

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: December 29, 2014

A. Hutton
Unify
J. Uberti
Google
M. Thomson
Mozilla
June 27, 2014

HTTP Connect - Tunnel Protocol For WebRTC
draft-hutton-httpbis-connect-protocol-00

Abstract

This document describes a mechanism to enable HTTP Clients to provide an indication within a HTTP Connect request as to which protocol will be used within the tunnel established to the Server identified by the target resource. The tunneled protocol is declared using the Tunnel-Protocol HTTP Request header field. Label usage relating to the use of HTTP Connect by WebRTC clients (e.g. turn, webrtc) are described in this document.

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 December 29, 2014.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
2. Use Cases	3
3. The Tunnel-Protocol HTTP Request Header Field	3
3.1. Header Field Values	3
3.2. Syntax	4
3.3. TURN as the Tunnel Protocol	4
3.4. ICE-TCP / WebRTC as the Tunnel Protocol	4
4. IANA Considerations	5
5. Security Considerations	5
6. References	5
6.1. Normative References	5
6.2. Informative References	6
Authors' Addresses	6

1. Introduction

The HTTP Connect method (Section 4.3.6 of [RFC7231]) requests that the recipient establish a tunnel to the destination origin server identified by the request-target and thereafter forward packets, in both directions, until the tunnel is closed. Such tunnels are commonly used to create end-to-end virtual connections, through one or more proxies, which may then be secured using TLS (Transport Layer Security, [RFC5246]).

The RTCWEB use cases and requirements document [I-D.ietf-rtcweb-use-cases-and-requirements] includes a requirement that a WebRTC Client must be able to send streams and data to a peer in the presence of Firewalls that only allow traffic via a HTTP Proxy, when Firewall policy allows WebRTC traffic. To facilitate this and to allow such a HTTP Proxy to be provided with an indication that WebRTC related real-time media is to be included in the tunnel this specification defines the Tunnel-Protocol Request header field and associated labels. This allows the proxy to identify the protocol being used in the tunnel as early as possible therefore enabling the proxy to make informed policy decisions. The type of policy decisions the proxy may make is not specified here but may include rejecting the request with a HTTP status code responses or prioritizing connections. As described in Section 4.3.6 of [RFC7231]

and 2xx response indicates consent for the client to switch to tunnel mode.

The HTTP Tunnel-Protocol header field may be used in conjunction with and complements the application layer next protocol extension [I-D.ietf-tls-applayerprotoneg] specified for TLS [RFC5246]". In the scenario where the HTTP Connect is used to establish a TLS tunnel then the HTTP Tunnel-Protocol may be used to carry the same next protocol label as carried within the TLS handshake. However, the Tunnel-Protocol is an indication rather a negotiation since the HTTP Proxy does not implement the tunneled protocol. ALPN Labels are already defined for TURN in [I-D.patil-tram-alpn] and WebRTC [I-D.thomson-rtcweb-alpn] and are re-used here.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Use Cases

The following two use cases are considered:

- o The WebRTC Client issues a HTTP CONNECT request to the HTTP proxy with the TURN server address in the Request URI.
- o The WebRTC Client issues a HTTP CONNECT request to the HTTP proxy with the TCP address of a WebRTC peer in the Request URI. This is used in the case of establishing ICE-TCP [RFC6544] with a WebRTC Peer.

3. The Tunnel-Protocol HTTP Request Header Field

The client MAY include the Tunnel-Protocol Request Header field in a HTTP Connect request to indicate the application layer protocol within the tunnel.

3.1. Header Field Values

Valid values for the protocol field are taken from the registry established in [I-D.ietf-tls-applayerprotoneg]. For the purposes of WebRTC, the values "webrtc" [I-D.thomson-rtcweb-alpn] and "turn" [I-D.patil-tram-alpn] are applicable.

3.2. Syntax

The ABNF (Augmented Backus-Naur Form) syntax for the Tunnel-Protocol header field is given below. It is based on the Generic Grammar defined in Section 2 of [RFC7230].

```
Tunnel-Protocol = "Tunnel-Protocol" ":" protocol | protocol-extension
```

```
protocol = "webrtc" | "turn"
```

```
protocol-extension = token
```

3.3. TURN as the Tunnel Protocol

The RTCWEB transports specification [I-D.ietf-rtcweb-transports] requires that a WebRTC client support the modes of TURN that uses TCP and TLS between the client and the TURN server in order to deal with firewalls blocking UDP traffic. In the case where HTTP Connect is used to establish a tunnel to the TURN server the client SHOULD include the "Tunnel-Protocol" header field with the value "turn" [I-D.patil-tram-alpn] as shown in the example below.

```
CONNECT turn_server.example.com:5349 HTTP/1.1
Host: turn_server.example.com:5349
Tunnel-Protocol: turn
```

3.4. ICE-TCP / WebRTC as the Tunnel Protocol

[I-D.ietf-rtcweb-transports] also requires that a WebRTC client support ICE-TCP [RFC6544] as a mechanism to allow webrtc applications to communicate to peers with public IP addresses across UDP-blocking firewalls without using a TURN server. In this case the client SHOULD include the "Tunnel-Protocol" header field with the value "webrtc" [I-D.thomson-rtcweb-alpn] as shown in the example below.

```
CONNECT 198.51.100.0:8999 HTTP/1.1
Host: 198.51.100.0:8999
Tunnel-Protocol: webrtc
```

Note: The protocol "c_webrtc" described in [I-D.thomson-rtcweb-alpn] is not relevant in this context and when used at the TLS layer the client SHOULD use "webrtc" in the Tunnel-Protocol header. OPEN ISSUE - Is this correct?

4. IANA Considerations

To Be Added

5. Security Considerations

In case of using HTTP CONNECT to a TURN server the security consideration of [RFC7231], Section-4.3.6] apply. It states that there "are significant risks in establishing a tunnel to arbitrary servers, particularly when the destination is a well-known or reserved TCP port that is not intended for Web traffic. Proxies that support CONNECT SHOULD restrict its use to a limited set of known ports or a configurable whitelist of safe request targets."

The Tunnel-Protocol request header field described in this document is an optional header and HTTP Proxies may of course not support the header and therefore ignore it. If the header is not present or ignored then the proxy has no explicit indication as to the purpose of the tunnel on which to provide consent, this is the generic case that exists without the Tunnel-Protocol header.

6. References

6.1. Normative References

[I-D.patil-tram-alpn]

Patil, P., Reddy, T., Salgueiro, G., and M. Petit-Huguenin, "Application Layer Protocol Negotiation (ALPN) for Session Traversal Utilities for NAT (STUN)", draft-patil-tram-alpn-00 (work in progress), April 2014.

[I-D.thomson-rtcweb-alpn]

Thomson, M., "Application Layer Protocol Negotiation for Web Real-Time Communications (WebRTC)", draft-thomson-rtcweb-alpn-00 (work in progress), April 2014.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC7230] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, June 2014.

[RFC7231] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, June 2014.

6.2. Informative References

- [I-D.ietf-rtcweb-transports]
Alvestrand, H., "Transports for RTCWEB", draft-ietf-rtcweb-transports-05 (work in progress), June 2014.
- [I-D.ietf-rtcweb-use-cases-and-requirements]
Holmberg, C., Hakansson, S., and G. Eriksson, "Web Real-Time Communication Use-cases and Requirements", draft-ietf-rtcweb-use-cases-and-requirements-14 (work in progress), February 2014.
- [I-D.ietf-tls-applayerprotoneg]
Friedl, S., Popov, A., Langley, A., and S. Emile, "Transport Layer Security (TLS) Application Layer Protocol Negotiation Extension", draft-ietf-tls-applayerprotoneg-05 (work in progress), March 2014.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC6544] Rosenberg, J., Keranen, A., Lowekamp, B., and A. Roach, "TCP Candidates with Interactive Connectivity Establishment (ICE)", RFC 6544, March 2012.

Authors' Addresses

Andrew Hutton
Unify
Technology Drive
Nottingham NG9 1LA
UK

Email: andrew.hutton@unify.com

Justin Uberti
Google
747 6th Ave S
Kirkland, WA 98033
US

Email: justin@uberti.name

Martin Thomson
Mozilla
331 E Evelyn Street
Mountain View, CA 94041
US

Email: martin.thomson@gmail.com

HTTP Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 9, 2016

M. Nottingham
Akamai
P. McManus
Mozilla
J. Reschke
greenbytes
March 8, 2016

HTTP Alternative Services
draft-ietf-httpbis-alt-svc-14

Abstract

This document specifies "Alternative Services" for HTTP, which allow an origin's resources to be authoritatively available at a separate network location, possibly accessed with a different protocol configuration.

Editorial Note (To be removed by RFC Editor)

Discussion of this draft takes place on the HTTPBIS 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>.

The changes in this draft are summarized in Appendix A.

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 September 9, 2016.

Copyright Notice

Copyright (c) 2016 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.

Table of Contents

1. Introduction	4
1.1. Notational Conventions	4
2. Alternative Services Concepts	5
2.1. Host Authentication	7
2.2. Alternative Service Caching	7
2.3. Requiring Server Name Indication	8
2.4. Using Alternative Services	8
3. The Alt-Svc HTTP Header Field	9
3.1. Caching Alt-Svc Header Field Values	11
4. The ALTSVC HTTP/2 Frame	12
5. The Alt-Used HTTP Header Field	14
6. The 421 Misdirected Request HTTP Status Code	14
7. IANA Considerations	15
7.1. Header Field Registrations	15
7.2. The ALTSVC HTTP/2 Frame Type	15
7.3. Alt-Svc Parameter Registry	15
7.3.1. Procedure	15
7.3.2. Registrations	16
8. Internationalization Considerations	16
9. Security Considerations	16
9.1. Changing Ports	16
9.2. Changing Hosts	17
9.3. Changing Protocols	17
9.4. Tracking Clients Using Alternative Services	18
9.5. Confusion Regarding Request Scheme	18
10. References	19
10.1. Normative References	19
10.2. Informative References	20
Appendix A. Change Log (to be removed by RFC Editor before publication)	20
A.1. Since draft-nottingham-httpbis-alt-svc-05	20
A.2. Since draft-ietf-httpbis-alt-svc-00	21
A.3. Since draft-ietf-httpbis-alt-svc-01	21
A.4. Since draft-ietf-httpbis-alt-svc-02	21
A.5. Since draft-ietf-httpbis-alt-svc-03	21
A.6. Since draft-ietf-httpbis-alt-svc-04	21
A.7. Since draft-ietf-httpbis-alt-svc-05	22
A.8. Since draft-ietf-httpbis-alt-svc-06	22
A.9. Since draft-ietf-httpbis-alt-svc-07	22
A.10. Since draft-ietf-httpbis-alt-svc-08	23
A.11. Since draft-ietf-httpbis-alt-svc-09	24
A.12. Since draft-ietf-httpbis-alt-svc-10	24
A.13. Since draft-ietf-httpbis-alt-svc-11	24
A.14. Since draft-ietf-httpbis-alt-svc-12	24
Appendix B. Acknowledgements	24

1. Introduction

HTTP [RFC7230] conflates the identification of resources with their location. In other words, "http://" and "https://" URIs are used to both name and find things to interact with.

In some cases, it is desirable to separate identification and location in HTTP; keeping the same identifier for a resource, but interacting with it at a different location on the network.

For example:

- o An origin server might wish to redirect a client to a different server when it is under load, or it has found a server in a location that is more local to the client.
- o An origin server might wish to offer access to its resources using a new protocol, such as HTTP/2 [RFC7540], or one using improved security, such as Transport Layer Security (TLS) [RFC5246].
- o An origin server might wish to segment its clients into groups of capabilities, such as those supporting Server Name Indication (SNI) (Section 3 of [RFC6066]), for operational purposes.

This specification defines a new concept in HTTP, "Alternative Services", that allows an origin server to nominate additional means of interacting with it on the network. It defines a general framework for this in Section 2, along with specific mechanisms for advertising their existence using HTTP header fields (Section 3) or HTTP/2 frames (Section 4), plus a way to indicate that an alternative service was used (Section 5).

It also endorses the status code 421 (Misdirected Request) (Section 6) that origin servers or their nominated alternatives can use to indicate that they are not authoritative for a given origin, in cases where the wrong location is used.

1.1. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

This document uses the Augmented BNF defined in [RFC5234] and updated by [RFC7405] along with the "#rule" extension defined in Section 7 of [RFC7230]. The rules below are defined in [RFC5234], [RFC7230], and [RFC7234]:

OWS = <OWS, see [RFC7230], Section 3.2.3>
delta-seconds = <delta-seconds; see [RFC7234], Section 1.2.1>
port = <port, see [RFC7230], Section 2.7>
quoted-string = <quoted-string, see [RFC7230], Section 3.2.6>
token = <token, see [RFC7230], Section 3.2.6>
uri-host = <uri-host, see [RFC7230], Section 2.7>

2. Alternative Services Concepts

This specification defines a new concept in HTTP, the "Alternative Service". When an origin [RFC6454] has resources that are accessible through a different protocol / host / port combination, it is said to have an alternative service available.

An alternative service can be used to interact with the resources on an origin server at a separate location on the network, possibly using a different protocol configuration. Alternative services are considered authoritative for an origin's resources, in the sense of [RFC7230], Section 9.1.

For example, an origin:

```
("http", "www.example.com", "80")
```

might declare that its resources are also accessible at the alternative service:

```
("h2", "new.example.com", "81")
```

By their nature, alternative services are explicitly at the granularity of an origin; they cannot be selectively applied to resources within an origin.

Alternative services do not replace or change the origin for any given resource; in general, they are not visible to the software "above" the access mechanism. The alternative service is essentially alternative routing information that can also be used to reach the origin in the same way that DNS CNAME or SRV records define routing information at the name resolution level. Each origin maps to a set of these routes -- the default route is derived from the origin itself and the other routes are introduced based on alternative-service information.

Furthermore, it is important to note that the first member of an alternative service tuple is different from the "scheme" component of an origin; it is more specific, identifying not only the major version of the protocol being used, but potentially communication options for that protocol.

This means that clients using an alternative service can change the host, port and protocol that they are using to fetch resources, but these changes MUST NOT be propagated to the application that is using HTTP; from that standpoint, the URI being accessed and all information derived from it (scheme, host, port) are the same as before.

Importantly, this includes its security context; in particular, when TLS [RFC5246] is used to authenticate, the alternative service will need to present a certificate for the origin's host name, not that of the alternative. Likewise, the Host header field ([RFC7230], Section 5.4) is still derived from the origin, not the alternative service (just as it would if a CNAME were being used).

The changes MAY, however, be made visible in debugging tools, consoles, etc.

Formally, an alternative service is identified by the combination of:

- o An Application Layer Protocol Negotiation (ALPN) protocol name, as per [RFC7301]
- o A host, as per [RFC3986], Section 3.2.2
- o A port, as per [RFC3986], Section 3.2.3

The ALPN protocol name is used to identify the application protocol or suite of protocols used by the alternative service. Note that for the purpose of this specification, an ALPN protocol name implicitly includes TLS in the suite of protocols it identifies, unless specified otherwise in its definition. In particular, the ALPN name "http/1.1", registered by Section 6 of [RFC7301], identifies HTTP/1.1 over TLS.

