

HTTP
Internet-Draft
Intended status: Standards Track
Expires: September 5, 2018

M. Nottingham
Fastly
P-H. Kamp
The Varnish Cache Project
March 4, 2018

Structured Headers for HTTP
draft-ietf-httpbis-header-structure-04

Abstract

This document describes a set of data types and parsing algorithms associated with them that are intended to make it easier and safer to define and handle HTTP header fields. It is intended for use by new specifications of HTTP header fields as well as revisions of existing header field specifications when doing so does not cause interoperability issues.

Note to Readers

RFC EDITOR: please remove this section before publication

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/> [1].

Working Group information can be found at <https://httpwg.github.io/> [2]; source code and issues list for this draft can be found at <https://github.com/httpwg/http-extensions/labels/header-structure> [3].

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 <https://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 September 5, 2018.

Internet-Draft

Structured Headers for HTTP

March 2018

Copyright Notice

Copyright (c) 2018 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 (<https://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.

Table of Contents

1.	Introduction	3
1.1.	Notational Conventions	4
2.	Specifying Structured Headers	4
3.	Parsing Text into Structured Headers	5
4.	Structured Header Data Types	6
4.1.	Dictionaries	6
4.2.	Lists	8
4.3.	Parameterised Lists	9
4.4.	Items	11
4.5.	Integers	11
4.6.	Floats	12
4.7.	Strings	13
4.8.	Identifiers	15
4.9.	Binary Content	16
5.	IANA Considerations	17
6.	Security Considerations	17
7.	References	17
7.1.	Normative References	17
7.2.	Informative References	18
7.3.	URIs	18
Appendix A.	Changes	18
A.1.	Since draft-ietf-httpbis-header-structure-03	18
A.2.	Since draft-ietf-httpbis-header-structure-02	18
A.3.	Since draft-ietf-httpbis-header-structure-01	19
A.4.	Since draft-ietf-httpbis-header-structure-00	19
	Authors' Addresses	19

1. Introduction

Specifying the syntax of new HTTP header fields is an onerous task; even with the guidance in [\[RFC7231\], Section 8.3.1](#), there are many decisions - and pitfalls - for a prospective HTTP header field author.

Once a header field is defined, bespoke parsers for it often need to be written, because each header has slightly different handling of what looks like common syntax.

This document introduces structured HTTP header field values (hereafter, Structured Headers) to address these problems. Structured Headers define a generic, abstract model for header field values, along with a concrete serialisation for expressing that model in textual HTTP headers, as used by HTTP/1 [\[RFC7230\]](#) and HTTP/2 [\[RFC7540\]](#).

HTTP headers that are defined as Structured Headers use the types defined in this specification to define their syntax and basic handling rules, thereby simplifying both their definition and parsing.

Additionally, future versions of HTTP can define alternative serialisations of the abstract model of Structured Headers, allowing headers that use it to be transmitted more efficiently without being redefined.

Note that it is not a goal of this document to redefine the syntax of existing HTTP headers; the mechanisms described herein are only intended to be used with headers that explicitly opt into them.

To specify a header field that uses Structured Headers, see [Section 2](#).

[Section 4](#) defines a number of abstract data types that can be used in

Structured Headers. Dictionaries and lists are only usable at the "top" level, while the remaining types can be specified appear at the top level or inside those structures.

Those abstract types can be serialised into textual headers - such as those used in HTTP/1 and HTTP/2 - using the algorithms described in [Section 3](#).

[1.1](#). Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

This document uses the Augmented Backus-Naur Form (ABNF) notation of [[RFC5234](#)], including the DIGIT, ALPHA and DQUOTE rules from that document. It also includes the OWS rule from [[RFC7230](#)].

[2](#). Specifying Structured Headers

A HTTP header that uses Structured Headers need to be defined to do so explicitly; recipients and generators need to know that the requirements of this document are in effect. The simplest way to do that is by referencing this document in its definition.

The field's definition will also need to specify the field-value's allowed syntax, in terms of the types described in [Section 4](#), along with their associated semantics.

A header field definition cannot relax or otherwise modify the requirements of this specification; doing so would preclude handling by generic software.

