Digest Headers

(was: Resource Digests, was: RFC 3230)

HTTPWG Interim 2021-06

draft-ietf-httpbis-digest-headers

[last interim slides] [latest editor copy]

Since February 2021 Interim

Mainly editorial changes, sitting in the editor's copy.

Thanks for the reviews and feedback.

See

https://tools.ietf.org/rfcdiff?url1=https://tools.ietf.org/id/draft-ietf-httpbis-digest-headers-04.txt&url2=https://httpwg.org/http-extensions/draft-ietf-httpbis-digest-headers.txt

Dealing with old algorithms (3) #1377

Current

Algorithm	Status
sha-256	standard
sha-512	standard
md5	deprecated
sha	deprecated
unixsum	deprecated
unixcksum	deprecated
crc32c	deprecated
adler32	deprecated

- standard (fine to use)
- deprecated (MUST NOT use)

What does Digest in requests mean?

Several issues related to digest and requests

#970, #1005, #1357, #1366

There is interest in being able to send a checksum on the actual content bytes.

Differing schools of thought

1. Digest should always be computed on **payload data***. Ignore notion of representation and therefore partial representation. Asymmetry between Request and Response messages.

2. Digest always talks in terms of **complete representation**. There should be symmetry between Request and Response messages.

* Or its called “Content” now?

Proposal: two different fields

`Content-Digest`: new header, always computed on the message content in both requests and responses, like `Content-MD5` see #1543

`Digest`: computed on the complete representation data retaining consistency with RFC3230; can support future methods standardizing partial representations in requests; it is useful.

Proposal: two different headers

`Content-Digest` questions:

- do we need a `Want-Content-Digest` header, like `Want-Digest`?
- caveats just like `Content-MD5`?

Towards WGLC

WGLC is two steps away

- move on with two headers (Digest, Content-Digest) and close all request representation issues
- cleanup Digest Algorithms table (help needed)

Thanks!

Roberto Polli - robipolli@gmail.com

Lucas Pardue - lucaspardue.24.7@gmail.com

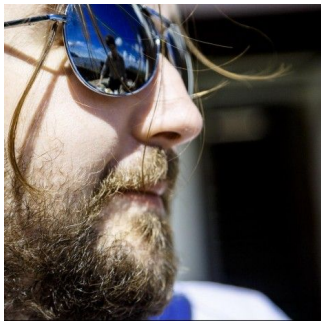

Example use case

A client uploads different ranges (chunks) of a large file. This supports a resumable upload model.

It wants to use a digest to help the upload process validate the integrity of [each chunk, the reassembled file, something else …]

Partial requests and http-semantics

Recently documented in http-semantics § 14.4 and 14.5

a. Servers MUST ignore Content-range in requests with a method **that does not define it**. “No request method in this specification is defined to support Content-Range”.
b. Partial PUT (with a Content-range) is supported by some. “though such support is inconsistent and depends on private agreements with user agents”

Digest of messages

Content-MD5 [RFC 1864 and 2616] used to allow digests of messages. “The Content-MD5 header field MAY be generated by an origin server or client to function as an integrity check of the entity-body”

RFC 7231 Appendix B - “The Content-MD5 header field has been **removed** because it was inconsistently implemented with respect to partial responses.”

Possible path forward

1. Digest applies to request representation.
2. No widespread standard usage of partial request. representation. De facto use of full representation, which is equivalent to payload data.
3. Some future Method can try to standardize a partial representation.
    a. It should probably also consider if an integrity check of the payload data is also useful. If so, new header.
4. Digest spec makes progress.

Backlog

id- prefix for digest-algorithms: should we strip id-sha-256? #885