Additionally, each alternative service MUST have:

- o A freshness lifetime, expressed in seconds; see Section 2.2

There are many ways that a client could discover the alternative service(s) associated with an origin. This document describes two such mechanisms: the "Alt-Svc" HTTP header field (Section 3) and the "ALTSVC" HTTP/2 frame type (Section 4).

The remainder of this section describes requirements that are common to alternative services, regardless of how they are discovered.

2.1. Host Authentication

Clients **MUST** have reasonable assurances that the alternative service is under control of and valid for the whole origin. This mitigates the attack described in Section 9.2.

For the purposes of this document, "reasonable assurances" can be established through use of a TLS-based protocol with the certificate checks defined in [RFC2818]. Clients **MAY** impose additional criteria for establishing reasonable assurances.

For example, if the origin's host is "www.example.com" and an alternative is offered on "other.example.com" with the "h2" protocol, and the certificate offered is valid for "www.example.com", the client can use the alternative. However, if either is offered with the "h2c" protocol, the client cannot use it, because there is no mechanism (at the time of the publication of this specification) in that protocol to establish the relationship between the origin and the alternative.

2.2. Alternative Service Caching

Mechanisms for discovering alternative services also associate a freshness lifetime with them; for example, the Alt-Svc header field uses the "ma" parameter.

Clients can choose to use an alternative service instead of the origin at any time when it is considered fresh; see Section 2.4 for specific recommendations.

Clients with existing connections to an alternative service do not need to stop using it when its freshness lifetime ends; the caching mechanism is intended for limiting how long an alternative service can be used for establishing new connections, not limiting the use of existing ones.

Alternative services are fully authoritative for the origin in question, including the ability to clear or update cached alternative service entries, extend freshness lifetimes, and any other authority the origin server would have.

When alternative services are used to send a client to the most optimal server, a change in network configuration can result in cached values becoming suboptimal. Therefore, clients **SHOULD** remove from cache all alternative services that lack the "persist" flag with the value "1" when they detect such a change, when information about network state is available.

2.3. Requiring Server Name Indication

A client **MUST NOT** use a TLS-based alternative service unless the client supports TLS Server Name Indication (SNI). This supports the conservation of IP addresses on the alternative service host.

Note that the SNI information provided in TLS by the client will be that of the origin, not the alternative (as will the Host HTTP header field value).

2.4. Using Alternative Services

By their nature, alternative services are **OPTIONAL**: clients do not need to use them. However, it is advantageous for clients to behave in a predictable way when alternative services are used by servers, to aid purposes like load balancing.

Therefore, if a client supporting this specification becomes aware of an alternative service, the client **SHOULD** use that alternative service for all requests to the associated origin as soon as it is available, provided the alternative service information is fresh (Section 2.2) and the security properties of the alternative service protocol are desirable, as compared to the existing connection. A viable alternative service is then treated in every way as the origin; this includes the ability to advertise alternative services.

If a client becomes aware of multiple alternative services, it chooses the most suitable according to its own criteria, keeping security properties in mind. For example, an origin might advertise multiple alternative services to notify clients of support for multiple versions of HTTP.

A client configured to use a proxy for a given request **SHOULD NOT** directly connect to an alternative service for this request, but instead route it through that proxy.

When a client uses an alternative service for a request, it can indicate this to the server using the Alt-Used header field (Section 5).

The client does not need to block requests on any existing connection; it can be used until the alternative connection is established. However, if the security properties of the existing connection are weak (for example, cleartext HTTP/1.1) then it might make sense to block until the new connection is fully available in order to avoid information leakage.

Furthermore, if the connection to the alternative service fails or is

unresponsive, the client MAY fall back to using the origin or another alternative service. Note, however, that this could be the basis of a downgrade attack, thus losing any enhanced security properties of the alternative service. If the connection to the alternative service does not negotiate the expected protocol (for example, ALPN fails to negotiate h2, or an Upgrade request to h2c is not accepted), the connection to the alternative service MUST be considered to have failed.

3. The Alt-Svc HTTP Header Field

An HTTP(S) origin server can advertise the availability of alternative services to clients by adding an Alt-Svc header field to responses.

```
Alt-Svc      = clear / 1#alt-value
clear        = %s"clear"; "clear", case-sensitive
alt-value    = alternative *( OWS ";" OWS parameter )
alternative  = protocol-id "=" alt-authority
protocol-id  = token ; percent-encoded ALPN protocol name
alt-authority = quoted-string ; containing [ uri-host ] ":" port
parameter    = token "=" ( token / quoted-string )
```

The field value consists either of a list of values, each of which indicates one alternative service, or the keyword "clear".

A field value containing the special value "clear" indicates that the origin requests all alternatives for that origin to be invalidated (including those specified in the same response, in case of an invalid reply containing both "clear" and alternative services).

ALPN protocol names are octet sequences with no additional constraints on format. Octets not allowed in tokens ([RFC7230], Section 3.2.6) MUST be percent-encoded as per Section 2.1 of [RFC3986]. Consequently, the octet representing the percent character "%" (hex 25) MUST be percent-encoded as well.

In order to have precisely one way to represent any ALPN protocol name, the following additional constraints apply:

1. Octets in the ALPN protocol name MUST NOT be percent-encoded if they are valid token characters except "%", and
2. When using percent-encoding, uppercase hex digits MUST be used.

With these constraints, recipients can apply simple string comparison to match protocol identifiers.

The "alt-authority" component consists of an OPTIONAL uri-host ("host" in Section 3.2.2 of [RFC3986]), a colon (":"), and a port number.

For example:

```
Alt-Svc: h2=":8000"
```

This indicates the "h2" protocol ([RFC7540]) on the same host using the indicated port 8000.

An example involving a change of host:

```
Alt-Svc: h2="new.example.org:80"
```

This indicates the "h2" protocol on the host "new.example.org", running on port 80. Note that the "quoted-string" syntax needs to be used because ":" is not an allowed character in "token".

Examples for protocol name escaping:

ALPN protocol name	protocol-id	Note
h2	h2	No escaping needed
w=x:y#z	w%3Dx%3Ay#z	"=" and ":" escaped
x%y	x%25y	"%" needs escaping

Alt-Svc MAY occur in any HTTP response message, regardless of the status code. Note that recipients of Alt-Svc can ignore the header field (and are required to in some situations; see Sections 2.1 and 6).

The Alt-Svc field value can have multiple values:

```
Alt-Svc: h2="alt.example.com:8000", h2=":443"
```

When multiple values are present, the order of the values reflects the server's preference (with the first value being the most preferred alternative).

The value(s) advertised by Alt-Svc can be used by clients to open a new connection to an alternative service. Subsequent requests can start using this new connection immediately, or can continue using the existing connection while the new connection is created.

When using HTTP/2 ([RFC7540]), servers SHOULD instead send an ALTSVC frame (Section 4). A single ALTSVC frame can be sent for a connection; a new frame is not needed for every request. Note that, despite this recommendation, Alt-Svc header fields remain valid in responses delivered over HTTP/2.

Each "alt-value" is followed by an OPTIONAL semicolon-separated list of additional parameters, each such "parameter" comprising a name and a value.

This specification defines two parameters: "ma" and "persist", defined in Section 3.1. Unknown parameters MUST be ignored. That is, the values (alt-value) they appear in MUST be processed as if the unknown parameter was not present.

New parameters can be defined in extension specifications (see Section 7.3 for registration details).

Note that all field elements that allow "quoted-string" syntax MUST be processed as per Section 3.2.6 of [RFC7230].

3.1. Caching Alt-Svc Header Field Values

When an alternative service is advertised using Alt-Svc, it is considered fresh for 24 hours from generation of the message. This can be modified with the 'ma' (max-age) parameter.

Syntax:

ma = delta-seconds; see [RFC7234], Section 1.2.1

The delta-seconds value indicates the number of seconds since the response was generated the alternative service is considered fresh for.

Alt-Svc: h2=":443"; ma=3600

See Section 4.2.3 of [RFC7234] for details of determining response age.

For example, a response:

```
HTTP/1.1 200 OK
Content-Type: text/html
Cache-Control: max-age=600
Age: 30
Alt-Svc: h2=":8000"; ma=60
```

indicates that an alternative service is available and usable for the next 60 seconds. However, the response has already been cached for 30 seconds (as per the Age header field value), so therefore the alternative service is only fresh for the 30 seconds from when this response was received, minus estimated transit time.

Note that the freshness lifetime for HTTP caching (here, 600 seconds) does not affect caching of Alt-Svc values.

When an Alt-Svc response header field is received from an origin, its value invalidates and replaces all cached alternative services for that origin.

By default, cached alternative services will be cleared when the client detects a network change. Alternative services that are intended to be longer-lived (such as those that are not specific to the client access network) can carry the "persist" parameter with a value "1" as a hint that the service is potentially useful beyond a network configuration change.

Syntax:

```
persist = "1"
```

For example:

```
Alt-Svc: h2=":443"; ma=2592000; persist=1
```

This specification only defines a single value for "persist". Clients MUST ignore "persist" parameters with values other than "1".

See Section 2.2 for general requirements on caching alternative services.

4. The ALTSVC HTTP/2 Frame

The ALTSVC HTTP/2 frame ([RFC7540], Section 4) advertises the availability of an alternative service to an HTTP/2 client.

The ALTSVC frame is a non-critical extension to HTTP/2. Endpoints

that do not support this frame will ignore it (as per the extensibility rules defined in Section 4.1 of [RFC7540]).

An ALTSVC frame from a server to a client on a stream other than stream 0 indicates that the conveyed alternative service is associated with the origin of that stream.

An ALTSVC frame from a server to a client on stream 0 indicates that the conveyed alternative service is associated with the origin contained in the Origin field of the frame. An association with an origin that the client does not consider authoritative for the current connection **MUST** be ignored.

The ALTSVC frame type is 0xa (decimal 10).

Origin-Len (16)	Origin? (*)	...
Alt-Svc-Field-Value (*)		

ALTSVC Frame Payload

The ALTSVC frame contains the following fields:

Origin-Len: An unsigned, 16-bit integer indicating the length, in octets, of the Origin field.

Origin: An OPTIONAL sequence of characters containing the ASCII serialization of an origin ([RFC6454], Section 6.2) that the alternative service is applicable to.

Alt-Svc-Field-Value: A sequence of octets (length determined by subtracting the length of all preceding fields from the frame length) containing a value identical to the Alt-Svc field value defined in Section 3 (ABNF production "Alt-Svc").

The ALTSVC frame does not define any flags.

The ALTSVC frame is intended for receipt by clients. A device acting as a server **MUST** ignore it.

An ALTSVC frame on stream 0 with empty (length 0) "Origin" information is invalid and **MUST** be ignored. An ALTSVC frame on a stream other than stream 0 containing non-empty "Origin" information is invalid and **MUST** be ignored.

The ALTSVC frame is processed hop-by-hop. An intermediary **MUST NOT**

forward ALTSVC frames, though it can use the information contained in ALTSVC frames in forming new ALTSVC frames to send to its own clients.

Receiving an ALTSVC frame is semantically equivalent to receiving an Alt-Svc header field. As a result, the ALTSVC frame causes alternative services for the corresponding origin to be replaced. Note that it would be unwise to mix the use of Alt-Svc header fields with the use of ALTSVC frames, as the sequence of receipt might be hard to predict.

5. The Alt-Used HTTP Header Field

The Alt-Used header field is used in requests to indicate the identity of the alternative service in use, just as the Host header field (Section 5.4 of [RFC7230]) identifies the host and port of the origin.

Alt-Used = uri-host [":" port]

Alt-Used is intended to allow alternative services to detect loops, differentiate traffic for purposes of load balancing, and generally to ensure that it is possible to identify the intended destination of traffic, since introducing this information after a protocol is in use has proven to be problematic.

When using an alternative service, clients SHOULD include an Alt-Used header field in all requests.

For example:

```
GET /thing HTTP/1.1
Host: origin.example.com
Alt-Used: alternate.example.net
```

6. The 421 Misdirected Request HTTP Status Code

The 421 (Misdirected Request) status code is defined in Section 9.1.2 of [RFC7540] to indicate that the current server instance is not authoritative for the requested resource. This can be used to indicate that an alternative service is not authoritative; see Section 2).

Clients receiving 421 (Misdirected Request) from an alternative service MUST remove the corresponding entry from its alternative service cache (see Section 2.2) for that origin. Regardless of the idempotency of the request method, they MAY retry the request, either at another alternative server, or at the origin.

An Alt-Svc header field in a 421 (Misdirected Request) response MUST be ignored.

7. IANA Considerations

7.1. Header Field Registrations

HTTP header fields are registered within the "Message Headers" registry maintained at <https://www.iana.org/assignments/message-headers/>.

This document defines the following HTTP header fields, so their associated registry entries shall be added according to the permanent registrations below (see [BCP90]):

Header Field Name	Protocol	Status	Reference
Alt-Svc	http	standard	Section 3
Alt-Used	http	standard	Section 5

The change controller is: "IETF (iesg@ietf.org) - Internet Engineering Task Force".

7.2. The ALTSVC HTTP/2 Frame Type

This document registers the ALTSVC frame type in the HTTP/2 Frame Types registry ([RFC7540], Section 11.2).

Frame Type: ALTSVC

Code: 0xa

Specification: Section 4 of this document

7.3. Alt-Svc Parameter Registry

The HTTP Alt-Svc Parameter Registry defines the name space for parameters. It will be created and maintained at (the suggested URI) <http://www.iana.org/assignments/http-alt-svc-parameters>.

7.3.1. Procedure

A registration MUST include the following fields:

- o Parameter Name

- o Pointer to specification text

Values to be added to this name space require Expert Review (see [RFC5226], Section 4.1).

7.3.2. Registrations

The HTTP Alt-Svc Parameter Registry is to be populated with the registrations below:

Alt-Svc Parameter	Reference
ma	Section 3.1
persist	Section 3.1

8. Internationalization Considerations

An internationalized domain name that appears in either the header field (Section 3) or the HTTP/2 frame (Section 4) MUST be expressed using A-labels ([RFC5890], Section 2.3.2.1).

9. Security Considerations

9.1. Changing Ports

Using an alternative service implies accessing an origin's resources on an alternative port, at a minimum. An attacker that can inject alternative services and listen at the advertised port is therefore able to hijack an origin. On certain servers, it is normal for users to be able to control some personal pages available on a shared port, and also to accept to requests on less-privileged ports.

For example, an attacker that can add HTTP response header fields to some pages can redirect traffic for an entire origin to a different port on the same host using the Alt-Svc header field; if that port is under the attacker's control, they can thus masquerade as the HTTP server.

This risk is mitigated by the requirements in Section 2.1.

On servers, this risk can also be reduced by restricting the ability to advertise alternative services, and restricting who can open a port for listening on that host.

9.2. Changing Hosts

When the host is changed due to the use of an alternative service, it presents an opportunity for attackers to hijack communication to an origin.