However, header field authors are encouraged to clearly state additional constraints upon the syntax, as well as the consequences

when those constraints are violated. Such additional constraints could include additional structure (e.g., a list of URLs [[RFC3986](#)] inside a string) that cannot be expressed using the primitives defined here.

For example:

Nottingham & Kamp Expires September 5, 2018 [Page 4]

Internet-Draft Structured Headers for HTTP March 2018

FooExample Header

The FooExample HTTP header field conveys a list of integers about how much Foo the sender has.

FooExample is a Structured header [RFCxxxx]. Its value MUST be a dictionary ([RFCxxxx], Section Y.Y).

The dictionary MUST contain:

- * Exactly one member whose key is "foo", and whose value is an integer ([RFCxxxx], Section Y.Y), indicating the number of foos in the message.
- * Exactly one member whose key is "barUrls", and whose value is a string ([RFCxxxx], Section Y.Y), conveying the Bar URLs for the message. See below for processing requirements.

If the parsed header field does not contain both, it MUST be ignored.

"foo" MUST be between 0 and 10, inclusive; other values MUST be ignored.

"barUrls" contains a space-separated list of URI-references ([RFC3986](#)),

[Section 4.1](#)):

barURLs = URI-reference *(1*SP URI-reference)

If a member of barURLs is not a valid URI-reference, it MUST be ignored.

If a member of barURLs is a relative reference ([\[RFC3986\], Section 4.2](#)), it MUST be resolved ([\[RFC3986\], Section 5](#)) before being used.

Note that empty header field values are not allowed by the syntax, and therefore parsing for them will fail.

[3](#). Parsing Text into Structured Headers

When a receiving implementation parses textual HTTP header fields (e.g., in HTTP/1 or HTTP/2) that are known to be Structured Headers, it is important that care be taken, as there are a number of edge cases that can cause interoperability or even security problems. This section specifies the algorithm for doing so.

Given an ASCII string `input_string` that represents the chosen header's field-value, return the parsed header value. When generating `input_string`, parsers MUST combine all instances of the target header field into one comma-separated field-value, as per [\[RFC7230\], Section 3.2.2](#); this assures that the header is processed correctly.

1. Discard any leading OWS from `input_string`.
2. If the field-value is defined to be a dictionary, let output be the result of Parsing a Dictionary from Text ([Section 4.1.1](#)).
3. If the field-value is defined to be a list, let output be the result of Parsing a List from Text ([Section 4.2.1](#)).
4. If the field-value is defined to be a parameterised list, let output be the result of Parsing a Parameterised List from Text ([Section 4.3.1](#)).
5. Otherwise, let output be the result of Parsing an Item from Text ([Section 4.4.1](#)).

6. Discard any leading OWS from `input_string`.
7. If `input_string` is not empty, fail parsing.
8. Otherwise, return output.

Note that in the case of lists, parameterised lists and dictionaries, this has the effect of coalescing all of the values for that field. However, for singular items, parsing will fail if more than instance of that header field is present.

If parsing fails, the entire header field's value MUST be discarded. This is intentionally strict, to improve interoperability and safety, and specifications referencing this document MUST NOT loosen this requirement.

Note that this has the effect of discarding any header field with non-ASCII characters in `input_string`.

[4.](#) Structured Header Data Types

This section defines the abstract value types that can be composed into Structured Headers, along with the textual HTTP serialisations of them.

[4.1.](#) Dictionaries

Dictionaries are unordered maps of key-value pairs, where the keys are identifiers ([Section 4.8](#)) and the values are items ([Section 4.4](#)). There can be between 1 and 1024 members, and keys are required to be unique.

In the textual HTTP serialisation, keys and values are separated by "=" (without whitespace), and key/value pairs are separated by a comma with optional whitespace. Duplicate keys MUST cause parsing to fail.

```
dictionary          = dictionary_member *1023( OWS "," OWS dictionary_member )
dictionary_member  = identifier "=" item
```

For example, a header field whose value is defined as a dictionary could look like:

```
ExampleDictHeader: foo=1.23, en="Applepie", da=*w4ZibGV0w6ZydGUK
```

Typically, a header field specification will define the semantics of individual keys, as well as whether their presence is required or optional. Recipients **MUST** ignore keys that are undefined or unknown, unless the header field's specification specifically disallows them.

