

API Request Limits

Last Modified on 03/09/2023 1:43 pm EST |

API access is only available to enterprise customers. If you are interested, please [contact us](#) for more information!

Is there a rate limit of API requests?

- Professional and Full Access accounts - 240 requests per minute

Note: These limits are Account limits, not user limits.

As you are approaching your request limit (<10% of your calls per minute remaining), you will see the following return within the headers:

X-SG-Warning: You are reaching your API request rate limit and may be throttled or blocked if you continue to make excessive requests.

If you are using Alchemer's Salesforce Integration (Salesforce SOAP API), and are concerned about Salesforce API limits, [click here](#) for additional information.

What will I see if I reach the Request Limit?

Should you reach an API Request Limit, users see one of the following two messages with the status code of 429:

"You have exceeded your requests per minute, see </help/api-request-limits> for more details"

"Please wait for other requests to complete." This message occurs when identical requests are made within a 60-second window.

If you routinely hit or exceed the above API request limits, your API access will be disabled.

What can I expect as far as API response time?

The time it takes for an API response is dependent on a number of factors including:

1. The call you are making.
2. The amount of data that is being returned.
3. The amount of data you have.
4. Server load.

Responses from the API can take anywhere from <1 second to 30 seconds. Requests will time out after 30 seconds returning a 500 error. After a call times out subsequent calls will return a 400 error for 1 minute while the throttling resets.

What does a 429 status code mean?

This status code is returned on two occasions:

- when too many requests are made per minute (see request limits). The status code will return for up to a minute after you have met your limit while throttling resets.
- When a user makes the same API call request more than once within a minute.

Do API returns cache?

Yes, get requests will cache for 60 seconds. If you make repeated API GET requests that are identical, the return will be cached and will thus return identical results. You will also see the following return within the headers:

```
X-SG-Cached: true
```

Identical **GET** requests should only be made after 60 seconds and cached on the users' systems. Excessive requests may be rejected by Alchemer. Alchemer reserves the right to terminate access to the API if we see an impact on other customers as a result of circumventing the policy.

Response Processing and API:

System data (Campaign, Survey, Response, etc.) is subject to response processing within the Alchemer platform, which varies based on a couple of factors, and ultimately the current server load. If you are looking to access data within the system, the time between when the data is submitted (even if submitted via the API) and when the data is available via the API can be upwards of 5 minutes.

Your code should account for these delays, as well as the request limits on your account.

Is there a limit to the size of an individual request?

Our web servers have a request-line limit of 8000 bytes (characters). We recommend imposing a tighter limit of perhaps 4000 characters to ensure that response timeouts (presently 20 seconds) are less likely to occur during high traffic time. It is important to note that for tests in browsers, each has their own limits (IE is notoriously short for some versions).

Is there a specific format for API requests?

Requests should be formatted as URL parameters or HTTP Headers.