For example, if an attacker can convince a user agent to send all traffic for "innocent.example.org" to "evil.example.com" by successfully associating it as an alternative service, they can masquerade as that origin. This can be done locally (see mitigations in Section 9.1) or remotely (e.g., by an intermediary as a man-in-the-middle attack).

This is the reason for the requirement in Section 2.1 that clients have reasonable assurances that the alternative service is under control of and valid for the whole origin; for example, presenting a certificate for the origin proves that the alternative service is authorized to serve traffic for the origin.

Note that this assurance is only as strong as the method used to authenticate the alternative service. In particular, when TLS authentication is used to do so, there are well-known exploits to make an attacker's certificate appear as legitimate.

Alternative services could be used to persist such an attack. For example, an intermediary could man-in-the-middle TLS-protected communication to a target, and then direct all traffic to an alternative service with a large freshness lifetime, so that the user agent still directs traffic to the attacker even when not using the intermediary.

Implementations MUST perform any certificate-pinning validation (such as [RFC7469]) on alternative services just as they would on direct connections to the origin. Implementations might also choose to add other requirements around which certificates are acceptable for alternative services.

9.3. Changing Protocols

When the ALPN protocol is changed due to the use of an alternative service, the security properties of the new connection to the origin can be different from that of the "normal" connection to the origin, because the protocol identifier itself implies this.

For example, if an "https://" URI has a protocol advertised that does not use some form of end-to-end encryption (most likely, TLS), it violates the expectations for security that the URI scheme implies. Therefore, clients cannot blindly use alternative services, but

instead evaluate the option(s) presented to assure that security requirements and expectations of specifications, implementations and end users are met.

9.4. Tracking Clients Using Alternative Services

Choosing an alternative service implies connecting to a new, server-supplied host name. By using unique names, servers could conceivably track client requests. Such tracking could follow users across multiple networks, when the "persist" flag is used.

Clients that wish to prevent requests from being correlated can decide not to use alternative services for multiple requests that would not otherwise be allowed to be correlated.

In a user agent, any alternative service information **MUST** be removed when origin-specific data is cleared (typically, when cookies [RFC6265] are cleared).

9.5. Confusion Regarding Request Scheme

Some server-side HTTP applications make assumptions about security based upon connection context; for example, equating being served upon port 443 with the use of an "https://" URI and the various security properties that implies.

This affects not only the security properties of the connection itself, but also the state of the client at the other end of it; for example, a Web browser treats "https://" URIs differently than "http://" URIs in many ways, not just for purposes of protocol handling.

Since one of the uses of Alternative Services is to allow a connection to be migrated to a different protocol and port, these applications can become confused about the security properties of a given connection, sending information (for example, cookies and content) that is intended for a secure context (such as an "https://" URI) to a client that is not treating it as one.

This risk can be mitigated in servers by using the URI scheme explicitly carried by the protocol (such as ":scheme" in HTTP/2 or the "absolute form" of the request target in HTTP/1.1) as an indication of security context, instead of other connection properties ([RFC7540], Section 8.1.2.3 and [RFC7230], Section 5.3.2).

When the protocol does not explicitly carry the scheme (as is usually the case for HTTP/1.1 over TLS), servers can mitigate this risk by either assuming that all requests have an insecure context, or by

refraining from advertising alternative services for insecure schemes (for example, HTTP).

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>>.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, DOI 10.17487/RFC2818, May 2000, <<http://www.rfc-editor.org/info/rfc2818>>.
- [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, <<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [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>>.
- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, DOI 10.17487/RFC5890, August 2010, <<http://www.rfc-editor.org/info/rfc5890>>.
- [RFC6066] Eastlake, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", RFC 6066, DOI 10.17487/RFC6066, January 2011, <<http://www.rfc-editor.org/info/rfc6066>>.
- [RFC6454] Barth, A., "The Web Origin Concept", RFC 6454, DOI 10.17487/RFC6454, December 2011, <<http://www.rfc-editor.org/info/rfc6454>>.
- [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>>.

- [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>>.
- [RFC7301] Friedl, S., Popov, A., Langley, A., and S. Emile, "Transport Layer Security (TLS) Application-Layer Protocol Negotiation Extension", RFC 7301, DOI 10.17487/RFC7301, July 2014, <<http://www.rfc-editor.org/info/rfc7301>>.
- [RFC7405] Kyzivat, P., "Case-Sensitive String Support in ABNF", RFC 7405, DOI 10.17487/RFC7405, December 2014, <<http://www.rfc-editor.org/info/rfc7405>>.
- [RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol version 2", RFC 7540, DOI 10.17487/RFC7540, May 2015, <<http://www.rfc-editor.org/info/rfc7540>>.

10.2. Informative References

- [BCP90] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", BCP 90, RFC 3864, September 2004, <<http://www.rfc-editor.org/info/bcp90>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.
- [RFC6265] Barth, A., "HTTP State Management Mechanism", RFC 6265, DOI 10.17487/RFC6265, April 2011, <<http://www.rfc-editor.org/info/rfc6265>>.
- [RFC7469] Evans, C., Palmer, C., and R. Sleevi, "Public Key Pinning Extension for HTTP", RFC 7469, DOI 10.17487/RFC7469, April 2015, <<http://www.rfc-editor.org/info/rfc7469>>.

Appendix A. Change Log (to be removed by RFC Editor before publication)

A.1. Since draft-nottingham-httpbis-alt-svc-05

This is the first version after adoption of draft-nottingham-httpbis-alt-svc-05 as Working Group work item. It only contains editorial changes.

A.2. Since draft-ietf-httpbis-alt-svc-00

Selected 421 as proposed status code for "Not Authoritative".

Changed header field syntax to use percent-encoding of ALPN protocol names (<<https://github.com/http2/http2-spec/issues/446>>).

A.3. Since draft-ietf-httpbis-alt-svc-01

Updated HTTP/1.1 references.

Renamed "Service" to "Alt-Svc-Used" and reduced information to a flag to address fingerprinting concerns (<<https://github.com/http2/http2-spec/issues/502>>).

Note that ALTSVC frame is preferred to Alt-Svc header field (<<https://github.com/http2/http2-spec/pull/503>>).

Incorporate ALTSRV frame (<<https://github.com/http2/http2-spec/pull/507>>).

Moved definition of status code 421 to HTTP/2.

Partly resolved <<https://github.com/httpwg/http-extensions/issues/5>>.

A.4. Since draft-ietf-httpbis-alt-svc-02

Updated ALPN reference.

Resolved <<https://github.com/httpwg/http-extensions/issues/2>>.

A.5. Since draft-ietf-httpbis-alt-svc-03

Renamed "Alt-Svc-Used" to "Alt-Used" (<<https://github.com/httpwg/http-extensions/issues/17>>).

Clarify ALTSVC Origin information requirements (<<https://github.com/httpwg/http-extensions/issues/19>>).

Remove/tune language with respect to tracking risks (see <<https://github.com/httpwg/http-extensions/issues/34>>).

A.6. Since draft-ietf-httpbis-alt-svc-04

Mention tracking by alt-svc host name in Security Considerations (<<https://github.com/httpwg/http-extensions/issues/36>>).

"421 (Not Authoritative)" -> "421 (Misdirected Request)".

Allow the frame to carry multiple indicator and use the same payload formats for both
(<https://github.com/httpwg/http-extensions/issues/37>).

A.7. Since draft-ietf-httpbis-alt-svc-05

Go back to specifying the origin in Alt-Used, but make it a "SHOULD"
(<https://github.com/httpwg/http-extensions/issues/34>).

Restore Origin field in ALT-SVC frame
(<https://github.com/httpwg/http-extensions/issues/38>).

A.8. Since draft-ietf-httpbis-alt-svc-06

Disallow use of alternative services when the protocol might not carry the scheme
(<https://github.com/httpwg/http-extensions/issues/12>).

Align opp-sec and alt-svc
(<https://github.com/httpwg/http-extensions/issues/33>).

alt svc frame on pushed (even and non-0) frame
(<https://github.com/httpwg/http-extensions/issues/44>).

"browser" -> "user agent"
(<https://github.com/httpwg/http-extensions/pull/61>).

ABNF for "parameter"
(<https://github.com/httpwg/http-extensions/issues/65>).

Updated HTTP/2 reference.

A.9. Since draft-ietf-httpbis-alt-svc-07

Alt-Svc alternative cache invalidation
(<https://github.com/httpwg/http-extensions/issues/16>).

Unexpected Alt-Svc frames
(<https://github.com/httpwg/http-extensions/issues/18>).

Associating Alt-Svc header with an origin
(<https://github.com/httpwg/http-extensions/issues/21>).

ALPN identifiers in Alt-Svc
(<https://github.com/httpwg/http-extensions/issues/43>).

Number of alternate services used
(<https://github.com/httpwg/http-extensions/issues/58>).

Proxy and .pac interaction

(<https://github.com/httpwg/http-extensions/issues/62>).

Need to define extensibility for alt-svc parameters

(<https://github.com/httpwg/http-extensions/issues/69>).

Persistence of alternates across network changes

(<https://github.com/httpwg/http-extensions/issues/71>).

Alt-Svc header with 421 status

(<https://github.com/httpwg/http-extensions/issues/75>).

Incorporate several editorial improvements suggested by Mike Bishop

(<https://github.com/httpwg/http-extensions/pull/77>,

<https://github.com/httpwg/http-extensions/pull/78>).

Alt-Svc response header field in HTTP/2 frame

(<https://github.com/httpwg/http-extensions/issues/87>).

A.10. Since draft-ietf-httpbis-alt-svc-08

Remove left over text about ext-params, applying to an earlier version of Alt-Used (see

<https://github.com/httpwg/http-extensions/issues/34>).

Conflicts between Alt-Svc and ALPN

(<https://github.com/httpwg/http-extensions/issues/72>).

Elevation of privilege

(<https://github.com/httpwg/http-extensions/issues/73>).

Alternates of alternates

(<https://github.com/httpwg/http-extensions/issues/74>).

Alt-Svc and Cert Pinning

(<https://github.com/httpwg/http-extensions/issues/76>).

Using alt-svc on localhost (no change to spec, see

<https://github.com/httpwg/http-extensions/issues/89>).

IANA procedure for alt-svc parameters

(<https://github.com/httpwg/http-extensions/issues/96>).

Alt-svc from https (1.1) to https (1.1)

(<https://github.com/httpwg/http-extensions/issues/91>).

Alt-svc vs the ability to convey the scheme inside the protocol

(<https://github.com/httpwg/http-extensions/issues/92>).

Reconciling MAY/can vs. SHOULD
(<https://github.com/httpwg/http-extensions/issues/101>).

Typo in alt-svc caching example
(<https://github.com/httpwg/http-extensions/issues/117>).

A.11. Since draft-ietf-httpbis-alt-svc-09

Editorial improvements
(<https://github.com/httpwg/http-extensions/issues/118>,
<https://github.com/httpwg/http-extensions/issues/119>,
<https://github.com/httpwg/http-extensions/issues/120>,
<https://github.com/httpwg/http-extensions/issues/121>,
<https://github.com/httpwg/http-extensions/issues/122>,
<https://github.com/httpwg/http-extensions/issues/123>,
<https://github.com/httpwg/http-extensions/issues/125>,
<https://github.com/httpwg/http-extensions/issues/126>).

A.12. Since draft-ietf-httpbis-alt-svc-10

Editorial improvements
(<https://github.com/httpwg/http-extensions/issues/130>).

Use RFC 7405 ABNF extension
(<https://github.com/httpwg/http-extensions/issues/131>).

A.13. Since draft-ietf-httpbis-alt-svc-11

Security considerations wrt system ports
(<https://github.com/httpwg/http-extensions/issues/139>).

A.14. Since draft-ietf-httpbis-alt-svc-12

Editorial changes triggered by <https://lists.w3.org/Archives/Public/ietf-http-wg/2016JanMar/0243.html>.

Reasonable Assurances and H2C
(<https://github.com/httpwg/http-extensions/issues/148>).

Appendix B. Acknowledgements

Thanks to Adam Langley, Bence Beky, Chris Lonvick, Eliot Lear, Erik Nygren, Guy Podjarny, Herve Ruellan, Lucas Pardue, Martin Thomson, Matthew Kerwin, Mike Bishop, Paul Hoffman, Richard Barnes, Richard Bradbury, Stephen Farrell, Stephen Ludin, and Will Chan for their feedback and suggestions.

The Alt-Svc header field was influenced by the design of the

Alternate-Protocol header field in SPDY.

Authors' Addresses

Mark Nottingham
Akamai

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

Patrick McManus
Mozilla

EMail: mcmanus@ducksong.com
URI: <https://mozillians.org/u/pmcmanus/>

Julian F. Reschke
greenbytes GmbH

EMail: julian.reschke@greenbytes.de
URI: <https://greenbytes.de/tech/webdav/>

HTTPbis Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 21, 2015

R. Peon
Google, Inc
H. Ruellan
Canon CRF
February 17, 2015

HPACK - Header Compression for HTTP/2
draft-ietf-httpbis-header-compression-12

Abstract

This specification defines HPACK, a compression format for efficiently representing HTTP header fields, to be used in HTTP/2.

Editorial Note (To be removed by RFC Editor)

Discussion of this draft takes place on the HTTPBIS working group mailing list (ietf-http-wg@w3.org), which is archived at [1].

Working Group information can be found at [2]; that specific to HTTP/2 are at [3].

The changes in this draft are summarized in Appendix D.2.

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 August 21, 2015.

Copyright Notice

Copyright (c) 2015 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.