[4.1.1](#). Parsing a Dictionary from Text

Given an ASCII string `input_string`, return a mapping of (identifier, item). `input_string` is modified to remove the parsed value.

1. Let `dictionary` be an empty, unordered mapping.
2. While `input_string` is not empty:
 1. Let `this_key` be the result of running Parse Identifier from Text ([Section 4.8.1](#)) with `input_string`.
 2. If `dictionary` already contains `this_key`, fail parsing.
 3. Consume a "=" from `input_string`; if none is present, fail parsing.
 4. Let `this_value` be the result of running Parse Item from Text ([Section 4.4.1](#)) with `input_string`.
 5. Add key `this_key` with value `this_value` to `dictionary`.
 6. If `dictionary` has more than 1024 members, fail parsing.
 7. Discard any leading OWS from `input_string`.
 8. If `input_string` is empty, return `dictionary`.
 9. Consume a COMMA from `input_string`; if no comma is present, fail parsing.

10. Discard any leading OWS from `input_string`.

11. If `input_string` is empty, fail parsing.
3. If `dictionary` is empty, fail parsing.
4. Return `dictionary`.

[4.2.](#) Lists

Lists are arrays of items ([Section 4.4](#)) with one to 1024 members.

In the textual HTTP serialisation, each member is separated by a comma and optional whitespace.

```
list = list_member 0*1023( OWS "," OWS list_member )
list_member = item
```

For example, a header field whose value is defined as a list of identifiers could look like:

```
ExampleIdListHeader: foo, bar, baz_45
```

[4.2.1.](#) Parsing a List from Text

Given an ASCII string `input_string`, return a list of items. `input_string` is modified to remove the parsed value.

1. Let `items` be an empty array.
2. While `input_string` is not empty:
 1. Let `item` be the result of running Parse Item from Text ([Section 4.4.1](#)) with `input_string`.
 2. Append `item` to `items`.
 3. If `items` has more than 1024 members, fail parsing.
 4. Discard any leading OWS from `input_string`.
 5. If `input_string` is empty, return `items`.
 6. Consume a COMMA from `input_string`; if no comma is present, fail parsing.
 7. Discard any leading OWS from `input_string`.

8. If `input_string` is empty, fail parsing.
3. If `items` is empty, fail parsing.
4. Return `items`.

[4.3.](#) Parameterised Lists

Parameterised Lists are arrays of a parameterised identifiers with 1 to 256 members.

A parameterised identifier is an identifier ([Section 4.8](#)) with up to 256 parameters, each parameter having a identifier and an optional value that is an item ([Section 4.4](#)). Ordering between parameters is not significant, and duplicate parameters MUST cause parsing to fail.

In the textual HTTP serialisation, each parameterised identifier is separated by a comma and optional whitespace. Parameters are delimited from each other using semicolons (";"), and equals ("=") delimits the parameter name from its value.

```
param_list = param_id 0*255( OWS "," OWS param_id )
param_id   = identifier 0*256( OWS ";" OWS identifier [ "=" item ] )
```

For example,

```
ExampleParamListHeader: abc_123;a=1;b=2; c, def_456, ghi;q="19";r=foo
```

[4.3.1.](#) Parsing a Parameterised List from Text

Given an ASCII string `input_string`, return a list of parameterised identifiers. `input_string` is modified to remove the parsed value.

1. Let `items` be an empty array.
2. While `input_string` is not empty:
 1. Let `item` be the result of running Parse Parameterised Identifier from Text ([Section 4.3.2](#)) with `input_string`.
 2. Append `item` to `items`.
 3. If `items` has more than 256 members, fail parsing.
 4. Discard any leading OWS from `input_string`.

5. If `input_string` is empty, return items.

6. Consume a COMMA from `input_string`; if no comma is present, fail parsing.
7. Discard any leading OWS from `input_string`.
8. If `input_string` is empty, fail parsing.
3. If `items` is empty, fail parsing.
4. Return items.

