

Entity Tags

By Akshat Kukreti

Etags (Entity tags)

- Part of HTTP
- Used for Cache validation: Is the cached copy of a web page still good?
- Efficient caching as the server does not need to send a full response in case of no change
- Unique ID for a resource found at a URL
- A new Etag is associated with a resource whenever the resource is modified.
- Generation: Collision-resistant hash functions
 - A hash of the last modified timestamp.

Etags (Entity tags) contd...

- Strong Etag
 - Matching Etags mean that the contents of the resources being matched are identical byte-for-byte.
 - “string”
- Weak Etag
 - Matching resources are similar but not byte-for-byte. Not good for byte serving.
 - W/“string”

Etag header fields

- If-Match
 - If-Match: (“*” | Etag string)
 - If Etag string matches an Etag value of a resource or if “*” is given and a current version of the resource exists, the resource is returned
 - If Etag string does not match or if “*” is given and no current version of the resource does not exist, the server responds with a status of **412 (Pre-condition failed)**

Etag header fields

- If-None-Match
 - If-None-Match(“*” | Etag string)
 - If Etag string matches the Etag for the resource, the server responds with a status of **304 (Not modified)**
 - If the Etag string is different from the Etag of the resource, the resource is returned.

Other header fields

- If-Modified-Since
 - If-Modified-Since: Date
 - If a resource entity has not been modified since Date, the response is a status **304 (Not Modified)**
 - If the resource entity has been modified since Date, the entity is returned.

Tracking with Entity tags

- Server can keep on returning the cached Etag whenever the client makes a request.
 - The Etag persists enabling the server to track the user even if cookies are disabled.
- Kissmetrics Etag Respawning (http://ashkansoltani.org/docs/respawn_redux.html)
 - Tracking user by using Etag as a persistent identifier.

Tracking with Entity tags

- Kissmetrics Etag Respawning
 - When the client makes the first request, a JavaScript generates a random value and returns it as an Etag and also sets a cookie with that value.
 - In the future when the client makes a request, a If-None-Match header with the Etag value is sent to the server.
 - The unique value can be used to track the user.