Table of Contents

1. Introduction	4
1.1. Overview	4
1.2. Conventions	5
1.3. Terminology	5
2. Compression Process Overview	6
2.1. Header List Ordering	6
2.2. Encoding and Decoding Contexts	6
2.3. Indexing Tables	6
2.3.1. Static Table	6
2.3.2. Dynamic Table	6
2.3.3. Index Address Space	7
2.4. Header Field Representation	8
3. Header Block Decoding	8
3.1. Header Block Processing	8
3.2. Header Field Representation Processing	9
4. Dynamic Table Management	9
4.1. Calculating Table Size	10
4.2. Maximum Table Size	10
4.3. Entry Eviction when Dynamic Table Size Changes	11
4.4. Entry Eviction when Adding New Entries	11
5. Primitive Type Representations	11
5.1. Integer Representation	11
5.2. String Literal Representation	13
6. Binary Format	14
6.1. Indexed Header Field Representation	14
6.2. Literal Header Field Representation	15
6.2.1. Literal Header Field with Incremental Indexing	15
6.2.2. Literal Header Field without Indexing	16
6.2.3. Literal Header Field never Indexed	17
6.3. Dynamic Table Size Update	18
7. Security Considerations	19
7.1. Probing Dynamic Table State	19
7.1.1. Applicability to HPACK and HTTP	20
7.1.2. Mitigation	20
7.1.3. Never Indexed Literals	21
7.2. Static Huffman Encoding	22

7.3. Memory Consumption	22
7.4. Implementation Limits	23
8. IANA Considerations	23
9. Acknowledgments	23
10. References	23
10.1. Normative References	23
10.2. Informative References	24
Appendix A. Static Table Definition	25
Appendix B. Huffman Code	26
Appendix C. Examples	32
C.1. Integer Representation Examples	33
C.1.1. Example 1: Encoding 10 Using a 5-bit Prefix	33
C.1.2. Example 2: Encoding 1337 Using a 5-bit Prefix	33
C.1.3. Example 3: Encoding 42 Starting at an Octet Boundary	34
C.2. Header Field Representation Examples	34
C.2.1. Literal Header Field with Indexing	34
C.2.2. Literal Header Field without Indexing	35
C.2.3. Literal Header Field never Indexed	36
C.2.4. Indexed Header Field	36
C.3. Request Examples without Huffman Coding	37
C.3.1. First Request	37
C.3.2. Second Request	38
C.3.3. Third Request	39
C.4. Request Examples with Huffman Coding	40
C.4.1. First Request	40
C.4.2. Second Request	41
C.4.3. Third Request	42
C.5. Response Examples without Huffman Coding	44
C.5.1. First Response	44
C.5.2. Second Response	46
C.5.3. Third Response	47
C.6. Response Examples with Huffman Coding	49
C.6.1. First Response	49
C.6.2. Second Response	51
C.6.3. Third Response	52
Appendix D. Change Log (to be removed by RFC Editor before publication)	54
D.1. Since draft-ietf-httpbis-header-compression-10	55
D.2. Since draft-ietf-httpbis-header-compression-09	55
D.3. Since draft-ietf-httpbis-header-compression-08	55
D.4. Since draft-ietf-httpbis-header-compression-07	55
D.5. Since draft-ietf-httpbis-header-compression-06	56
D.6. Since draft-ietf-httpbis-header-compression-05	56
D.7. Since draft-ietf-httpbis-header-compression-04	56
D.8. Since draft-ietf-httpbis-header-compression-03	57
D.9. Since draft-ietf-httpbis-header-compression-02	57
D.10. Since draft-ietf-httpbis-header-compression-01	57
D.11. Since draft-ietf-httpbis-header-compression-00	57

1. Introduction

In HTTP/1.1 (see [RFC7230]), header fields are not compressed. As Web pages have grown to require dozens to hundreds of requests, the redundant header fields in these requests unnecessarily consume bandwidth, measurably increasing latency.

SPDY [SPDY] initially addressed this redundancy by compressing header fields using the DEFLATE [DEFLATE] format, which proved very effective at efficiently representing the redundant header fields. However, that approach exposed a security risk as demonstrated by the CRIME attack (see [CRIME]).

This specification defines HPACK, a new compressor for header fields which eliminates redundant header fields, limits vulnerability to known security attacks, and which has a bounded memory requirement for use in constrained environments. Potential security concerns for HPACK are described in Section 7.

The HPACK format is intentionally simple and inflexible. Both characteristics reduce the risk of interoperability or security issues due to implementation error. No extensibility mechanisms are defined; changes to the format are only possible by defining a complete replacement.

1.1. Overview

The format defined in this specification treats a list of header fields as an ordered collection of name-value pairs that can include duplicate pairs. Names and values are considered to be opaque sequences of octets, and the order of header fields is preserved after being compressed and decompressed.

Encoding is informed by header field tables that map header fields to indexed values. These header field tables can be incrementally updated as new header fields are encoded or decoded.

In the encoded form, a header field is represented either literally or as a reference to a header field in one of the header field tables. Therefore, a list of header fields can be encoded using a mixture of references and literal values.

Literal values are either encoded directly or using a static Huffman code.

The encoder is responsible for deciding which header fields to insert as new entries in the header field tables. The decoder executes the modifications to the header field tables prescribed by the encoder,

reconstructing the list of header fields in the process. This enables decoders to remain simple and interoperate with a wide variety of encoders.

Examples illustrating the use of these different mechanisms to represent header fields are available in Appendix C.

1.2. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

All numeric values are in network byte order. Values are unsigned unless otherwise indicated. Literal values are provided in decimal or hexadecimal as appropriate.

1.3. Terminology

This specification uses the following terms:

Header Field: A name-value pair. Both the name and value are treated as opaque sequences of octets.

Dynamic Table: The dynamic table (see Section 2.3.2) is a table that associates stored header fields with index values. This table is dynamic and specific to an encoding or decoding context.

Static Table: The static table (see Section 2.3.1) is a table that statically associates header fields that occur frequently with index values. This table is ordered, read-only, always accessible, and may be shared amongst all encoding or decoding contexts.

Header List: A header list is an ordered collection of header fields that are encoded jointly, and can contain duplicate header fields. A complete list of header fields contained in an HTTP/2 header block is a header list.

Header Field Representation: A header field can be represented in encoded form either as a literal or as an index (see Section 2.4).

Header Block: An ordered list of header field representations which, when decoded, yields a complete header list.

2. Compression Process Overview

This specification does not describe a specific algorithm for an encoder. Instead, it defines precisely how a decoder is expected to operate, allowing encoders to produce any encoding that this definition permits.

2.1. Header List Ordering

HPACK preserves the ordering of header fields inside the header list. An encoder **MUST** order header field representations in the header block according to their ordering in the original header list. A decoder **MUST** order header fields in the decoded header list according to their ordering in the header block.

2.2. Encoding and Decoding Contexts

To decompress header blocks, a decoder only needs to maintain a dynamic table (see Section 2.3.2) as a decoding context. No other dynamic state is needed.

When used for bidirectional communication, such as in HTTP, the encoding and decoding dynamic tables maintained by an endpoint are completely independent. I.e., the request and response dynamic tables are separate.

2.3. Indexing Tables

HPACK uses two tables for associating header fields to indexes. The static table (see Section 2.3.1) is predefined and contains common header fields (most of them with an empty value). The dynamic table (see Section 2.3.2) is dynamic and can be used by the encoder to index header fields repeated in the encoded header lists.

These two tables are combined into a single address space for defining index values (see Section 2.3.3).

2.3.1. Static Table

The static table consists of a predefined static list of header fields. Its entries are defined in Appendix A.

2.3.2. Dynamic Table

The dynamic table consists of a list of header fields maintained in first-in, first-out order. The first and newest entry in a dynamic table is at the lowest index, and the oldest entry of a dynamic table is at the highest index.

The dynamic table is initially empty. Entries are added as each header block is decompressed.

The dynamic table can contain duplicate entries (i.e., entries with the same name and same value). Therefore, duplicate entries **MUST NOT** be treated as an error by a decoder.

The encoder decides how to update the dynamic table and as such can control how much memory is used by the dynamic table. To limit the memory requirements of the decoder, the dynamic table size is strictly bounded (see Section 4.2).

The decoder updates the dynamic table during the processing of a list of header field representations (see Section 3.2).

2.3.3. Index Address Space

The static table and the dynamic table are combined into a single index address space.

Indices between 1 and the length of the static table (inclusive) refer to elements in the static table (see Section 2.3.1).

Indices strictly greater than the length of the static table refer to elements in the dynamic table (see Section 2.3.2). The length of the static table is subtracted to find the index into the dynamic table.

Indices strictly greater than the sum of the lengths of both tables **MUST** be treated as a decoding error.

For a static table size of s and a dynamic table size of k , the following diagram shows the entire valid index address space.




Figure 1: Index Address Space

2.4. Header Field Representation

An encoded header field can be represented either as an index or as a literal.

An indexed representation defines a header field as a reference to an entry in either the static table or the dynamic table (see Section 6.1).

A literal representation defines a header field by specifying its name and value. The header field name can be represented literally or as a reference to an entry in either the static table or the dynamic table. The header field value is represented literally.

Three different literal representations are defined:

- o A literal representation that adds the header field as a new entry at the beginning of the dynamic table (see Section 6.2.1).
- o A literal representation that does not add the header field to the dynamic table (see Section 6.2.2).
- o A literal representation that does not add the header field to the dynamic table, with the additional stipulation that this header field always use a literal representation, in particular when re-encoded by an intermediary (see Section 6.2.3). This representation is intended for protecting header field values that are not to be put at risk by compressing them (see Section 7.1.3 for more details).

The selection of one of these literal representations can be guided by security considerations, in order to protect sensitive header field values (see Section 7.1).

The literal representation of a header field name or of a header field value can encode the sequence of octets either directly or using a static Huffman code (see Section 5.2).

3. Header Block Decoding

3.1. Header Block Processing

A decoder processes a header block sequentially to reconstruct the original header list.

A header block is the concatenation of header field representations. The different possible header field representations are described in Section 6.

Once a header field is decoded and added to the reconstructed header list, the header field cannot be removed. A header field added to the header list can be safely passed to the application.

By passing the resulting header fields to the application, a decoder can be implemented with minimal transitory memory commitment in addition to the dynamic table.

3.2. Header Field Representation Processing

The processing of a header block to obtain a header list is defined in this section. To ensure that the decoding will successfully produce a header list, a decoder **MUST** obey the following rules.

All the header field representations contained in a header block are processed in the order in which they appear, as specified below. Details on the formatting of the various header field representations, and some additional processing instructions are found in Section 6.

An `_indexed representation_` entails the following actions:

- o The header field corresponding to the referenced entry in either the static table or dynamic table is appended to the decoded header list.

A `_literal representation_` that is `_not added_` to the dynamic table entails the following action:

- o The header field is appended to the decoded header list.

A `_literal representation_` that is `_added_` to the dynamic table entails the following actions:

- o The header field is appended to the decoded header list.
- o The header field is inserted at the beginning of the dynamic table. This insertion could result in the eviction of previous entries in the dynamic table (see Section 4.4).

4. Dynamic Table Management

To limit the memory requirements on the decoder side, the dynamic table is constrained in size.

4.1. Calculating Table Size

The size of the dynamic table is the sum of the size of its entries.

The size of an entry is the sum of its name's length in octets (as defined in Section 5.2), its value's length in octets, plus 32.

The size of an entry is calculated using the length of its name and value without any Huffman encoding applied.

Note: The additional 32 octets account for an estimated overhead associated with an entry. For example, an entry structure using two 64-bit pointers to reference the name and the value of the entry, and two 64-bit integers for counting the number of references to the name and value would have 32 octets of overhead.

4.2. Maximum Table Size

Protocols that use HPACK determine the maximum size that the encoder is permitted to use for the dynamic table. In HTTP/2, this value is determined by the `SETTINGS_HEADER_TABLE_SIZE` setting (see Section 6.5.2 of [HTTP2]).

An encoder can choose to use less capacity than this maximum size (see Section 6.3), but the chosen size **MUST** stay lower than or equal to the maximum set by the protocol.

A change in the maximum size of the dynamic table is signaled via an encoding context update (see Section 6.3). This encoding context update **MUST** occur at the beginning of the first header block following the change to the dynamic table size. In HTTP/2, this follows a settings acknowledgment (see Section 6.5.3 of [HTTP2]).

Multiple updates to the maximum table size can occur between the transmission of two header blocks. In the case that this size is changed more than once in this interval, the smallest maximum table size that occurs in that interval **MUST** be signaled in an encoding context update. The final maximum size is always signaled, resulting in at most two encoding context updates. This ensures that the decoder is able to perform eviction based on reductions in dynamic table size (see Section 4.3).

This mechanism can be used to completely clear entries from the dynamic table by setting a maximum size of 0, which can subsequently be restored.

4.3. Entry Eviction when Dynamic Table Size Changes

Whenever the maximum size for the dynamic table is reduced, entries are evicted from the end of the dynamic table until the size of the dynamic table is less than or equal to the maximum size.

4.4. Entry Eviction when Adding New Entries

Before a new entry is added to the dynamic table, entries are evicted from the end of the dynamic table until the size of the dynamic table is less than or equal to (maximum size - new entry size), or until the table is empty.

If the size of the new entry is less than or equal to the maximum size, that entry is added to the table. It is not an error to attempt to add an entry that is larger than the maximum size; an attempt to add an entry larger than the maximum size causes the table to be emptied of all existing entries, and results in an empty table.

A new entry can reference the name of an entry in the dynamic table that will be evicted when adding this new entry into the dynamic table. Implementations are cautioned to avoid deleting the referenced name if the referenced entry is evicted from the dynamic table prior to inserting the new entry.

5. Primitive Type Representations

HPACK encoding uses two primitive types: unsigned variable length integers, and strings of octets.

5.1. Integer Representation

Integers are used to represent name indexes, header field indexes or string lengths. An integer representation can start anywhere within an octet. To allow for optimized processing, an integer representation always finishes at the end of an octet.

An integer is represented in two parts: a prefix that fills the current octet and an optional list of octets that are used if the integer value does not fit within the prefix. The number of bits of the prefix (called N) is a parameter of the integer representation.

If the integer value is small enough, i.e., strictly less than $2^N - 1$, it is encoded within the N -bit prefix.




Figure 2: Integer Value Encoded within the Prefix (shown for N = 5)

Otherwise, all the bits of the prefix are set to 1 and the value, decreased by 2^N-1 , is encoded using a list of one or more octets. The most significant bit of each octet is used as a continuation flag: its value is set to 1 except for the last octet in the list. The remaining bits of the octets are used to encode the decreased value.




Figure 3: Integer Value Encoded after the Prefix (shown for N = 5)

Decoding the integer value from the list of octets starts by reversing the order of the octets in the list. Then, for each octet, its most significant bit is removed. The remaining bits of the octets are concatenated and the resulting value is increased by 2^N-1 to obtain the integer value.

The prefix size, N, is always between 1 and 8 bits. An integer starting at an octet-boundary will have an 8-bit prefix.

Pseudo-code to represent an integer I is as follows:

```

if I <  $2^N - 1$ , encode I on N bits
else
    encode ( $2^N - 1$ ) on N bits
    I = I - ( $2^N - 1$ )
    while I >= 128
        encode (I % 128 + 128) on 8 bits
        I = I / 128
    encode I on 8 bits

```

Pseudo-code to decode an integer I is as follows:

```

decode I from the next N bits
if I < 2^N - 1, return I
else
    M = 0
    repeat
        B = next octet
        I = I + (B & 127) * 2^M
        M = M + 7
    while B & 128 == 128
    return I

```

Examples illustrating the encoding of integers are available in Appendix C.1.

This integer representation allows for values of indefinite size. It is also possible for an encoder to send a large number of zero values, which can waste octets and could be used to overflow integer values. Integer encodings that exceed an implementation limits - in value or octet length - MUST be treated as a decoding error. Different limits can be set for each of the different uses of integers, based on implementation constraints.

5.2. String Literal Representation

Header field names and header field values can be represented as literal strings. A literal string is encoded as a sequence of octets, either by directly encoding the literal string's octets, or by using a Huffman code (see [HUFFMAN]).




Figure 4: String Literal Representation

A literal string representation contains the following fields:

H: A one bit flag, H, indicating whether or not the octets of the string are Huffman encoded.

String Length: The number of octets used to encode the string literal, encoded as an integer with 7-bit prefix (see Section 5.1).

String Data: The encoded data of the string literal. If H is '0', then the encoded data is the raw octets of the string literal. If H is '1', then the encoded data is the Huffman encoding of the string literal.

