

HTTP Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 26, 2017

P-H. Kamp
The Varnish Cache Project
April 24, 2017

HTTP Header Common Structure
draft-ietf-httpbis-header-structure-01

Abstract

An abstract data model for HTTP headers, "Common Structure", and a HTTP/1 serialization of it, generalized from current HTTP headers.

Note to Readers

Discussion of this draft takes place on the HTTP working group mailing list (ietf-http-wg@w3.org), which is archived at <https://lists.w3.org/Archives/Public/ietf-http-wg/> .

Working Group information can be found at <http://httpwg.github.io/> ; source code and issues list for this draft can be found at <https://github.com/httpwg/http-extensions/labels/header-structure> .

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on October 26, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents

(<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

1. Introduction

The HTTP protocol does not impose any structure or datamodel on the information in HTTP headers, the HTTP/1 serialization is the datamodel: An ASCII string without control characters.

HTTP header definitions specify how the string must be formatted and while families of similar headers exist, it still requires an uncomfortable large number of bespoke parser and validation routines to process HTTP traffic correctly.

In order to improve performance HTTP/2 and HPACK uses naive text-compression, which incidentally decoupled the on-the-wire serialization from the data model.

During the development of HPACK it became evident that significantly bigger gains were available if semantic compression could be used, most notably with timestamps. However, the lack of a common data structure for HTTP headers would make semantic compression one long list of special cases.

Parallel to this, various proposals for how to fulfill data-transportation needs, and to a lesser degree to impose some kind of order on HTTP headers, at least going forward, were floated.

All of these proposals, JSON, CBOR etc. run into the same basic problem: Their serialization is incompatible with [RFC 7230](#)'s [\[RFC7230\]](#) ABNF definition of 'field-value'.

For binary formats, such as CBOR, a wholesale base64/85 reserialization would be needed, with negative results for both debugability and bandwidth.

For textual formats, such as JSON, the format must first be neutered to not violate field-value's ABNF, and then workarounds added to

reintroduce the features just lost, for instance UNICODE strings.

The post-surgery format is no longer JSON, and it experience indicates that almost-but-not-quite compatibility is worse than no compatibility.

This proposal starts from the other end, and builds and generalizes a data structure definition from existing HTTP headers, which means that HTTP/1 serialization and 'field-value' compatibility is built in.

If all future HTTP headers are defined to fit into this Common Structure we have at least halted the proliferation of bespoke parsers and started to pave the road for semantic compression serializations of HTTP traffic.

1.1. Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [BCP 14](#), [RFC 2119](#) [[RFC2119](#)].

2. Definition of HTTP Header Common Structure

The data model of Common Structure is an ordered sequence of named dictionaries. Please see [Appendix A](#) for how this model was derived.

The definition of the data model is on purpose abstract, uncoupled from any protocol serialization or programming environment representation, it is meant as the foundation on which all such manifestations of the model can be built.

Common Structure in ABNF (Slightly bastardized relative to [RFC5234](#) [[RFC5234](#)]):

```
import token from RFC7230
import DIGIT from RFC5234

common-structure = 1* ( identifier dictionary )

dictionary = * ( identifier [ value ] )
```

```
value = identifier /
        integer /
        number /
        ascii-string /
        unicode-string /
        blob /
        timestamp /
        common-structure
```

Recursion is included as a way to support deep and more general data structures, but its use is highly discouraged and where it is

used the depth of recursion SHALL always be explicitly limited in the specifications of the HTTP headers which allow it.

```
identifier = token [ "/" token ]
```

```
integer = ["-"] 1*19 DIGIT
```

Integers SHALL be in the range +/- 2⁶³-1 (= +/- 9223372036854775807)

```
number = ["-"] DIGIT '.' 1*14DIGIT /
         ["-"] 2DIGIT '.' 1*13DIGIT /
         ["-"] 3DIGIT '.' 1*12DIGIT /
         ... /
         ["-"] 12DIGIT '.' 1*3DIGIT /
         ["-"] 13DIGIT '.' 1*2DIGIT /
         ["-"] 14DIGIT '.' 1DIGIT
```

The limit of 15 significant digits is chosen so that numbers can be correctly represented by IEEE754 64 bit binary floating point.

```
ascii-string = * %x20-7e
```

This is intended to be an efficient, "safe" and uncomplicated string type, for uses where the string content is culturally neutral or where it will not be user visible.

```
unicode-string = * UNICODE
```

```
UNICODE = <U+0000-U+D7FF / U+E000-U+10FFFF>
```

UNICODE nicked from [draft-seantek-unicode-in-abnf-02](#)

Unicode-strings are unrestricted because there is no sane and/or culturally neutral way to subset or otherwise make unicode "safe", and Unicode is still evolving new and interesting code points.

