

HTTP Conditional GET

HTTP Conditional GET

A *conditional GET* is an HTTP GET request that may return an HTTP 304 response (instead of HTTP 200). An HTTP 304 response indicates that the resource has not been modified since the previous GET, and so the resource is not returned to the client in such a response. See [RFC 7232](#) for details.

There are at least two (not completely independent) approaches to conditional GET:

1. Last-Modified / If-Modified-Since
2. ETag / If-None-Match

In both cases, the value of a response header is used as the value of a subsequent request header. For example, note the Last-Modified and ETag headers in the response to this HEAD request for InCommon metadata:

```
$ MD_LOCATION=http://md.incommon.org/InCommon/InCommon-metadata.xml
$ /usr/bin/curl --silent --head $MD_LOCATION
HTTP/1.1 200 OK
Date: Tue, 30 Dec 2014 19:25:39 GMT
Server: Apache
Last-Modified: Mon, 29 Dec 2014 20:24:24 GMT
ETag: "110328-b28945-50b60a9050e00"
Accept-Ranges: bytes
Content-Length: 11700549
Connection: close
Content-Type: application/samlmetadata+xml
```

If we take the value of the Last-Modified header from the previous response as the value of the If-Modified-Since header in the following request, we receive a 304 response (and no content) from the server:

```
$ /usr/bin/curl --silent --head $MD_LOCATION \
  --header 'If-Modified-Since: Mon, 29 Dec 2014 20:24:24 GMT'
HTTP/1.1 304 Not Modified
Date: Tue, 30 Dec 2014 19:26:20 GMT
Server: Apache
Connection: close
ETag: "110328-b28945-50b60a9050e00"
```

Similarly, if we take the value of the ETag header from the previous response as the value of the If-None-Match header in the following request, we again receive a 304 response:

```
$ /usr/bin/curl --silent --head $MD_LOCATION \
  --header 'If-None-Match: "110328-b28945-50b60a9050e00"'
HTTP/1.1 304 Not Modified
Date: Tue, 30 Dec 2014 19:26:58 GMT
Server: Apache
Connection: close
ETag: "110328-b28945-50b60a9050e00"
```

The use of conditional GET has significant benefits, on both the client and the server (and the intervening network as well). On the InCommon metadata server, roughly 3/4 of all metadata requests result in HTTP 304. That translates into many thousands of metadata requests per day that conveniently avoid the unnecessary overhead of metadata refresh. For a file whose size is large and growing, that represents a significant cost savings.

Conditional GET has security benefits as well. Since requests that result in HTTP 304 are issued virtually without penalty, a client can request metadata more frequently than absolutely necessary. In the case of InCommon metadata, which is produced daily, hourly requests will result in just one HTTP 200 response in a typical 24-hour period. If, however, InCommon Operations signs metadata more than once per day (which happens on occasion), or more importantly, a key in metadata is compromised, necessitating an immediate production run, the fact that clients are attempting to refresh metadata hourly has significant potential benefit.

Tools and Tips

The command-line examples above suggest a tool based on `curl` is possible. In fact, [a tool that implements HTTP conditional GET](#) can be downloaded from GitHub. It's a bash script called `cget` that caches the HTTP response header along with the resource. In this way, subsequent requests can provide the appropriate request headers. If the server supports conditional GET, and the resource has not changed since the previous GET (as indicated by HTTP 304), the script accesses the resource from cache.

Let's use the script to illustrate HTTP conditional GET (as we did with `curl` above). Here's how to fetch and cache a metadata file:

```
$ echo $MD_LOCATION
http://md.incommon.org/InCommon/InCommon-metadata.xml
$ cget.sh -H $MD_LOCATION
HTTP/1.1 200 OK
Date: Tue, 30 Dec 2014 19:28:30 GMT
Server: Apache
Last-Modified: Mon, 29 Dec 2014 20:24:24 GMT
ETag: "110328-b28945-50b60a9050e00"
Accept-Ranges: bytes
Content-Length: 11700549
Connection: close
Content-Type: application/samlmetadata+xml
```

Subsequent requests will produce HTTP 304 responses as long as the metadata file does not change:

```
$ cget.sh -H $MD_LOCATION
HTTP/1.1 304 Not Modified
Date: Tue, 30 Dec 2014 19:29:01 GMT
Server: Apache
Connection: close
ETag: "110328-b28945-50b60a9050e00"
```

Later versions of Shibboleth (at least IdP 2.2 and SP 2.4) implement HTTP conditional GET (and more) so the above script is not particularly useful unless you're running something other than Shibboleth. For instance, simpleSAMLphp does everything *except* HTTP conditional GET, so users of simpleSAMLphp might find the above script useful.

| File | Modified |
|------|----------|
|------|----------|

No files shared here yet.