String literals which use Huffman encoding are encoded with the Huffman code defined in Appendix B (see examples for requests in Appendix C.4 and for responses in Appendix C.6). The encoded data is the bitwise concatenation of the codes corresponding to each octet of the string literal.

As the Huffman encoded data doesn't always end at an octet boundary, some padding is inserted after it, up to the next octet boundary. To prevent this padding to be misinterpreted as part of the string literal, the most significant bits of the code corresponding to the EOS (end-of-string) symbol are used.

Upon decoding, an incomplete code at the end of the encoded data is to be considered as padding and discarded. A padding strictly longer than 7 bits MUST be treated as a decoding error. A padding not corresponding to the most significant bits of the code for the EOS symbol MUST be treated as a decoding error. A Huffman encoded string literal containing the EOS symbol MUST be treated as a decoding error.

6. Binary Format

This section describes the detailed format of each of the different header field representations, plus the encoding context update instruction.

6.1. Indexed Header Field Representation

An indexed header field representation identifies an entry in either the static table or the dynamic table (see Section 2.3).

An indexed header field representation causes a header field to be added to the decoded header list, as described in Section 3.2.




Figure 5: Indexed Header Field

An indexed header field starts with the '1' 1-bit pattern, followed by the index of the matching header field, represented as an integer with a 7-bit prefix (see Section 5.1).

The index value of 0 is not used. It MUST be treated as a decoding error if found in an indexed header field representation.

6.2. Literal Header Field Representation

A literal header field representation contains a literal header field value. Header field names are either provided as a literal or by reference to an existing table entry, either from the static table or the dynamic table (see Section 2.3).

This specification defines three forms of literal header field representations; with indexing, without indexing, and never indexed.

6.2.1. Literal Header Field with Incremental Indexing

A literal header field with incremental indexing representation results in appending a header field to the decoded header list and inserting it as a new entry into the dynamic table.




Figure 6: Literal Header Field with Incremental Indexing - Indexed Name




Figure 7: Literal Header Field with Incremental Indexing - New Name

A literal header field with incremental indexing representation starts with the '01' 2-bit pattern.

If the header field name matches the header field name of an entry stored in the static table or the dynamic table, the header field name can be represented using the index of that entry. In this case, the index of the entry is represented as an integer with a 6-bit prefix (see Section 5.1). This value is always non-zero.

Otherwise, the header field name is represented as a literal string (see Section 5.2). A value 0 is used in place of the 6-bit index, followed by the header field name.

Either form of header field name representation is followed by the header field value represented as a literal string (see Section 5.2).

6.2.2. Literal Header Field without Indexing

A literal header field without indexing representation results in appending a header field to the decoded header list without altering the dynamic table.




Figure 8: Literal Header Field without Indexing - Indexed Name




Figure 9: Literal Header Field without Indexing - New Name

A literal header field without indexing representation starts with the '0000' 4-bit pattern.

If the header field name matches the header field name of an entry stored in the static table or the dynamic table, the header field name can be represented using the index of that entry. In this case, the index of the entry is represented as an integer with a 4-bit prefix (see Section 5.1). This value is always non-zero.

Otherwise, the header field name is represented as a literal string (see Section 5.2). A value 0 is used in place of the 4-bit index, followed by the header field name.

Either form of header field name representation is followed by the header field value represented as a literal string (see Section 5.2).

6.2.3. Literal Header Field never Indexed

A literal header field never indexed representation results in appending a header field to the decoded header list without altering the dynamic table. Intermediaries MUST use the same representation for encoding this header field.




Figure 10: Literal Header Field never Indexed - Indexed Name




Figure 11: Literal Header Field never Indexed - New Name

A literal header field never indexed representation starts with the '0001' 4-bit pattern.

When a header field is represented as a literal header field never indexed, it **MUST** always be encoded with this specific literal representation. In particular, when a peer sends a header field that it received represented as a literal header field never indexed, it **MUST** use the same representation to forward this header field.

This representation is intended for protecting header field values that are not to be put at risk by compressing them (see Section 7.1 for more details).

The encoding of the representation is identical to the literal header field without indexing (see Section 6.2.2).

6.3. Dynamic Table Size Update

A dynamic table size update signals a change to the size of the dynamic table.




Figure 12: Maximum Dynamic Table Size Change

A dynamic table size update starts with the '001' 3-bit pattern, followed by the new maximum size, represented as an integer with a 5-bit prefix (see Section 5.1).

The new maximum size MUST be lower than or equal to the last value of the maximum size of the dynamic table. A value that exceeds this limit MUST be treated as a decoding error. In HTTP/2, this limit is the last value of the `SETTINGS_HEADER_TABLE_SIZE` parameter (see Section 6.5.2 of [HTTP2]) received from the decoder and acknowledged by the encoder (see Section 6.5.3 of [HTTP2]).

Reducing the maximum size of the dynamic table can cause entries to be evicted (see Section 4.3).

7. Security Considerations

This section describes potential areas of security concern with HPACK:

- o Use of compression as a length-based oracle for verifying guesses about secrets that are compressed into a shared compression context.
- o Denial of service resulting from exhausting processing or memory capacity at a decoder.

7.1. Probing Dynamic Table State

HPACK reduces the length of header field encodings by exploiting the redundancy inherent in protocols like HTTP. The ultimate goal of this is to reduce the amount of data that is required to send HTTP requests or responses.

The compression context used to encode header fields can be probed by an attacker who can both define header fields to be encoded and transmitted and observe the length of those fields once they are encoded. When an attacker can do both, they can adaptively modify requests in order to confirm guesses about the dynamic table state. If a guess is compressed into a shorter length, the attacker can observe the encoded length and infer that the guess was correct.

This is possible even over the Transport Layer Security Protocol (TLS, see [TLS12]), because while TLS provides confidentiality protection for content, it only provides a limited amount of protection for the length of that content.

Note: Padding schemes only provide limited protection against an attacker with these capabilities, potentially only forcing an increased number of guesses to learn the length associated with a given guess. Padding schemes also work directly against compression by increasing the number of bits that are transmitted.

Attacks like CRIME [CRIME] demonstrated the existence of these general attacker capabilities. The specific attack exploited the fact that DEFLATE [DEFLATE] removes redundancy based on prefix matching. This permitted the attacker to confirm guesses a character at a time, reducing an exponential-time attack into a linear-time attack.

7.1.1. Applicability to HPACK and HTTP

HPACK mitigates but does not completely prevent attacks modeled on CRIME [CRIME] by forcing a guess to match an entire header field value, rather than individual characters. An attacker can only learn whether a guess is correct or not, so is reduced to a brute force guess for the header field values.

The viability of recovering specific header field values therefore depends on the entropy of values. As a result, values with high entropy are unlikely to be recovered successfully. However, values with low entropy remain vulnerable.

Attacks of this nature are possible any time that two mutually distrustful entities control requests or responses that are placed onto a single HTTP/2 connection. If the shared HPACK compressor permits one entity to add entries to the dynamic table, and the other to access those entries, then the state of the table can be learned.

Having requests or responses from mutually distrustful entities occurs when an intermediary either:

- o sends requests from multiple clients on a single connection toward an origin server, or
- o takes responses from multiple origin servers and places them on a shared connection toward a client.

Web browsers also need to assume that requests made on the same connection by different web origins [ORIGIN] are made by mutually distrustful entities.

7.1.2. Mitigation

Users of HTTP that require confidentiality for header fields can use values with entropy sufficient to make guessing infeasible. However, this is impractical as a general solution because it forces all users of HTTP to take steps to mitigate attacks. It would impose new constraints on how HTTP is used.

Rather than impose constraints on users of HTTP, an implementation of HPACK can instead constrain how compression is applied in order to limit the potential for dynamic table probing.

An ideal solution segregates access to the dynamic table based on the entity that is constructing header fields. Header field values that are added to the table are attributed to an entity, and only the entity that created a particular value can extract that value.

To improve compression performance of this option, certain entries might be tagged as being public. For example, a web browser might make the values of the Accept-Encoding header field available in all requests.

An encoder without good knowledge of the provenance of header fields might instead introduce a penalty for a header field with many different values, such that a large number of attempts to guess a header field value results in the header field no more being compared to the dynamic table entries in future messages, effectively preventing further guesses.

Note: Simply removing entries corresponding to the header field from the dynamic table can be ineffectual if the attacker has a reliable way of causing values to be reinstalled. For example, a request to load an image in a web browser typically includes the Cookie header field (a potentially highly valued target for this sort of attack), and web sites can easily force an image to be loaded, thereby refreshing the entry in the dynamic table.

This response might be made inversely proportional to the length of the header field value. Marking a header field as not using the dynamic table any more might occur for shorter values more quickly or with higher probability than for longer values.

7.1.3. Never Indexed Literals

Implementations can also choose to protect sensitive header fields by not compressing them and instead encoding their value as literals.

Refusing to generate an indexed representation for a header field is only effective if compression is avoided on all hops. The never indexed literal (see Section 6.2.3) can be used to signal to intermediaries that a particular value was intentionally sent as a literal.

An intermediary **MUST NOT** re-encode a value that uses the never indexed literal representation with another representation that would

index it. If HPACK is used for re-encoding, the never indexed literal representation MUST be used.

The choice to use a never indexed literal representation for a header field depends on several factors. Since HPACK doesn't protect against guessing an entire header field value, short or low-entropy values are more readily recovered by an adversary. Therefore, an encoder might choose not to index values with low entropy.

An encoder might also choose not to index values for header fields that are considered to be highly valuable or sensitive to recovery, such as the Cookie or Authorization header fields.

On the contrary, an encoder might prefer indexing values for header fields that have little or no value if they were exposed. For instance, a User-Agent header field does not commonly vary between requests and is sent to any server. In that case, confirmation that a particular User-Agent value has been used provides little value.

Note that these criteria for deciding to use a never indexed literal representation will evolve over time as new attacks are discovered.

7.2. Static Huffman Encoding

There is no currently known attack against a static Huffman encoding. A study has shown that using a static Huffman encoding table created an information leakage, however this same study concluded that an attacker could not take advantage of this information leakage to recover any meaningful amount of information (see [PETAL]).

7.3. Memory Consumption

An attacker can try to cause an endpoint to exhaust its memory. HPACK is designed to limit both the peak and state amounts of memory allocated by an endpoint.

The amount of memory used by the compressor is limited by the protocol using HPACK through the definition of the maximum size of the dynamic table. In HTTP/2, this value is controlled by the decoder through the setting parameter SETTINGS_HEADER_TABLE_SIZE (see Section 6.5.2 of [HTTP2]). This limit takes into account both the size of the data stored in the dynamic table, plus a small allowance for overhead.

A decoder can limit the amount of state memory used by setting an appropriate value for the maximum size of the dynamic table. In HTTP/2, this is realized by setting an appropriate value for the SETTINGS_HEADER_TABLE_SIZE parameter. An encoder can limit the

amount of state memory it uses by signaling lower dynamic table size than the decoder allows (see Section 6.3).

The amount of temporary memory consumed by an encoder or decoder can be limited by processing header fields sequentially. An implementation does not need to retain a complete list of header fields. Note however that it might be necessary for an application to retain a complete header list for other reasons; even though HPACK does not force this to occur, application constraints might make this necessary.

7.4. Implementation Limits

An implementation of HPACK needs to ensure that large values for integers, long encoding for integers, or long string literals do not create security weaknesses.

An implementation has to set a limit for the values it accepts for integers, as well as for the encoded length (see Section 5.1). In the same way, it has to set a limit to the length it accepts for string literals (see Section 5.2).

8. IANA Considerations

This document has no IANA actions.

9. Acknowledgments

This specification includes substantial input from the following individuals:

- o Mike Bishop, Jeff Pinner, Julian Reschke, Martin Thomson (substantial editorial contributions).
- o Johnny Graettinger (Huffman code statistics).

10. References

10.1. Normative References

- [HTTP2] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol version 2", draft-ietf-httpbis-http2-17 (work in progress), February 2015.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, June 2014.

10.2. Informative References

- [CANONICAL] Schwartz, E. and B. Kallick, "Generating a canonical prefix encoding", Communications of the ACM Volume 7 Issue 3, pp. 166-169, March 1964, <<https://dl.acm.org/citation.cfm?id=363991>>.
- [CRIME] Rizzo, J. and T. Duong, "The CRIME Attack", September 2012, <https://docs.google.com/a/twist.com/presentation/d/11eBmGiHbYcHR9gL5nDyZChu_lCa2GizeuOfaLU2HOU>.
- [DEFLATE] Deutsch, P., "DEFLATE Compressed Data Format Specification version 1.3", RFC 1951, May 1996.
- [HUFFMAN] Huffman, D., "A Method for the Construction of Minimum Redundancy Codes", Proceedings of the Institute of Radio Engineers Volume 40, Number 9, pp. 1098-1101, September 1952, <<http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=4051119>>.
- [ORIGIN] Barth, A., "The Web Origin Concept", RFC 6454, December 2011.
- [PETAL] Tan, J. and J. Nahata, "PETAL: Preset Encoding Table Information Leakage", April 2013, <<http://www.pdl.cmu.edu/PDL-FTP/associated/CMU-PDL-13-106.pdf>>.
- [SPDY] Belshe, M. and R. Peon, "SPDY Protocol", draft-mbelshe-httpbis-spdy-00 (work in progress), February 2012.
- [TLS12] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

Appendix A. Static Table Definition

The static table (see Section 2.3.1) consists in a predefined and unchangeable list of header fields.

The static table was created from the most frequent header fields used by popular web sites, with the addition of HTTP/2-specific pseudo-header fields (see Section 8.1.2.1 of [HTTP2]). For header fields with a few frequent values, an entry was added for each of these frequent values. For other header fields, an entry was added with an empty value.

The following table lists the predefined header fields that make-up the static table.

Index	Header Name	Header Value
1	:authority	
2	:method	GET
3	:method	POST
4	:path	/
5	:path	/index.html
6	:scheme	http
7	:scheme	https
8	:status	200
9	:status	204
10	:status	206
11	:status	304
12	:status	400
13	:status	404
14	:status	500
15	accept-charset	
16	accept-encoding	gzip, deflate
17	accept-language	
18	accept-ranges	
19	accept	
20	access-control-allow-origin	
21	age	
22	allow	
23	authorization	
24	cache-control	
25	content-disposition	
26	content-encoding	
27	content-language	
28	content-length	
29	content-location	
30	content-range	

31	content-type
32	cookie
33	date
34	etag
35	expect
36	expires
37	from
38	host
39	if-match
40	if-modified-since
41	if-none-match
42	if-range
43	if-unmodified-since
44	last-modified
45	link
46	location
47	max-forwards
48	proxy-authenticate
49	proxy-authorization
50	range
51	referer
52	refresh
53	retry-after
54	server
55	set-cookie
56	strict-transport-security
57	transfer-encoding
58	user-agent
59	vary
60	via
61	www-authenticate

Table 1: Static Table Entries

Table 1 gives the index of each entry in the static table.

Appendix B. Huffman Code

The following Huffman code is used when encoding string literals with a Huffman coding (see Section 5.2).