Users of unicode-string SHALL be prepared for the full gammut of glyph-gymnastics in order to avoid U+1F4A9 U+08 U+1F574.

```
blob = * %0x00-ff
```

Blobs are intended primarily for cryptographic data, but can be used for any otherwise unsatisfied needs.

```
timestamp = number
```

A timestamp counts seconds since the UNIX time_t epoch, including the "invisible leap-seconds" misfeature.

Kamp

Expires October 26, 2017

[Page 4]

Internet-Draft

HTTP Header Common Structure

April 2017

[3.](#) HTTP/1 Serialization of HTTP Header Common Structure

In ABNF:

```
import OWS from RFC7230
import HEXDIG, DQUOTE from RFC5234
import EmbeddedUnicodeChar from RFC5137

h1-common-structure-header =
    h1-common-structure-legacy-header /
    h1-common-structure-self-identifying-header

h1-common-structure-legacy-header =
    field-name ":" OWS h1-common-structure
```

Only white-listed legacy headers (see [Section 8](#)) can use this format.

```
h1-common-structure-self-identifying-header:
    field-name ":" OWS ">" h1-common-structure "<"

h1-common-structure = h1-element * ("," h1-element)

h1-element = identifier * (";" identifier ["=" h1-value])

h1-value = identifier /
    integer /
    number /
    h1-ascii-string /
    h1-unicode-string /
    h1-blob /
    h1-timestamp /
    ">" h1-common-structure "<"
```

```
h1-ascii-string = DQUOTE *(
    ( "\" DQUOTE ) /
    ( "\" "\"" ) /
    0x20-21 /
    0x23-5B /
    0x5D-7E
) DQUOTE
```

```
h1-unicode-string = DQUOTE *(
    ( "\" DQUOTE )
    ( "\" "\"" ) /
    EmbeddedUnicodeChar /
    0x20-21 /
    0x23-5B /
    0x5D-7E /
) DQUOTE
```

The dim prospects of ever getting a majority of HTTP1 paths 8-bit clean makes UTF-8 unviable as H1 serialization. Given that very little of the information in HTTP headers is presented to users in the first place, improving H1 and HPACK efficiency by inventing a more efficient [RFC5137](#) compliant escape-sequences seems unwarranted.

```
h1-blob = ":" base64 ":"
# XXX: where to import base64 from ?
```

```
h1-timestamp = number
```

XXX: Allow OWS in parsers, but not in generators ?

In programming environments which do not define a native representation or serialization of Common Structure, the HTTP/1 serialization should be used.

[4.](#) When to use Common Structure Parser

All future standardized and all private HTTP headers using Common Structure should self identify as such. In the HTTP/1 serialization

by making the first character ">" and the last "<". (These two characters are deliberately "the wrong way" to not clash with existing usages.)

Legacy HTTP headers which fit into Common Structure, are marked as such in the IANA Message Header Registry (see [Section 8](#)), and a snapshot of the registry can be used to trigger parsing according to Common Structure of these headers.

[5.](#) Desired Normative Effects

All new HTTP headers SHOULD use the Common Structure if at all possible.

[6.](#) Open/Outstanding issues to resolve

[6.1.](#) Single/Multiple Headers

Should we allow splitting common structure data over multiple headers ?

Pro:

Avoids size restrictions, easier on-the-fly editing

Contra:

Cannot act on any such header until all headers have been received.

We must define where headers can be split (between identifier and dictionary ?, in the middle of dictionaries ?)

Most on-the-fly editing is hackish at best.