[4.3.2.](#) Parsing a Parameterised Identifier from Text

Given an ASCII string `input_string`, return a identifier with an mapping of parameters. `input_string` is modified to remove the parsed value.

1. Let `primary_identifier` be the result of Parsing a Identifier from Text ([Section 4.8.1](#)) from `input_string`.
2. Let `parameters` be an empty, unordered mapping.
3. In a loop:
 1. Discard any leading OWS from `input_string`.
 2. If the first character of `input_string` is not ";", exit the loop.
 3. Consume a ";" character from the beginning of `input_string`.
 4. Discard any leading OWS from `input_string`.
 5. let `param_name` be the result of Parsing a Identifier from Text ([Section 4.8.1](#)) from `input_string`.
 6. If `param_name` is already present in `parameters`, fail parsing.

7. Let param_value be a null value.
8. If the first character of input_string is "=":
 1. Consume the "=" character at the beginning of input_string.
 2. Let param_value be the result of Parsing an Item from Text ([Section 4.4.1](#)) from input_string.

9. If parameters has more than 255 members, fail parsing.
10. Add param_name to parameters with the value param_value.
4. Return the tuple (primary_identifier, parameters).

[4.4.](#) Items

An item is can be a integer ([Section 4.5](#)), float ([Section 4.6](#)), string ([Section 4.7](#)), identifier ([Section 4.8](#)) or binary content ([Section 4.9](#)).

item = integer / float / string / identifier / binary

[4.4.1.](#) Parsing an Item from Text

Given an ASCII string input_string, return an item. input_string is modified to remove the parsed value.

1. Discard any leading OWS from input_string.
2. If the first character of input_string is a "-" or a DIGIT, process input_string as a number ([Section 4.5.1](#)) and return the result.
3. If the first character of input_string is a DQUOTE, process input_string as a string ([Section 4.7.1](#)) and return the result.
4. If the first character of input_string is "*", process input_string as binary content ([Section 4.9.1](#)) and return the result.

5. If the first character of `input_string` is an lcalpha, process `input_string` as a identifier ([Section 4.8.1](#)) and return the result.
6. Otherwise, fail parsing.

[4.5.](#) Integers

Abstractly, integers have a range of $-9,223,372,036,854,775,808$ to $9,223,372,036,854,775,807$ inclusive (i.e., a 64-bit signed integer).

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

Parsers that encounter an integer outside the range defined above MUST fail parsing. Therefore, the value "9223372036854775808" would

be invalid. Likewise, values that do not conform to the ABNF above are invalid, and MUST fail parsing.

For example, a header whose value is defined as a integer could look like:

```
ExampleIntegerHeader: 42
```

[4.5.1.](#) Parsing a Number from Text

NOTE: This algorithm parses both Integers and Floats [Section 4.6](#), and returns the corresponding structure.

1. If the first character of `input_string` is not "-" or a DIGIT, fail parsing.
2. Let `input_number` be the result of consuming `input_string` up to (but not including) the first character that is not in DIGIT, "-", and ".".
3. If `input_number` contains ".", parse it as a floating point number and let `output_number` be the result.
4. Otherwise, parse `input_number` as an integer and let `output_number` be the result.

5. Return output_number.

[4.6.](#) Floats

Abstractly, floats are integers with a fractional part. They have a maximum of fifteen digits available to be used in both of the parts, as reflected in the ABNF below; this allows them to be stored as IEEE 754 double precision numbers (binary64) ([\[IEEE754\]](#)).

The textual HTTP serialisation of floats allows a maximum of fifteen digits between the integer and fractional part, with at least one required on each side, along with an optional "-" indicating negative numbers.

```
float    = ["-"] (
    DIGIT "." 1*14DIGIT /
    2DIGIT "." 1*13DIGIT /
    3DIGIT "." 1*12DIGIT /
    4DIGIT "." 1*11DIGIT /
    5DIGIT "." 1*10DIGIT /
    6DIGIT "." 1*9DIGIT /
    7DIGIT "." 1*8DIGIT /
    8DIGIT "." 1*7DIGIT /
    9DIGIT "." 1*6DIGIT /
    10DIGIT "." 1*5DIGIT /
    11DIGIT "." 1*4DIGIT /
    12DIGIT "." 1*3DIGIT /
    13DIGIT "." 1*2DIGIT /
    14DIGIT "." 1DIGIT )
```

Values that do not conform to the ABNF above are invalid, and MUST fail parsing.