This Huffman code was generated from statistics obtained on a large sample of HTTP headers. It is a canonical Huffman code (see [CANONICAL]) with some tweaking to ensure that no symbol has a unique code length.

Each row in the table defines the code used to represent a symbol:

sym: The symbol to be represented. It is the decimal value of an octet, possibly prepended with its ASCII representation. A specific symbol, "EOS", is used to indicate the end of a string literal.

code as bits: The Huffman code for the symbol represented as a base-2 integer, aligned on the most significant bit (MSB).

code as hex: The Huffman code for the symbol, represented as a hexadecimal integer, aligned on the least significant bit (LSB).

len: The number of bits for the code representing the symbol.

As an example, the code for the symbol 47 (corresponding to the ASCII character "/") consists in the 6 bits "0", "1", "1", "0", "0", "0". This corresponds to the value 0x18 (in hexadecimal) encoded in 6 bits.

sym	code as bits aligned to MSB				code as hex aligned to LSB	len in bits
(0)	11111111	11000			1ff8	[13]
(1)	11111111	11111111	1011000		7fffd8	[23]
(2)	11111111	11111111	11111110	0010	fffffe2	[28]
(3)	11111111	11111111	11111110	0011	fffffe3	[28]
(4)	11111111	11111111	11111110	0100	fffffe4	[28]
(5)	11111111	11111111	11111110	0101	fffffe5	[28]
(6)	11111111	11111111	11111110	0110	fffffe6	[28]
(7)	11111111	11111111	11111110	0111	fffffe7	[28]
(8)	11111111	11111111	11111110	1000	fffffe8	[28]
(9)	11111111	11111111	11101010		ffffea	[24]
(10)	11111111	11111111	11111111	111100	3ffffffc	[30]
(11)	11111111	11111111	11111110	1001	fffffe9	[28]
(12)	11111111	11111111	11111110	1010	fffffea	[28]
(13)	11111111	11111111	11111111	111101	3ffffffd	[30]
(14)	11111111	11111111	11111110	1011	fffffeb	[28]
(15)	11111111	11111111	11111110	1100	fffffec	[28]
(16)	11111111	11111111	11111110	1101	fffffed	[28]
(17)	11111111	11111111	11111110	1110	fffffee	[28]
(18)	11111111	11111111	11111110	1111	fffffef	[28]
(19)	11111111	11111111	11111111	0000	ffffff0	[28]
(20)	11111111	11111111	11111111	0001	ffffff1	[28]
(21)	11111111	11111111	11111111	0010	ffffff2	[28]
(22)	11111111	11111111	11111111	111110	3ffffffe	[30]
(23)	11111111	11111111	11111111	0011	ffffff3	[28]
(24)	11111111	11111111	11111111	0100	ffffff4	[28]
(25)	11111111	11111111	11111111	0101	ffffff5	[28]

(26)	11111111 11111111 11111111 0110	ffffff6 [28]
(27)	11111111 11111111 11111111 0111	ffffff7 [28]
(28)	11111111 11111111 11111111 1000	ffffff8 [28]
(29)	11111111 11111111 11111111 1001	ffffff9 [28]
(30)	11111111 11111111 11111111 1010	ffffffa [28]
(31)	11111111 11111111 11111111 1011	ffffffb [28]
' ' (32)	010100	14 [6]
'!' (33)	11111110 00	3f8 [10]
'"' (34)	11111110 01	3f9 [10]
'#' (35)	11111111 1010	ffa [12]
'\$' (36)	11111111 11001	1ff9 [13]
'%' (37)	010101	15 [6]
'&' (38)	11111000	f8 [8]
' ' (39)	11111111 010	7fa [11]
'(' (40)	11111110 10	3fa [10]
')' (41)	11111110 11	3fb [10]
'*' (42)	11111001	f9 [8]
'+' (43)	11111111 011	7fb [11]
',' (44)	11111010	fa [8]
'-' (45)	010110	16 [6]
'.' (46)	010111	17 [6]
'/' (47)	011000	18 [6]
'0' (48)	00000	0 [5]
'1' (49)	00001	1 [5]
'2' (50)	00010	2 [5]
'3' (51)	011001	19 [6]
'4' (52)	011010	1a [6]
'5' (53)	011011	1b [6]
'6' (54)	011100	1c [6]
'7' (55)	011101	1d [6]
'8' (56)	011110	1e [6]
'9' (57)	011111	1f [6]
':' (58)	1011100	5c [7]
';' (59)	11111011	fb [8]
'<' (60)	11111111 1111100	7ffc [15]
'=' (61)	100000	20 [6]
'>' (62)	11111111 1011	ffb [12]
'?' (63)	11111111 00	3fc [10]
'@' (64)	11111111 11010	1ffa [13]
'A' (65)	100001	21 [6]
'B' (66)	1011101	5d [7]
'C' (67)	1011110	5e [7]
'D' (68)	1011111	5f [7]
'E' (69)	1100000	60 [7]
'F' (70)	1100001	61 [7]
'G' (71)	1100010	62 [7]
'H' (72)	1100011	63 [7]
'I' (73)	1100100	64 [7]

'J' (74)	1100101	65 [7]
'K' (75)	1100110	66 [7]
'L' (76)	1100111	67 [7]
'M' (77)	1101000	68 [7]
'N' (78)	1101001	69 [7]
'O' (79)	1101010	6a [7]
'P' (80)	1101011	6b [7]
'Q' (81)	1101100	6c [7]
'R' (82)	1101101	6d [7]
'S' (83)	1101110	6e [7]
'T' (84)	1101111	6f [7]
'U' (85)	1110000	70 [7]
'V' (86)	1110001	71 [7]
'W' (87)	1110010	72 [7]
'X' (88)	11111100	fc [8]
'Y' (89)	1110011	73 [7]
'Z' (90)	11111101	fd [8]
'[' (91)	11111111 11011	1fffb [13]
'\' (92)	11111111 11111110 000	7fff0 [19]
']' (93)	11111111 11100	1fffc [13]
'^' (94)	11111111 111100	3fffc [14]
'_' (95)	100010	22 [6]
'`' (96)	11111111 1111101	7fffd [15]
'a' (97)	00011	3 [5]
'b' (98)	100011	23 [6]
'c' (99)	00100	4 [5]
'd' (100)	100100	24 [6]
'e' (101)	00101	5 [5]
'f' (102)	100101	25 [6]
'g' (103)	100110	26 [6]
'h' (104)	100111	27 [6]
'i' (105)	00110	6 [5]
'j' (106)	1110100	74 [7]
'k' (107)	1110101	75 [7]
'l' (108)	101000	28 [6]
'm' (109)	101001	29 [6]
'n' (110)	101010	2a [6]
'o' (111)	00111	7 [5]
'p' (112)	101011	2b [6]
'q' (113)	1110110	76 [7]
'r' (114)	101100	2c [6]
's' (115)	01000	8 [5]
't' (116)	01001	9 [5]
'u' (117)	101101	2d [6]
'v' (118)	1110111	77 [7]
'w' (119)	1111000	78 [7]
'x' (120)	1111001	79 [7]
'y' (121)	1111010	7a [7]

'z' (122)	1111011			7b [7]
'{' (123)	11111111 1111110			7ffe [15]
' ' (124)	11111111 100			7fc [11]
'}' (125)	11111111 111101			3ffd [14]
'~' (126)	11111111 11101			1ffd [13]
(127)	11111111 11111111 11111111 1100			ffffffc [28]
(128)	11111111 11111110 0110			fffe6 [20]
(129)	11111111 11111111 010010			3ffd2 [22]
(130)	11111111 11111110 0111			fffe7 [20]
(131)	11111111 11111110 1000			fffe8 [20]
(132)	11111111 11111111 010011			3ffd3 [22]
(133)	11111111 11111111 010100			3ffd4 [22]
(134)	11111111 11111111 010101			3ffd5 [22]
(135)	11111111 11111111 1011001			7ffd9 [23]
(136)	11111111 11111111 010110			3ffd6 [22]
(137)	11111111 11111111 1011010			7ffdfa [23]
(138)	11111111 11111111 1011011			7ffddb [23]
(139)	11111111 11111111 1011100			7fffdc [23]
(140)	11111111 11111111 1011101			7ffdd [23]
(141)	11111111 11111111 1011110			7ffdde [23]
(142)	11111111 11111111 11101011			ffffeb [24]
(143)	11111111 11111111 1011111			7ffddf [23]
(144)	11111111 11111111 11101100			ffffec [24]
(145)	11111111 11111111 11101101			ffffed [24]
(146)	11111111 11111111 010111			3ffd7 [22]
(147)	11111111 11111111 1100000			7ffe0 [23]
(148)	11111111 11111111 11101110			ffffee [24]
(149)	11111111 11111111 1100001			7ffe1 [23]
(150)	11111111 11111111 1100010			7ffe2 [23]
(151)	11111111 11111111 1100011			7ffe3 [23]
(152)	11111111 11111111 1100100			7ffe4 [23]
(153)	11111111 11111110 11100			1fffdc [21]
(154)	11111111 11111111 011000			3ffd8 [22]
(155)	11111111 11111111 1100101			7ffe5 [23]
(156)	11111111 11111111 011001			3ffd9 [22]
(157)	11111111 11111111 1100110			7ffe6 [23]
(158)	11111111 11111111 1100111			7ffe7 [23]
(159)	11111111 11111111 11101111			ffffef [24]
(160)	11111111 11111111 011010			3ffdfa [22]
(161)	11111111 11111110 11101			1ffdd [21]
(162)	11111111 11111110 1001			fffe9 [20]
(163)	11111111 11111111 011011			3ffddb [22]
(164)	11111111 11111111 011100			3fffdc [22]
(165)	11111111 11111111 1101000			7ffe8 [23]
(166)	11111111 11111111 1101001			7ffe9 [23]
(167)	11111111 11111110 11110			1ffde [21]
(168)	11111111 11111111 1101010			7fffea [23]
(169)	11111111 11111111 011101			3ffdd [22]

(170)	11111111 11111111 011110	3ffffde [22]
(171)	11111111 11111111 11110000	ffffff0 [24]
(172)	11111111 11111110 11111	1fffdff [21]
(173)	11111111 11111111 011111	3fffdff [22]
(174)	11111111 11111111 1101011	7ffffeb [23]
(175)	11111111 11111111 1101100	7ffffec [23]
(176)	11111111 11111111 00000	1ffffe0 [21]
(177)	11111111 11111111 00001	1ffffe1 [21]
(178)	11111111 11111111 100000	3ffffe0 [22]
(179)	11111111 11111111 00010	1ffffe2 [21]
(180)	11111111 11111111 1101101	7ffffed [23]
(181)	11111111 11111111 100001	3ffffe1 [22]
(182)	11111111 11111111 1101110	7ffffee [23]
(183)	11111111 11111111 1101111	7ffffef [23]
(184)	11111111 11111110 1010	fffea [20]
(185)	11111111 11111111 100010	3ffffe2 [22]
(186)	11111111 11111111 100011	3ffffe3 [22]
(187)	11111111 11111111 100100	3ffffe4 [22]
(188)	11111111 11111111 1110000	7fffff0 [23]
(189)	11111111 11111111 100101	3ffffe5 [22]
(190)	11111111 11111111 100110	3ffffe6 [22]
(191)	11111111 11111111 1110001	7fffff1 [23]
(192)	11111111 11111111 11111000 00	3ffffe0 [26]
(193)	11111111 11111111 11111000 01	3ffffe1 [26]
(194)	11111111 11111110 1011	fffeb [20]
(195)	11111111 11111110 001	7ffff1 [19]
(196)	11111111 11111111 100111	3ffffe7 [22]
(197)	11111111 11111111 1110010	7fffff2 [23]
(198)	11111111 11111111 101000	3ffffe8 [22]
(199)	11111111 11111111 11110110 0	1fffffec [25]
(200)	11111111 11111111 11111000 10	3ffffe2 [26]
(201)	11111111 11111111 11111000 11	3ffffe3 [26]
(202)	11111111 11111111 11111001 00	3ffffe4 [26]
(203)	11111111 11111111 11111011 110	7ffffde [27]
(204)	11111111 11111111 11111011 111	7ffffdf [27]
(205)	11111111 11111111 11111001 01	3ffffe5 [26]
(206)	11111111 11111111 11110001	fffff1 [24]
(207)	11111111 11111111 11110110 1	1ffffed [25]
(208)	11111111 11111110 010	7fff2 [19]
(209)	11111111 11111111 00011	1ffffe3 [21]
(210)	11111111 11111111 11111001 10	3ffffe6 [26]
(211)	11111111 11111111 11111100 000	7ffffe0 [27]
(212)	11111111 11111111 11111100 001	7ffffe1 [27]
(213)	11111111 11111111 11111001 11	3ffffe7 [26]
(214)	11111111 11111111 11111100 010	7ffffe2 [27]
(215)	11111111 11111111 11110010	fffff2 [24]
(216)	11111111 11111111 00100	1ffffe4 [21]
(217)	11111111 11111111 00101	1ffffe5 [21]

(218)	11111111	11111111	11111010	00	3ffffe8	[26]
(219)	11111111	11111111	11111010	01	3ffffe9	[26]
(220)	11111111	11111111	11111111	1101	ffffffd	[28]
(221)	11111111	11111111	11111100	011	7ffffe3	[27]
(222)	11111111	11111111	11111100	100	7ffffe4	[27]
(223)	11111111	11111111	11111100	101	7ffffe5	[27]
(224)	11111111	11111111	1100		fffec	[20]
(225)	11111111	11111111	11110011		fffff3	[24]
(226)	11111111	11111111	1101		fffed	[20]
(227)	11111111	11111111	00110		1ffffe6	[21]
(228)	11111111	11111111	101001		3ffffe9	[22]
(229)	11111111	11111111	00111		1ffffe7	[21]
(230)	11111111	11111111	01000		1ffffe8	[21]
(231)	11111111	11111111	1110011		7fffff3	[23]
(232)	11111111	11111111	101010		3fffea	[22]
(233)	11111111	11111111	101011		3fffeb	[22]
(234)	11111111	11111111	11110111	0	1ffffee	[25]
(235)	11111111	11111111	11110111	1	1ffffef	[25]
(236)	11111111	11111111	11110100		fffff4	[24]
(237)	11111111	11111111	11110101		fffff5	[24]
(238)	11111111	11111111	11111010	10	3ffffea	[26]
(239)	11111111	11111111	1110100		7fffff4	[23]
(240)	11111111	11111111	11111010	11	3ffffeb	[26]
(241)	11111111	11111111	11111100	110	7ffffe6	[27]
(242)	11111111	11111111	11111011	00	3ffffec	[26]
(243)	11111111	11111111	11111011	01	3ffffed	[26]
(244)	11111111	11111111	11111100	111	7ffffe7	[27]
(245)	11111111	11111111	11111101	000	7ffffe8	[27]
(246)	11111111	11111111	11111101	001	7ffffe9	[27]
(247)	11111111	11111111	11111101	010	7ffffea	[27]
(248)	11111111	11111111	11111101	011	7ffffeb	[27]
(249)	11111111	11111111	11111111	1110	ffffffe	[28]
(250)	11111111	11111111	11111101	100	7ffffec	[27]
(251)	11111111	11111111	11111101	101	7ffffed	[27]
(252)	11111111	11111111	11111101	110	7ffffee	[27]
(253)	11111111	11111111	11111101	111	7ffffef	[27]
(254)	11111111	11111111	11111110	000	7fffff0	[27]
(255)	11111111	11111111	11111011	10	3ffffee	[26]
EOS (256)	11111111	11111111	11111111	111111	3fffffff	[30]

Appendix C. Examples

A number of examples are worked through here, covering integer encoding, header field representation, and the encoding of whole lists of header fields, for both requests and responses, and with and without Huffman coding.