[7.](#) Future Work

[7.1.](#) Redefining existing headers for better performance

The HTTP/1 serializations self-identification mechanism makes it possible to extend the definition of existing [Appendix A.5](#) headers into Common Structure.

For instance one could imagine:

```
Date: >1475061449.201<
```

Which would be faster to parse and validate than the current definition of the Date header and more precise too.

Some kind of signal/negotiation mechanism would be required to make this work in practice.

[7.2.](#) Define a validation dictionary

A machine-readable specification of the legal contents of HTTP headers would go a long way to improve efficiency and security in HTTP implementations.

[8.](#) IANA Considerations

The IANA Message Header Registry will be extended with an additional field named "Common Structure" which can have the values "True", "False" or "Unknown".

The RFC723x headers listed in [Appendix A.4](#) will get the value "True" in the new field.

The RFC723x headers listed in [Appendix A.5](#) will get the value "False" in the new field.

All other existing entries in the registry will be set to "Unknown" until and if the owner of the entry requests otherwise.

[9.](#) Security Considerations

Unique dictionary keys are required to reduce the risk of smuggling attacks.

[10.](#) References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5137] Klensin, J., "ASCII Escaping of Unicode Characters", [BCP 137](#), [RFC 5137](#), DOI 10.17487/RFC5137, February 2008, <<http://www.rfc-editor.org/info/rfc5137>>.
- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), DOI 10.17487/RFC5234, January 2008, <<http://www.rfc-editor.org/info/rfc5234>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", [RFC 7230](#), DOI 10.17487/RFC7230, June 2014, <<http://www.rfc-editor.org/info/rfc7230>>.

10.2. Informative References

- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", [RFC 7231](#), DOI 10.17487/RFC7231, June 2014, <<http://www.rfc-editor.org/info/rfc7231>>.
- [RFC7232] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests", [RFC 7232](#), DOI 10.17487/RFC7232, June 2014, <<http://www.rfc-editor.org/info/rfc7232>>.
- [RFC7233] Fielding, R., Ed., Lafon, Y., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Range Requests", [RFC 7233](#), DOI 10.17487/RFC7233, June 2014, <<http://www.rfc-editor.org/info/rfc7233>>.
- [RFC7234] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Caching", [RFC 7234](#), DOI 10.17487/RFC7234, June 2014, <<http://www.rfc-editor.org/info/rfc7234>>.
- [RFC7235] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Authentication", [RFC 7235](#), DOI 10.17487/RFC7235, June 2014,

<<http://www.rfc-editor.org/info/rfc7235>>.

Kamp

Expires October 26, 2017

[Page 9]

Internet-Draft

HTTP Header Common Structure

April 2017

[RFC7239] Petersson, A. and M. Nilsson, "Forwarded HTTP Extension", [RFC 7239](#), DOI 10.17487/RFC7239, June 2014, <<http://www.rfc-editor.org/info/rfc7239>>.

[RFC7694] Reschke, J., "Hypertext Transfer Protocol (HTTP) Client-Initiated Content-Encoding", [RFC 7694](#), DOI 10.17487/RFC7694, November 2015, <<http://www.rfc-editor.org/info/rfc7694>>.

[Appendix A](#). Do HTTP headers have any common structure ?

Several proposals have been floated in recent years to use some preexisting structured data serialization or other for HTTP headers, to impose some sanity.

None of these proposals have gained traction and no obvious candidate data serializations have been left unexamined.

This effort tries to tackle the question from the other side, by asking if there is a common structure in existing HTTP headers we can generalize for this purpose.

[A.1](#). Survey of HTTP header structure

The RFC723x family of HTTP/1 standards control 49 entries in the IANA Message Header Registry, and they share two common motifs.

The majority of RFC723x HTTP headers are lists. A few of them are ordered, ('Content-Encoding'), some are unordered ('Connection') and some are ordered by 'q=%f' weight parameters ('Accept')

In most cases, the list elements are some kind of identifier, usually derived from ABNF 'token' as defined by [[RFC7230](#)].

A subgroup of headers, mostly related to MIME, uses what one could call a 'qualified token'::

```
qualified-token = token-or-asterix [ "/" token-or-asterix ]
```

The second motif is parameterized list elements. The best known is the "q=0.5" weight parameter, but other parameters exist as well.

