLimit-Limit.
X-RateLimit-Remaining	Number of API calls you can make right now before reaching the rate limit <number-limit> in the last <number-sec> seconds.
X-RateLimit-ToWait-Sec	The wait period (in seconds) before you can make the next API call without being blocked by the rate limiting rule.
X-ConcurrencyLimit-Limit	Number of API calls you are allowed to run simultaneously.
X-ConcurrencyLimit-Running	Number of API calls that are running right now (including the one identified in the current HTTP response header).
X-Powered-By	You can track API usage by user using this header which includes a unique ID for each subscription and user. This capability must be enabled for your account by Qualys. Learn more

The HTTP status code “OK” (example: “HTTP/1.1 200 OK”) is returned in the header for normal (not blocked) API calls. The HTTP status code “Conflict” (example: “HTTP/1.1 409 Conflict”) is returned for API calls that were blocked.

Sample HTTP Response Headers

Sample 1: Normal API call (API call not blocked)

Returned from API call using HTTP authentication.

```
< HTTP/1.1 200 OK
< Date: Fri, 28 Apr 2017 05:28:39 GMT
< Server: Qualys
< X-Frame-Options: SAMEORIGIN
< X-RateLimit-Limit: 300
< X-RateLimit-Window-Sec: 3600
< X-Concurrency-Limit-Limit: 50
< X-Concurrency-Limit-Running: 0
< X-RateLimit-ToWait-Sec: 0
< X-RateLimit-Remaining: 287
< X-Qualys-Application-Version: QWEB-8.10.0.0-SNAPSHOT-20170427151441#3811
< X-Server-Virtual-Host: qualysapi.qualys.com
< X-Server-Http-Host: qualysapi.qualys.com
< Transfer-Encoding: chunked
< Content-Type: text/xml;charset=UTF-8
```

Sample 2: API Call Blocked - Rate Limit exceeded

Returned from API call using HTTP authentication.

```
< HTTP/1.1 409 Conflict
< Date: Fri, 28 Apr 2017 06:32:34 GMT
< Server: Qualys
< X-Frame-Options: SAMEORIGIN
< X-RateLimit-Limit: 1
< X-RateLimit-Window-Sec: 3600
< X-Concurrency-Limit-Limit: 5
< X-Concurrency-Limit-Running: 0
< X-RateLimit-ToWait-Sec: 981
< X-RateLimit-Remaining: 0
< X-Qualys-Application-Version: QWEB-8.10.0.0-SNAPSHOT-20170427151441#3811
< X-Server-Virtual-Host: qualysapi.qualys.com
< X-Server-Http-Host: qualysapi.qualys.com
< Transfer-Encoding: chunked
< Content-Type: text/xml;charset=UTF-8
```

Sample 3: API V2 Call Blocked - Concurrency Limit exceeded

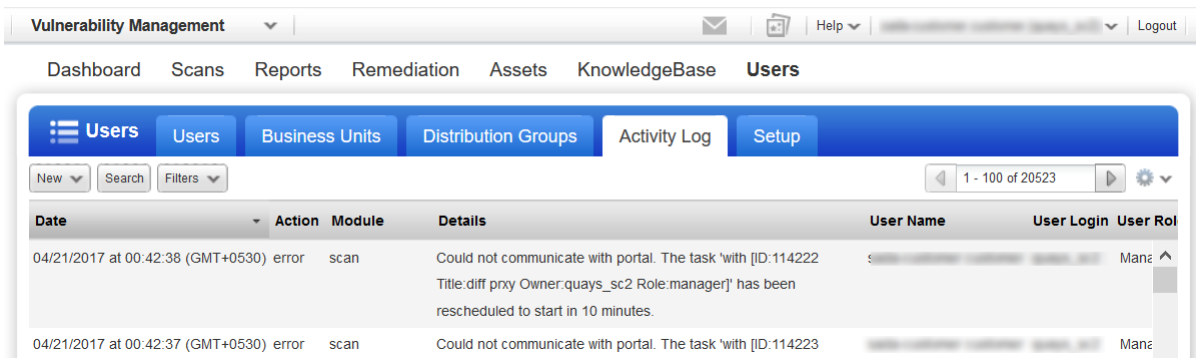
Returned from API V2 call using API V2 session authentication.

```
< HTTP/1.1 409 Conflict
< Date: Fri, 28 Apr 2017 05:45:20 GMT
< Server: Qualys
< X-Frame-Options: SAMEORIGIN
< X-RateLimit-Limit: 10
< X-RateLimit-Window-Sec: 10
< X-Concurrency-Limit-Limit: 1
< X-Concurrency-Limit-Running: 1
< X-Qualys-Application-Version: QWEB-8.10.0.0-SNAPSHOT-20170427151441#3811
< X-Server-Virtual-Host: qualysapi.qualys.com
< X-Server-Http-Host: qualysapi.qualys.com
< Transfer-Encoding: chunked
< Content-Type: text/xml;charset=UTF-8
```

Good to Know - In the case where the concurrency limit has been reached, no information about rate limits will appear in the HTTP headers.