C.1. Integer Representation Examples

This section shows the representation of integer values in details (see Section 5.1).

C.1.1. Example 1: Encoding 10 Using a 5-bit Prefix

The value 10 is to be encoded with a 5-bit prefix.

- o 10 is less than 31 ($2^5 - 1$) and is represented using the 5-bit prefix.

0	1	2	3	4	5	6	7	
X	X	X	0	1	0	1	0	10 stored on 5 bits

C.1.2. Example 2: Encoding 1337 Using a 5-bit Prefix

The value I=1337 is to be encoded with a 5-bit prefix.

1337 is greater than 31 ($2^5 - 1$).

The 5-bit prefix is filled with its max value (31).

$I = 1337 - (2^5 - 1) = 1306$.

I (1306) is greater than or equal to 128, the while loop body executes:

$I \% 128 == 26$

$26 + 128 == 154$

154 is encoded in 8 bits as: 10011010

I is set to 10 ($1306 / 128 == 10$)

I is no longer greater than or equal to 128, the while loop terminates.

I, now 10, is encoded in 8 bits as: 00001010.

The process ends.

0	1	2	3	4	5	6	7	
X	X	X	1	1	1	1	1	Prefix = 31, I = 1306
1	0	0	1	1	0	1	0	1306>=128, encode(154), I=1306/128
0	0	0	0	1	0	1	0	10<128, encode(10), done

C.1.3. Example 3: Encoding 42 Starting at an Octet Boundary

The value 42 is to be encoded starting at an octet-boundary. This implies that a 8-bit prefix is used.

- o 42 is less than 255 ($2^8 - 1$) and is represented using the 8-bit prefix.

0	1	2	3	4	5	6	7	
0	0	1	0	1	0	1	0	42 stored on 8 bits

C.2. Header Field Representation Examples

This section shows several independent representation examples.

C.2.1. Literal Header Field with Indexing

The header field representation uses a literal name and a literal value. The header field is added to the dynamic table.

Header list to encode:

custom-key: custom-header

Hex dump of encoded data:

400a 6375 7374 6f6d 2d6b 6579 0d63 7573	@.custom-key.cus
746f 6d2d 6865 6164 6572	tom-header

Decoding process:

40		== Literal indexed ==
0a		Literal name (len = 10)
6375 7374 6f6d 2d6b 6579		custom-key
0d		Literal value (len = 13)
6375 7374 6f6d 2d68 6561 6465 72		custom-header
		-> custom-key: custom-head\
		er

Dynamic Table (after decoding):

```
[ 1] (s = 55) custom-key: custom-header
      Table size: 55
```

Decoded header list:

custom-key: custom-header

C.2.2. Literal Header Field without Indexing

The header field representation uses an indexed name and a literal value. The header field is not added to the dynamic table.

Header list to encode:

:path: /sample/path

Hex dump of encoded data:

040c 2f73 616d 706c 652f 7061 7468		../sample/path
------------------------------------	--	----------------

Decoding process:

04		== Literal not indexed ==
		Indexed name (idx = 4)
		:path
0c		Literal value (len = 12)
2f73 616d 706c 652f 7061 7468		/sample/path
		-> :path: /sample/path

Dynamic table (after decoding): empty.

Decoded header list:

:path: /sample/path

C.2.3. Literal Header Field never Indexed

The header field representation uses a literal name and a literal value. The header field is not added to the dynamic table, and must use the same representation if re-encoded by an intermediary.

Header list to encode:

password: secret

Hex dump of encoded data:

1008 7061 7373 776f 7264 0673 6563 7265		..password.secre
74		t

Decoding process:

10		== Literal never indexed ==
08		Literal name (len = 8)
7061 7373 776f 7264		password
06		Literal value (len = 6)
7365 6372 6574		secret
		-> password: secret

Dynamic table (after decoding): empty.

Decoded header list:

password: secret

C.2.4. Indexed Header Field

The header field representation uses an indexed header field, from the static table.

Header list to encode:

:method: GET

Hex dump of encoded data:

82		.
----	--	---

Decoding process:

82		== Indexed - Add ==
		idx = 2
		-> :method: GET

Dynamic table (after decoding): empty.

Decoded header list:

:method: GET

C.3. Request Examples without Huffman Coding

This section shows several consecutive header lists, corresponding to HTTP requests, on the same connection.

C.3.1. First Request

Header list to encode:

:method: GET
:scheme: http
:path: /
:authority: www.example.com

Hex dump of encoded data:

8286 8441 0f77 7777 2e65 7861 6d70 6c65	...A.www.example
2e63 6f6d	.com

Decoding process:

82	== Indexed - Add ==
	idx = 2
	-> :method: GET
86	== Indexed - Add ==
	idx = 6
	-> :scheme: http
84	== Indexed - Add ==
	idx = 4
	-> :path: /
41	== Literal indexed ==
	Indexed name (idx = 1)
	:authority
0f	Literal value (len = 15)
7777 772e 6578 616d 706c 652e 636f 6d	www.example.com
	-> :authority: www.example\
	.com

Dynamic Table (after decoding):

[1] (s = 57) :authority: www.example.com
Table size: 57

Decoded header list:

```
:method: GET
:scheme: http
:path: /
:authority: www.example.com
```

C.3.2. Second Request

Header list to encode:

```
:method: GET
:scheme: http
:path: /
:authority: www.example.com
cache-control: no-cache
```

Hex dump of encoded data:

```
8286 84be 5808 6e6f 2d63 6163 6865      | ....X.no-cache
```

Decoding process:

82	== Indexed - Add == idx = 2
86	-> :method: GET == Indexed - Add == idx = 6
84	-> :scheme: http == Indexed - Add == idx = 4
be	-> :path: / == Indexed - Add == idx = 62
58	-> :authority: www.example\ .com == Literal indexed == Indexed name (idx = 24)
08	cache-control
6e6f 2d63 6163 6865	Literal value (len = 8) no-cache -> cache-control: no-cache

Dynamic Table (after decoding):

```
[ 1] (s = 53) cache-control: no-cache
[ 2] (s = 57) :authority: www.example.com
      Table size: 110
```

Decoded header list:

```
:method: GET
:scheme: http
:path: /
:authority: www.example.com
cache-control: no-cache
```

C.3.3. Third Request

Header list to encode:

```
:method: GET
:scheme: https
:path: /index.html
:authority: www.example.com
custom-key: custom-value
```

Hex dump of encoded data:

```
8287 85bf 400a 6375 7374 6f6d 2d6b 6579 | ....@.custom-key
0c63 7573 746f 6d2d 7661 6c75 65      | .custom-value
```

Decoding process:

```
82      | == Indexed - Add ==
          |     idx = 2
          |     -> :method: GET
87      | == Indexed - Add ==
          |     idx = 7
          |     -> :scheme: https
85      | == Indexed - Add ==
          |     idx = 5
          |     -> :path: /index.html
bf      | == Indexed - Add ==
          |     idx = 63
          |     -> :authority: www.example\
          |     .com
40      | == Literal indexed ==
0a      |     Literal name (len = 10)
6375 7374 6f6d 2d6b 6579 |     custom-key
0c      |     Literal value (len = 12)
6375 7374 6f6d 2d76 616c 7565 |     custom-value
          |     -> custom-key: custom-valu\
          |     e
```


Dynamic Table (after decoding):

```
[ 1] (s = 54) custom-key: custom-value
[ 2] (s = 53) cache-control: no-cache
[ 3] (s = 57) :authority: www.example.com
      Table size: 164
```

Decoded header list:

```
:method: GET
:scheme: https
:path: /index.html
:authority: www.example.com
custom-key: custom-value
```

C.4. Request Examples with Huffman Coding

This section shows the same examples as the previous section, but using Huffman encoding for the literal values.

C.4.1. First Request

Header list to encode:

```
:method: GET
:scheme: http
:path: /
:authority: www.example.com
```

Hex dump of encoded data:

```
8286 8441 8cf1 e3c2 e5f2 3a6b a0ab 90f4 | ...A.....:k....
ff                                     | .
```

Decoding process:

82 86 84 41 8c f1e3 c2e5 f23a 6ba0 ab90 f4ff	== Indexed - Add == idx = 2 -> :method: GET == Indexed - Add == idx = 6 -> :scheme: http == Indexed - Add == idx = 4 -> :path: / == Literal indexed == Indexed name (idx = 1) :authority Literal value (len = 12) Huffman encoded::k..... Decoded: www.example.com -> :authority: www.example\ .com
---	---

Dynamic Table (after decoding):

```
[ 1] (s = 57) :authority: www.example.com
      Table size: 57
```

Decoded header list:

```
:method: GET
:scheme: http
:path: /
:authority: www.example.com
```

C.4.2. Second Request

Header list to encode:

```
:method: GET
:scheme: http
:path: /
:authority: www.example.com
cache-control: no-cache
```

Hex dump of encoded data:

8286 84be 5886 a8eb 1064 9cbfX....d..
-------------------------------	--------------

Decoding process:

82 86 84 be 58 86 a8eb 1064 9cbf	== Indexed - Add == idx = 2 -> :method: GET == Indexed - Add == idx = 6 -> :scheme: http == Indexed - Add == idx = 4 -> :path: / == Indexed - Add == idx = 62 -> :authority: www.example\ .com == Literal indexed == Indexed name (idx = 24) cache-control Literal value (len = 6) Huffman encoded: ...d.. Decoded: no-cache -> cache-control: no-cache
--	--

Dynamic Table (after decoding):

```
[ 1] (s = 53) cache-control: no-cache
[ 2] (s = 57) :authority: www.example.com
      Table size: 110
```

Decoded header list:

```
:method: GET
:scheme: http
:path: /
:authority: www.example.com
cache-control: no-cache
```

C.4.3. Third Request

Header list to encode:

```
:method: GET
:scheme: https
:path: /index.html
:authority: www.example.com
custom-key: custom-value
```

Hex dump of encoded data:

```
8287 85bf 4088 25a8 49e9 5ba9 7d7f 8925 | ....@.%.I.[.}..%
a849 e95b b8e8 b4bf | .I.[....
```

Decoding process:

```
82 | == Indexed - Add ==
   |     idx = 2
87 | -> :method: GET
   | == Indexed - Add ==
   |     idx = 7
   | -> :scheme: https
85 | == Indexed - Add ==
   |     idx = 5
   | -> :path: /index.html
bf | == Indexed - Add ==
   |     idx = 63
   | -> :authority: www.example\
   |     .com
40 | == Literal indexed ==
88 |     Literal name (len = 8)
   |     Huffman encoded:
25a8 49e9 5ba9 7d7f | %.I.[.}.
   |     Decoded:
   | custom-key
89 |     Literal value (len = 9)
   |     Huffman encoded:
25a8 49e9 5bb8 e8b4 bf | %.I.[....
   |     Decoded:
   | custom-value
   | -> custom-key: custom-valu\
   |     e
```

Dynamic Table (after decoding):

```
[ 1] (s = 54) custom-key: custom-value
[ 2] (s = 53) cache-control: no-cache
[ 3] (s = 57) :authority: www.example.com
      Table size: 164
```

Decoded header list:

```
:method: GET
:scheme: https
:path: /index.html
:authority: www.example.com
custom-key: custom-value
```

C.5. Response Examples without Huffman Coding

This section shows several consecutive header lists, corresponding to HTTP responses, on the same connection. The HTTP/2 setting parameter SETTINGS_HEADER_TABLE_SIZE is set to the value of 256 octets, causing some evictions to occur.

C.5.1. First Response

Header list to encode:

```
:status: 302
cache-control: private
date: Mon, 21 Oct 2013 20:13:21 GMT
location: https://www.example.com
```

Hex dump of encoded data:

4803	3330	3258	0770	7269	7661	7465	611d		H.302X.privatea.
4d6f	6e2c	2032	3120	4f63	7420	3230	3133		Mon, 21 Oct 2013
2032	303a	3133	3a32	3120	474d	546e	1768		20:13:21 GMTn.h
7474	7073	3a2f	2f77	7777	2e65	7861	6d70		ttps://www.examp
6c65	2e63	6f6d							le.com

Decoding process:

48 03 3330 32 58 07 7072 6976 6174 65 61 1d 4d6f 6e2c 2032 3120 4f63 7420 3230 3133 2032 303a 3133 3a32 3120 474d 54 6e 17 6874 7470 733a 2f2f 7777 772e 6578 616d 706c 652e 636f 6d	== Literal indexed == Indexed name (idx = 8) :status Literal value (len = 3) 302 -> :status: 302 == Literal indexed == Indexed name (idx = 24) cache-control Literal value (len = 7) private -> cache-control: private == Literal indexed == Indexed name (idx = 33) date Literal value (len = 29) Mon, 21 Oct 2013 20:13:21 GMT -> date: Mon, 21 Oct 2013 \ 20:13:21 GMT == Literal indexed == Indexed name (idx = 46) location Literal value (len = 23) https://www.exam ple.com -> location: https://www.e\ xample.com
---	--

Dynamic Table (after decoding):

```
[ 1] (s = 63) location: https://www.example.com
[ 2] (s = 65) date: Mon, 21 Oct 2013 20:13:21 GMT
[ 3] (s = 52) cache-control: private
[ 4] (s = 42) :status: 302
    Table size: 222
```

Decoded header list:

```
:status: 302
cache-control: private
date: Mon, 21 Oct 2013 20:13:21 GMT
location: https://www.example.com
```

C.5.2. Second Response

The (":status", "302") header field is evicted from the dynamic table to free space to allow adding the (":status", "307") header field.

Header list to encode:

```
:status: 307
cache-control: private
date: Mon, 21 Oct 2013 20:13:21 GMT
location: https://www.example.com
```

Hex dump of encoded data:

```
4803 3330 37c1 c0bf | H.307...
```

Decoding process:

48		== Literal indexed ==
		Indexed name (idx = 8)
		:status
03		Literal value (len = 3)
3330 37		307
		- evict: :status: 302
		-> :status: 307
c1		== Indexed - Add ==
		idx = 65
		-> cache-control: private
c0		== Indexed - Add ==
		idx = 64
		-> date: Mon, 21 Oct 2013 \
		20:13:21 GMT
bf		== Indexed - Add ==
		idx = 63
		-> location: https://www.e\
		xample.com

Dynamic Table (after decoding):

```
[ 1] (s = 42) :status: 307
[ 2] (s = 63) location: https://www.example.com
[ 3] (s = 65) date: Mon, 21 Oct 2013 20:13:21 GMT
[ 4] (s = 52) cache-control: private
Table size: 222
```

Decoded header list:

```
:status: 307
cache-control: private
date: Mon, 21 Oct 2013 20:13:21 GMT
location: https://www.example.com
```

C.5.3. Third Response

Several header fields are evicted from the dynamic table during the processing of this header list.