Generalizing from these motifs, our candidate "Common Structure" data model becomes an ordered list of named dictionaries.

In pidgin ABNF, ignoring white-space for the sake of clarity, the HTTP/1.1 serialization of Common Structure is something like:

Kamp

Expires October 26, 2017

[Page 10]

Internet-Draft

HTTP Header Common Structure

April 2017

token-or-asterix = token from [RFC7230](#), but also allowing "*"

qualified-token = token-or-asterix ["/" token-or-asterix]

field-name, see [RFC7230](#)

Common-Structure-Header = field-name ":" 1#named-dictionary

named-dictionary = qualified-token [*("; " param)]

param = token ["=" value]

value = we'll get back to this in a moment.

Nineteen out of the RFC723x's 48 headers, almost 40%, can already be parsed using this definition, and none the rest have requirements which could not be met by this data model. See [Appendix A.4](#) and [Appendix A.5](#) for the full survey details.

[A.2.](#) Survey of values in HTTP headers

Surveying the datatypes of HTTP headers, standardized as well as private, the following picture emerges:

[A.2.1.](#) Numbers

Integer and floating point are both used. Range and precision is mostly unspecified in controlling documents.

Scientific notation (9.192631770e9) does not seem to be used anywhere.

The ranges used seem to be minus several thousand to plus a couple of billions, the high end almost exclusively being POSIX `time_t` timestamps.

[A.2.2.](#) Timestamps

RFC723x text format, but POSIX `time_t` represented as integer or floating point is not uncommon. ISO8601 have also been spotted.

[A.2.3.](#) Strings

The vast majority are pure ASCII strings, with either no escapes, `%xx` URL-like escapes or C-style back-slash escapes, possibly with the addition of `\uxxxx` UNICODE escapes.

Kamp

Expires October 26, 2017

[Page 11]

Internet-Draft

HTTP Header Common Structure

April 2017

Where non-ASCII character sets are used, they are almost always implicit, rather than explicit. UTF8 and ISO-8859-1 seem to be most common.

[A.2.4.](#) Binary blobs

Often used for cryptographic data. Usually in base64 encoding, sometimes `""`-quoted more often not. base85 encoding is also seen, usually quoted.

[A.2.5.](#) Identifiers

Seems to almost always fit in the RFC723x 'token' definition.

[A.3.](#) Is this actually a useful thing to generalize ?

The number one wishlist item seems to be UNICODE strings, with a big side order of not having to write a new parser routine every time somebody comes up with a new header.

Having a common parser would indeed be a good thing, and having an underlying data model which makes it possible define a compressed serialization, rather than rely on serialization to text followed by text compression (ie: HPACK) seems like a good idea too.

However, when using a datamodel and a parser general enough to transport useful data, it will have to be followed by a validation step, which checks that the data also makes sense.

Today validation, such as it is, is often done by the bespoke parsers.

This then is probably where the next big potential for improvement lies:

Ideally a machine readable "data dictionary" which makes it possible to copy that text out of RFCs, run it through a code generator which spits out validation code which operates on the output of the common parser.

But history has been particularly unkind to that idea.

Most attempts studied as part of this effort, have sunk under complexity caused by reaching for generality, but where scope has been wisely limited, it seems to be possible.

So file that idea under "future work".

[A.4.](#) RFC723x headers with "common structure"

- o Accept [\[RFC7231\], Section 5.3.2](#)
- o Accept-Charset [\[RFC7231\], Section 5.3.3](#)
- o Accept-Encoding [\[RFC7231\], Section 5.3.4](#), [\[RFC7694\], Section 3](#)
- o Accept-Language [\[RFC7231\], Section 5.3.5](#)
- o Age [\[RFC7234\], Section 5.1](#)
- o Allow [\[RFC7231\], Section 7.4.1](#)
- o Connection [\[RFC7230\], Section 6.1](#)
- o Content-Encoding [\[RFC7231\], Section 3.1.2.2](#)