For example, a header whose value is defined as a float could look like:

```
ExampleFloatHeader: 4.5
```

See [Section 4.5.1](#) for the parsing algorithm for floats.

4.7. Strings

Abstractly, strings are up to 1024 printable ASCII [[RFC0020](#)] characters (i.e., the range 0x20 to 0x7E). Note that this excludes tabs, newlines and carriage returns.

The textual HTTP serialisation of strings uses a backslash ("\") to escape double quotes and backslashes in strings.

```
string    = DQUOTE 0*1024(char) DQUOTE
char      = unescaped / escape ( DQUOTE / "\" )
unescaped = %x20-21 / %x23-5B / %x5D-7E
escape    = "\"
```

For example, a header whose value is defined as a string could look like:

```
ExampleStringHeader: "hello world"
```

Note that strings only use DQUOTE as a delimiter; single quotes do not delimit strings. Furthermore, only DQUOTE and "\" can be escaped; other sequences MUST cause parsing to fail.

Unicode is not directly supported in Structured Headers, because it causes a number of interoperability issues, and - with few exceptions - header values do not require it.

When it is necessary for a field value to convey non-ASCII string content, binary content ([Section 4.9](#)) SHOULD be specified, along with a character encoding (preferably, UTF-8).

4.7.1. Parsing a String from Text

Given an ASCII string `input_string`, return an unquoted string.

input_string is modified to remove the parsed value.

1. Let output_string be an empty string.
2. If the first character of input_string is not DQUOTE, fail parsing.
3. Discard the first character of input_string.
4. While input_string is not empty:
 1. Let char be the result of removing the first character of input_string.
 2. If char is a backslash ("\"):
 1. If input_string is now empty, fail parsing.
 2. Else:
 1. Let next_char be the result of removing the first character of input_string.
 2. If next_char is not DQUOTE or "\", fail parsing.
 3. Append next_char to output_string.
 3. Else, if char is DQUOTE, return output_string.
 4. Else, append char to output_string.
 5. If output_string contains more than 1024 characters, fail parsing.
5. Otherwise, fail parsing.

[4.8.](#) Identifiers

Identifiers are short (up to 256 characters) textual identifiers; their abstract model is identical to their expression in the textual

HTTP serialisation.

```
identifier = lcalpha *255( lcalpha / DIGIT / "_" / "-" / "*" / "/" )  
lcalpha    = %x61-7A ; a-z
```

Note that identifiers can only contain lowercase letters.

For example, a header whose value is defined as a identifier could look like:

ExampleIdHeader: foo/bar

[4.8.1.](#) Parsing a Identifier from Text

Given an ASCII string `input_string`, return a identifier. `input_string` is modified to remove the parsed value.

1. If the first character of `input_string` is not `lcalpha`, fail parsing.
2. Let `output_string` be an empty string.
3. While `input_string` is not empty:
 1. Let `char` be the result of removing the first character of `input_string`.
 2. If `char` is not one of `lcalpha`, `DIGIT`, `"_"`, `"-"`, `"*"` or `"/"`:
 1. Prepend `char` to `input_string`.
 2. Return `output_string`.
 3. Append `char` to `output_string`.
 4. If `output_string` contains more than 256 characters, fail parsing.
4. Return `output_string`.

[4.9.](#) Binary Content

Arbitrary binary content up to 16384 bytes in size can be conveyed in Structured Headers.

The textual HTTP serialisation encodes the data using Base 64 Encoding [\[RFC4648\], Section 4](#), and surrounds it with a pair of asterisks ("`*`") to delimit from other content.

The encoded data is required to be padded with "=", as per [\[RFC4648\], Section 3.2](#). It is RECOMMENDED that parsers reject encoded data that is not properly padded, although this might not be possible with some base64 implementations.

Likewise, encoded data is required to have pad bits set to zero, as per [\[RFC4648\], Section 3.5](#). It is RECOMMENDED that parsers fail on encoded data that has non-zero pad bits, although this might not be possible with some base64 implementations.