Qualys Activity Log

The Qualys Activity Log shows details about user activities including actions taken using the Qualys user interface and API. Just log into Qualys portal, select VM from the module picker, and go to Users > Activity Log.

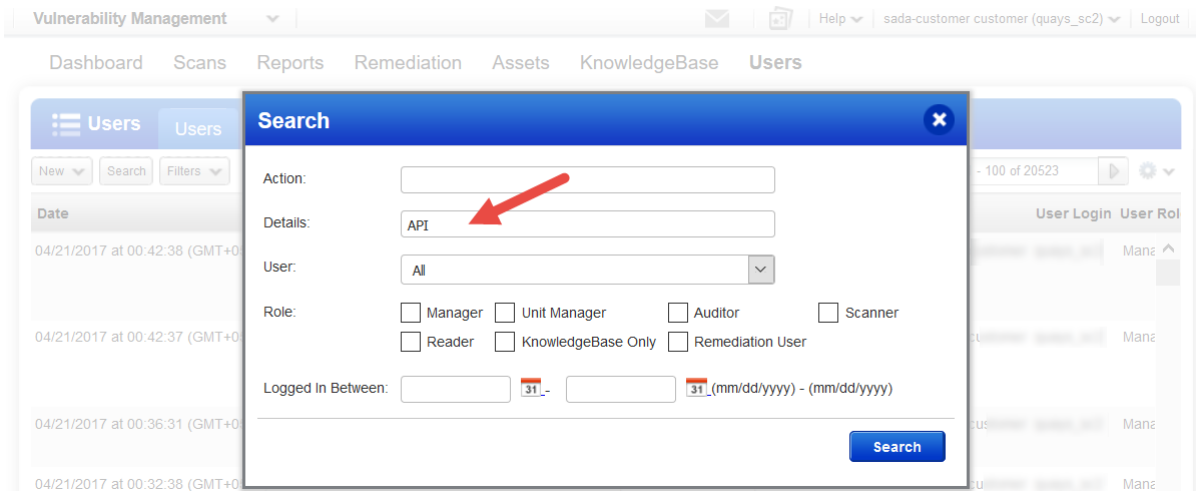


The screenshot shows the Qualys Activity Log interface. At the top, there is a navigation bar with 'Vulnerability Management' and 'Users' selected. Below the navigation bar, there are tabs for 'Users', 'Business Units', 'Distribution Groups', 'Activity Log', and 'Setup'. The 'Activity Log' tab is active. Below the tabs, there is a search bar and a table of activity entries. The table has columns for Date, Action, Module, Details, User Name, User Login, and User Role. The first entry shows an error at 04/21/2017 at 00:42:38 (GMT+0530) for a scan action, with details: 'Could not communicate with portal. The task 'with [ID:114222 Title:diff prxy Owner:quays_sc2 Role:manager]' has been rescheduled to start in 10 minutes.'

Date	Action	Module	Details	User Name	User Login	User Role
04/21/2017 at 00:42:38 (GMT+0530)	error	scan	Could not communicate with portal. The task 'with [ID:114222 Title:diff prxy Owner:quays_sc2 Role:manager]' has been rescheduled to start in 10 minutes.			Manag
04/21/2017 at 00:42:37 (GMT+0530)	error	scan	Could not communicate with portal. The task 'with [ID:114223			Manag

Viewing API related activity logs

Use the search option to find API logs. For example you can search for API in Details.



The screenshot shows the Qualys Activity Log search interface. A search modal is open, allowing users to filter activity logs. The modal has a search bar and several filters: Action, Details, User, Role, and Logged In Between. The 'Details' field is highlighted with a red arrow and contains the text 'API'. The 'User' field is set to 'All'. The 'Role' field has several checkboxes: Manager, Unit Manager, Auditor, Scanner, Reader, KnowledgeBase Only, and Remediation User. The 'Logged In Between' field has two date pickers. A 'Search' button is located at the bottom right of the modal.