- o Content-Language [\[RFC7231\], Section 3.1.3.2](#)
- o Content-Length [\[RFC7230\], Section 3.3.2](#)
- o Content-Type [\[RFC7231\], Section 3.1.1.5](#)
- o Expect [\[RFC7231\], Section 5.1.1](#)
- o Max-Forwards [\[RFC7231\], Section 5.1.2](#)
- o MIME-Version [\[RFC7231\], Appendix A.1](#)
- o TE [\[RFC7230\], Section 4.3](#)
- o Trailer [\[RFC7230\], Section 4.4](#)
- o Transfer-Encoding [\[RFC7230\], Section 3.3.1](#)
- o Upgrade [\[RFC7230\], Section 6.7](#)
- o Vary [\[RFC7231\], Section 7.1.4](#)

A.5. RFC723x headers with "uncommon structure"

1 of the RFC723x headers is only reserved, and therefore have no structure at all:

- o Close [\[RFC7230\], Section 8.1](#)

5 of the RFC723x headers are HTTP dates:

- o Date [\[RFC7231\], Section 7.1.1.2](#)
- o Expires [\[RFC7234\], Section 5.3](#)
- o If-Modified-Since [\[RFC7232\], Section 3.3](#)
- o If-Unmodified-Since [\[RFC7232\], Section 3.4](#)
- o Last-Modified [\[RFC7232\], Section 2.2](#)

24 of the RFC723x headers use bespoke formats which only a single or

in rare cases two headers share:

- o Accept-Ranges [\[RFC7233\], Section 2.3](#)
 - * bytes-unit / other-range-unit
- o Authorization [\[RFC7235\], Section 4.2](#)
- o Proxy-Authorization [\[RFC7235\], Section 4.4](#)
 - * credentials
- o Cache-Control [\[RFC7234\], Section 5.2](#)
 - * 1#cache-directive
- o Content-Location [\[RFC7231\], Section 3.1.4.2](#)
 - * absolute-URI / partial-URI
- o Content-Range [\[RFC7233\], Section 4.2](#)
 - * byte-content-range / other-content-range
- o ETag [\[RFC7232\], Section 2.3](#)
 - * entity-tag
- o Forwarded [\[RFC7239\]](#)
 - * 1#forwarded-element
- o From [\[RFC7231\], Section 5.5.1](#)
 - * mailbox
- o If-Match [\[RFC7232\], Section 3.1](#)

- o If-None-Match [\[RFC7232\], Section 3.2](#)
 - * "*" / 1#entity-tag

- o If-Range [\[RFC7233\], Section 3.2](#)
 - * entity-tag / HTTP-date
- o Host [\[RFC7230\], Section 5.4](#)
 - * uri-host [":" port]
- o Location [\[RFC7231\], Section 7.1.2](#)
 - * URI-reference
- o Pragma [\[RFC7234\], Section 5.4](#)
 - * 1#pragma-directive
- o Range [\[RFC7233\], Section 3.1](#)
 - * byte-ranges-specifier / other-ranges-specifier
- o Referer [\[RFC7231\], Section 5.5.2](#)
 - * absolute-URI / partial-URI
- o Retry-After [\[RFC7231\], Section 7.1.3](#)
 - * HTTP-date / delay-seconds
- o Server [\[RFC7231\], Section 7.4.2](#)
- o User-Agent [\[RFC7231\], Section 5.5.3](#)
 - * product *(RWS (product / comment))
- o Via [\[RFC7230\], Section 5.7.1](#)
 - * 1#(received-protocol RWS received-by [RWS comment])
- o Warning [\[RFC7234\], Section 5.5](#)
 - * 1#warning-value
- o Proxy-Authenticate [\[RFC7235\], Section 4.3](#)

- o WWW-Authenticate [\[RFC7235\], Section 4.1](#)
 - * 1#challenge

[Appendix B](#). Changes

[B.1](#). Since [draft-ietf-httpbis-header-structure-00](#)

Added signed 64bit integer type.

Drop UTF8, and settle on [BCP137](#) [\[RFC5137\]](#)::EmbeddedUnicodeChar for h1-unicode-string.

Change h1_blob delimiter to ":" since "'" is valid t_char

Author's Address

Poul-Henning Kamp
The Varnish Cache Project

Email: phk@varnish-cache.org