This specification does not relax the requirements in [\[RFC4648\], Section 3.1](#) and 3.3; therefore, parsers MUST fail on characters outside the base64 alphabet, and on line feeds in encoded data.

```
binary = "*" 0*21846(base64) "*"
base64 = ALPHA / DIGIT / "+" / "/" / "="
```

For example, a header whose value is defined as binary content could look like:

```
ExampleBinaryHeader: *cHJldGVuZCB0aGlzIGlzIGJpbmFyeSBjb250ZW50Lg*
```

[4.9.1.](#) Parsing Binary Content from Text

Given an ASCII string `input_string`, return binary content.
`input_string` is modified to remove the parsed value.

1. If the first character of `input_string` is not "*", fail parsing.
2. Discard the first character of `input_string`.
3. Let `b64_content` be the result of removing content of `input_string` up to but not including the first instance of the character "*".
If there is not a "*" character before the end of `input_string`, fail parsing.
4. Consume the "*" character at the beginning of `input_string`.

5. If `b64_content` is has more than 21846 characters, fail parsing.

6. Let `binary_content` be the result of Base 64 Decoding [[RFC4648](#)] `b64_content`, synthesising padding if necessary (note the requirements about recipient behaviour in [Section 4.9](#)).
7. Return `binary_content`.

[5.](#) IANA Considerations

This draft has no actions for IANA.

[6.](#) Security Considerations

TBD

[7.](#) References

[7.1.](#) Normative References

- [RFC0020] Cerf, V., "ASCII format for network interchange", STD 80, [RFC 20](#), DOI 10.17487/RFC0020, October 1969, <<https://www.rfc-editor.org/info/rfc20>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", [RFC 4648](#), DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/info/rfc4648>>.
- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), DOI 10.17487/RFC5234, January 2008, <<https://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, <<https://www.rfc-editor.org/info/rfc7230>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

Nottingham & Kamp

Expires September 5, 2018

[Page 17]

Internet-Draft

Structured Headers for HTTP

March 2018

[7.2.](#) Informative References

[IEEE754] IEEE, "IEEE Standard for Floating-Point Arithmetic", 2008, <<http://grouper.ieee.org/groups/754/>>.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.

[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, <<https://www.rfc-editor.org/info/rfc7231>>.

[RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", [RFC 7540](#), DOI 10.17487/RFC7540, May 2015, <<https://www.rfc-editor.org/info/rfc7540>>.

[7.3.](#) URIs

[1] <https://lists.w3.org/Archives/Public/ietf-http-wg/>

[2] <https://httpwg.github.io/>

[3] <https://github.com/httpwg/http-extensions/labels/header-structure>

[Appendix A.](#) Changes

[A.1.](#) Since [draft-ietf-httpbis-header-structure-03](#)

- o Strengthen language around failure handling.

[A.2.](#) Since [draft-ietf-httpbis-header-structure-02](#)

- o Split Numbers into Integers and Floats.
- o Define number parsing.
- o Tighten up binary parsing and give it an explicit end delimiter.
- o Clarify that mappings are unordered.
- o Allow zero-length strings.
- o Improve string parsing algorithm.

Nottingham & Kamp

Expires September 5, 2018

[Page 18]

Internet-Draft

Structured Headers for HTTP

March 2018

- o Improve limits in algorithms.
- o Require parsers to combine header fields before processing.
- o Throw an error on trailing garbage.

[A.3.](#) Since [draft-ietf-httpbis-header-structure-01](#)

- o Replaced with [draft-nottingham-structured-headers](#).

[A.4.](#) Since [draft-ietf-httpbis-header-structure-00](#)

- o Added signed 64bit integer type.
- o Drop UTF8, and settle on [BCP137](#) ::EmbeddedUnicodeChar for h1-unicode-string.
- o Change h1_blob delimiter to ":" since "" is valid t_char

Authors' Addresses

Mark Nottingham
Fastly

Email: mnot@mnot.net

URI: <https://www.mnot.net/>

Poul-Henning Kamp
The Varnish Cache Project

Email: phk@varnish-cache.org