

The QUERY Method

What are we doing?

GET

- Safe
- Idempotent
- Body has “no defined semantics”

POST

- Unsafe
- Potentially non-idempotent
- Body is expected
 - Semantics per target resource

What are we doing?

GET

- Safe
- Idempotent
- Cacheable
- Body has “no defined semantics”

QUERY

POST

- Potentially unsafe
- Potentially non-idempotent
- Not cacheable
- Body is expected
 - Semantics per target resource

Key Updates since IETF 120

- Define use of Content-Location and Location headers:
 - **Content-Location:** where you might retrieve *these results* in the future
 - **Location:** where you might retrieve updated answers to *this query* in the future
- Guidance around generating URIs for C-L/L headers
 - Don't embed sensitive information
- Register Accept-Query header for discovery purposes
 - Header was already present; IANA section added

#1917: Conditional QUERY

What does a If-None-Match or If-Modified-Since header mean on a QUERY request?

- Only perform the query if the state of (...the data...?) has changed
- Only return the results if the results have changed

Current text says the query is performed conditionally; proposal is to always perform the query and operate on the results.

#2860 & #2934: Accept-Query

Accept-Query header in response indicates

- ...that resource would have accepted a QUERY request
- ...what media types the resource can process in QUERY

Questions:

- Should we define a way to accept any media type?
- Should we use the format of Accept or a newly-minted SF format?

#2894 & #2924: Media-type semantics

How does a client learn what the server expects in the QUERY body? Accept-Query provides media type, but that's insufficient.

- *A priori* knowledge, probably from API specs
 - ...so this spec doesn't need to say anything?
- Defined by media-type, probably implying new media types

In other words, should this look more like POST or like PATCH? If we specify them, should we define one or more, like jsonpath?

#2896: Normalization is scary

Cache key for QUERY requests includes a “normalized” version of the request body.

For certain input formats, reliable normalization is an intractable problem.

Do we need to change this or say more?

#2897 & #2901: Direct vs. indirect responses

Server can respond to QUERY with 2XX, results and optional Location/Content-Location fields (direct)

-OR-

Server can respond with 3XX and Location field, which client subsequently fetches (indirect)

These have different performance and caching implications.

- Any guidance for when to use each?
- Should clients or downstream caches be able to express a preference?
- Is query is executed on receipt of QUERY or receipt of GET?

#2898: CORS?

Probably nothing to say here, since CORS is not an HTTP protocol concept.

#2899: NORMALIZE method

Issue requests separate method of converting QUERY body to a GET URL without actually performing the query, as well as control over how query body is encoded in the URL.

Editors believe this issue can be closed with no action.