Action:	Details:	User:	Role:	Logged In Between:
	API	All	<input type="checkbox"/> Manager <input type="checkbox"/> Unit Manager <input type="checkbox"/> Auditor <input type="checkbox"/> Scanner <input type="checkbox"/> Reader <input type="checkbox"/> KnowledgeBase Only <input type="checkbox"/> Remediation User	

Concurrency Limit error - For an API call that exceeded the Concurrency Limit the log details is in the format: "API blocked (concurrency): <API name>"

The screenshot shows the 'Users' section of the Vulnerability Management interface. The 'Activity Log' tab is selected. A table lists activity entries with columns for Date, Action, Module, Details, and User Name. A red dotted arrow points to a specific entry in the Details column.

Date	Action	Module	Details	User Name
04/28/2017 at 11:18:55 (GMT+0530)	request	auth	API: /api/2.0/fo/asset/hostvm/detection/index.php	[redacted]
04/28/2017 at 11:18:26 (GMT+0530)	request	auth	API: /api/2.0/fo/asset/hostvm/detection/index.php	[redacted]
04/28/2017 at 11:18:00 (GMT+0530)	request	auth	API: /api/2.0/fo/asset/hostvm/detection/index.php	[redacted]
04/28/2017 at 11:16:09 (GMT+0530)	request	auth	API blocked (concurrency): /api/2.0/fo/asset/group/index.php	[redacted]
04/28/2017 at 11:16:08 (GMT+0530)	request	auth	API: /api/2.0/fo/asset/group/index.php	[redacted]
04/28/2017 at 11:15:20 (GMT+0530)	request	auth	API blocked (concurrency): /api/2.0/fo/asset/group/index.php	[redacted]
04/28/2017 at 11:15:19 (GMT+0530)	request	auth	API: /api/2.0/fo/asset/group/index.php	[redacted]

Rate Limit error - For an API call that exceeded the Rate Limit the details entry is the format: "API blocked (rate): <API name>"

The screenshot shows the 'Users' section of the Vulnerability Management interface. The 'Activity Log' tab is selected. A table lists activity entries with columns for Date, Action, Module, Details, and User Name. A red dotted arrow points to a specific entry in the Details column.

Date	Action	Module	Details	User Name
04/28/2017 at 12:02:35 (GMT+0530)	request	auth	API blocked (rate): /api/2.0/fo/asset/hostvm/detection/index.php	[redacted]
04/28/2017 at 12:02:34 (GMT+0530)	request	auth	API blocked (rate): /api/2.0/fo/asset/hostvm/detection/index.php	[redacted]
04/28/2017 at 11:18:55 (GMT+0530)	request	auth	API: /api/2.0/fo/asset/hostvm/detection/index.php	[redacted]

Tell me about Activity Log permissions -The logs you see depends on your user role. Managers see actions performed by all user. Unit Managers see actions performed by all users within their business unit. Scanners and Readers see their own actions only. Auditors see compliance actions performed by all users.

The screenshot shows the 'Users' section of the Vulnerability Management interface. The 'Recent API Calls' tab is selected and highlighted in yellow. A table lists activity entries with columns for Date and Module.

Date	Module
04/21/2017 at 00:42:38 (GMT+0530)	error scan

Viewing Recent API Calls

The API Calls List includes APIs subject to the API limits (all APIs except "session" V2 API). By default the service displays API calls submitted (by users) and/or updated (by Qualys) in the past week.

API Processes						
API	User Login	Incident Signature	Incident Created Date	State	Submitted	Last Updated
/api/2.0/fo/asset /group/index.php				Finished	04/28/2017 at 11:16:08 (GMT+0530)	04/28/2017 at 11:16:10 (GMT+0530)
/api/2.0/fo/asset /group/index.php				Blocked (Concurrency)	04/28/2017 at 11:16:09 (GMT+0530)	04/28/2017 at 11:16:09 (GMT+0530)
/api/2.0/fo/asset				Finished	04/28/2017 at 11:15:19 (GMT+0530)	04/28/2017 at 11:15:21 (GMT+0530)

Tell me about Recent API Calls permissions - Managers see all API calls performed by all users in the subscription. Unit Managers, Scanners and Readers see all VM API calls, and PC and WAS API calls when the corresponding modules are enabled for their account.

State - The value for State will be one of the following: Queued, Running, Expired, Finished (means API call completed successfully), Blocked (Rate) or Blocked (Concurrency).

User Login - This is the user who made the API call. The value “-” appears here when you are restricted from viewing this information due to your account permissions and all these conditions are true:

- 1) Your user role is Unit Manager, Scanner or Reader,
- 2) The user who performed the API call is not in your business unit, and
- 3) Your subscription has this user permission selected: “Restrict view of user information for users outside of business unit”.

Last updated: January 30, 2018