Header list to encode:

```
:status: 200
cache-control: private
date: Mon, 21 Oct 2013 20:13:22 GMT
location: https://www.example.com
content-encoding: gzip
set-cookie: foo=ASDJKHQKBZXOQWEOPIUAXQWEOIU; max-age=3600; version=1
```

Hex dump of encoded data:

88c1	611d	4d6f	6e2c	2032	3120	4f63	7420		..a.Mon, 21 Oct
3230	3133	2032	303a	3133	3a32	3220	474d		2013 20:13:22 GM
54c0	5a04	677a	6970	7738	666f	6f3d	4153		T.Z.gzipw8foo=AS
444a	4b48	514b	425a	584f	5157	454f	5049		DJKHQKBZXOQWEOPI
5541	5851	5745	4f49	553b	206d	6178	2d61		UAXQWEOIU; max-a
6765	3d33	3630	303b	2076	6572	7369	6f6e		ge=3600; version
3d31									=1

Decoding process:

88	== Indexed - Add == idx = 8 -> :status: 200
c1	== Indexed - Add == idx = 65 -> cache-control: private
61	== Literal indexed == Indexed name (idx = 33) date
1d	Literal value (len = 29)
4d6f 6e2c 2032 3120 4f63 7420 3230 3133	Mon, 21 Oct 2013
2032 303a 3133 3a32 3220 474d 54	20:13:22 GMT
	- evict: cache-control: pr\ivate
	-> date: Mon, 21 Oct 2013 \20:13:22 GMT
c0	== Indexed - Add == idx = 64 -> location: https://www.e\sample.com
5a	== Literal indexed == Indexed name (idx = 26) content-encoding
04	Literal value (len = 4)
677a 6970	gzip
	- evict: date: Mon, 21 Oct\2013 20:13:21 GMT
	-> content-encoding: gzip
77	== Literal indexed == Indexed name (idx = 55) set-cookie
38	Literal value (len = 56)
666f 6f3d 4153 444a 4b48 514b 425a 584f	foo=ASDJKHQKBZXO
5157 454f 5049 5541 5851 5745 4f49 553b	QWEOPIUAXQWEOIU;
206d 6178 2d61 6765 3d33 3630 303b 2076	max-age=3600; v
6572 7369 6f6e 3d31	ersion=1
	- evict: location: https://\www.example.com
	- evict: :status: 307
	-> set-cookie: foo=ASDJKHQ\KBZXOQWEOPIUAXQWEOIU; ma\ x-age=3600; version=1

Dynamic Table (after decoding):

```
[ 1] (s = 98) set-cookie: foo=ASDJKHQKBZXOQWEOPIUAXQWEIOIU; max-age\
    =3600; version=1
[ 2] (s = 52) content-encoding: gzip
[ 3] (s = 65) date: Mon, 21 Oct 2013 20:13:22 GMT
    Table size: 215
```

Decoded header list:

```
:status: 200
cache-control: private
date: Mon, 21 Oct 2013 20:13:22 GMT
location: https://www.example.com
content-encoding: gzip
set-cookie: foo=ASDJKHQKBZXOQWEOPIUAXQWEIOIU; max-age=3600; version=1
```

C.6. Response Examples with Huffman Coding

This section shows the same examples as the previous section, but using Huffman encoding for the literal values. The HTTP/2 setting parameter `SETTINGS_HEADER_TABLE_SIZE` is set to the value of 256 octets, causing some evictions to occur. The eviction mechanism uses the length of the decoded literal values, so the same evictions occurs as in the previous section.

C.6.1. First Response

Header list to encode:

```
:status: 302
cache-control: private
date: Mon, 21 Oct 2013 20:13:21 GMT
location: https://www.example.com
```

Hex dump of encoded data:

4882	6402	5885	aec3	771a	4b61	96d0	7abe		H.d.X...w.Ka..z.
9410	54d4	44a8	2005	9504	0b81	66e0	82a6		..T.D.f...
2dlb	ff6e	919d	29ad	1718	63c7	8f0b	97c8		-..n..)....c.....
e9ae	82ae	43d3						C.

Decoding process:

48	82	6402	58	85	aec3 771a 4b	61	96	d07a be94 1054 d444 a820 0595 040b 8166 e082 a62d 1bff	6e	91	9d29 ad17 1863 c78f 0b97 c8e9 ae82 ae43 d3	<pre> == Literal indexed == Indexed name (idx = 8) :status Literal value (len = 2) Huffman encoded: d. Decoded: 302 -> :status: 302 == Literal indexed == Indexed name (idx = 24) cache-control Literal value (len = 5) Huffman encoded: ..w.K Decoded: private -> cache-control: private == Literal indexed == Indexed name (idx = 33) date Literal value (len = 22) Huffman encoded: .z...T.D.f ...-... Decoded: Mon, 21 Oct 2013 20:13:21 \ GMT -> date: Mon, 21 Oct 2013 \ 20:13:21 GMT == Literal indexed == Indexed name (idx = 46) location Literal value (len = 17) Huffman encoded: .)...c.....C . Decoded: https://www.example.com -> location: https://www.e\ xample.com </pre>
----	----	------	----	----	--------------	----	----	---	----	----	---	---

Dynamic Table (after decoding):

```
[ 1] (s = 63) location: https://www.example.com
[ 2] (s = 65) date: Mon, 21 Oct 2013 20:13:21 GMT
[ 3] (s = 52) cache-control: private
[ 4] (s = 42) :status: 302
    Table size: 222
```

Decoded header list:

```
:status: 302
cache-control: private
date: Mon, 21 Oct 2013 20:13:21 GMT
location: https://www.example.com
```

C.6.2. Second Response

The (":status", "302") header field is evicted from the dynamic table to free space to allow adding the (":status", "307") header field.

Header list to encode:

```
:status: 307
cache-control: private
date: Mon, 21 Oct 2013 20:13:21 GMT
location: https://www.example.com
```

Hex dump of encoded data:

```
4883 640e ffc1 c0bf          | H.d.....
```

Decoding process:

48		== Literal indexed ==
		Indexed name (idx = 8)
		:status
83		Literal value (len = 3)
		Huffman encoded:
640e ff		d..
		Decoded:
		307
		- evict: :status: 302
		-> :status: 307
c1		== Indexed - Add ==
		idx = 65
		-> cache-control: private
c0		== Indexed - Add ==
		idx = 64
		-> date: Mon, 21 Oct 2013 \
		20:13:21 GMT
bf		== Indexed - Add ==
		idx = 63
		-> location: https://www.e\
		xample.com

Dynamic Table (after decoding):

```
[ 1] (s = 42) :status: 307
[ 2] (s = 63) location: https://www.example.com
[ 3] (s = 65) date: Mon, 21 Oct 2013 20:13:21 GMT
[ 4] (s = 52) cache-control: private
      Table size: 222
```

Decoded header list:

```
:status: 307
cache-control: private
date: Mon, 21 Oct 2013 20:13:21 GMT
location: https://www.example.com
```

C.6.3. Third Response

Several header fields are evicted from the dynamic table during the processing of this header list.

Header list to encode:

```
:status: 200
cache-control: private
date: Mon, 21 Oct 2013 20:13:22 GMT
location: https://www.example.com
content-encoding: gzip
set-cookie: foo=ASDJKHQBZXOQWEOPIUAXQWEOIU; max-age=3600; version=1
```

Hex dump of encoded data:

88c1 6196 d07a be94 1054 d444 a820 0595	..a...z...T.D. ...
040b 8166 e084 a62d 1bff c05a 839b d9ab	...f...-...Z....
77ad 94e7 821d d7f2 e6c7 b335 dfdf cd5b	w.....5...[
3960 d5af 2708 7f36 72c1 ab27 0fb5 291f	9\'...\'..6r...\'..).
9587 3160 65c0 03ed 4ee5 b106 3d50 07	..1\'e...N...=P.

Decoding process:

88	== Indexed - Add ==
	idx = 8
	-> :status: 200
c1	== Indexed - Add ==
	idx = 65
	-> cache-control: private
61	== Literal indexed ==
	Indexed name (idx = 33)
	date
96	Literal value (len = 22)
	Huffman encoded:
d07a be94 1054 d444 a820 0595 040b 8166	.z...T.D.f
e084 a62d 1bff	...-...
	Decoded:
	Mon, 21 Oct 2013 20:13:22 \
	GMT
	- evict: cache-control: pr\
	ivate
	-> date: Mon, 21 Oct 2013 \
	20:13:22 GMT
c0	== Indexed - Add ==
	idx = 64
	-> location: https://www.e\
	xample.com
5a	== Literal indexed ==
	Indexed name (idx = 26)
	content-encoding
83	Literal value (len = 3)
	Huffman encoded:

9bd9 ab 77 ad 94e7 821d d7f2 e6c7 b335 dfdf cd5b 3960 d5af 2708 7f36 72c1 ab27 0fb5 291f 9587 3160 65c0 03ed 4ee5 b106 3d50 07	... Decoded: gzip - evict: date: Mon, 21 Oct\ 2013 20:13:21 GMT -> content-encoding: gzip == Literal indexed == Indexed name (idx = 55) set-cookie Literal value (len = 45) Huffman encoded:5...[9\ ..'...6r..'...) 1'e...N...=P. Decoded: foo=ASDJKHQKBZXOQWEOPIUAXQ\ WEOIU; max-age=3600; versi\ on=1 - evict: location: https://\ /www.example.com - evict: :status: 307 -> set-cookie: foo=ASDJKHQ\ KBZXOQWEOPIUAXQWEOIU; ma\ x-age=3600; version=1
---	--

Dynamic Table (after decoding):

```
[ 1] (s = 98) set-cookie: foo=ASDJKHQKBZXOQWEOPIUAXQWEOIU; max-age\  

=3600; version=1  

[ 2] (s = 52) content-encoding: gzip  

[ 3] (s = 65) date: Mon, 21 Oct 2013 20:13:22 GMT  

Table size: 215
```

Decoded header list:

```
:status: 200  

cache-control: private  

date: Mon, 21 Oct 2013 20:13:22 GMT  

location: https://www.example.com  

content-encoding: gzip  

set-cookie: foo=ASDJKHQKBZXOQWEOPIUAXQWEOIU; max-age=3600; version=1
```

Appendix D. Change Log (to be removed by RFC Editor before publication)

- D.1. Since draft-ietf-httpbis-header-compression-10
- o Editorial corrections for taking into account IETF LC comments.
 - * Added links to security sections.
 - * Made spec more independent of HTTP/2.
 - * Expanded security section about never indexed literal usage.
 - o Removed most usages of 'name-value pair' instead of header field.
 - o Changed 'header table' to 'header field table'.
- D.2. Since draft-ietf-httpbis-header-compression-09
- o Renamed header table to dynamic table.
 - o Updated integer representation.
 - o Editorial corrections.
- D.3. Since draft-ietf-httpbis-header-compression-08
- o Removed the reference set.
 - o Removed header emission.
 - o Explicit handling of several SETTINGS_HEADER_TABLE_SIZE parameter changes.
 - o Changed header set to header list, and forced ordering.
 - o Updated examples.
 - o Exchanged header and static table positions.
- D.4. Since draft-ietf-httpbis-header-compression-07
- o Removed old text on index value of 0.
 - o Added clarification for signalling of maximum table size after a SETTINGS_HEADER_TABLE_SIZE update.
 - o Rewrote security considerations.
 - o Many editorial clarifications or improvements.

- o Added convention section.
- o Reworked document's outline.
- o Updated static table. Entry 16 has now "gzip, deflate" for value.
- o Updated Huffman table, using data set provided by Google.

D.5. Since draft-ietf-httpbis-header-compression-06

- o Updated format to include literal headers that must never be compressed.
- o Updated security considerations.
- o Moved integer encoding examples to the appendix.
- o Updated Huffman table.
- o Updated static header table (adding and removing status values).
- o Updated examples.

D.6. Since draft-ietf-httpbis-header-compression-05

- o Regenerated examples.
- o Only one Huffman table for requests and responses.
- o Added maximum size for dynamic table, independent of SETTINGS_HEADER_TABLE_SIZE.
- o Added pseudo-code for integer decoding.
- o Improved examples (removing unnecessary removals).

D.7. Since draft-ietf-httpbis-header-compression-04

- o Updated examples: take into account changes in the spec, and show more features.
- o Use 'octet' everywhere instead of having both 'byte' and 'octet'.
- o Added reference set emptying.
- o Editorial changes and clarifications.
- o Added "host" header to the static table.

- o Ordering for list of values (either NULL- or comma-separated).
- D.8. Since draft-ietf-httpbis-header-compression-03
- o A large number of editorial changes; changed the description of evicting/adding new entries.
 - o Removed substitution indexing
 - o Changed 'initial headers' to 'static headers', as per issue #258
 - o Merged 'request' and 'response' static headers, as per issue #259
 - o Changed text to indicate that new headers are added at index 0 and expire from the largest index, as per issue #233
- D.9. Since draft-ietf-httpbis-header-compression-02
- o Corrected error in integer encoding pseudocode.
- D.10. Since draft-ietf-httpbis-header-compression-01
- o Refactored of Header Encoding Section: split definitions and processing rule.
 - o Backward incompatible change: Updated reference set management as per issue #214. This changes how the interaction between the reference set and eviction works. This also changes the working of the reference set in some specific cases.
 - o Backward incompatible change: modified initial header list, as per issue #188.
 - o Added example of 32 octets entry structure (issue #191).
 - o Added Header Set Completion section. Reflowed some text. Clarified some writing which was awkward. Added text about duplicate header entry encoding. Clarified some language w.r.t Header Set. Changed x-my-header to mynewheader. Added text in the HeaderEmission section indicating that the application may also be able to free up memory more quickly. Added information in Security Considerations section.
- D.11. Since draft-ietf-httpbis-header-compression-00
- Fixed bug/omission in integer representation algorithm.
- Changed the document title.

Header matching text rewritten.

Changed the definition of header emission.

Changed the name of the setting which dictates how much memory the compression context should use.

Removed "specific use cases" section

Corrected erroneous statement about what index can be contained in one octet

Added descriptions of opcodes

Removed security claims from introduction.

Authors' Addresses

Roberto Peon
Google, Inc

EMail: fenix@google.com

Herve Ruellan
Canon CRF

EMail: herve.ruellan@crf.canon.fr

HTTPbis Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 15, 2015

M. Belshe
Twist
R. Peon
Google, Inc
M. Thomson, Ed.
Mozilla
February 11, 2015

Hypertext Transfer Protocol version 2
draft-ietf-httpbis-http2-17

Abstract

This specification describes an optimized expression of the semantics of the Hypertext Transfer Protocol (HTTP). HTTP/2 enables a more efficient use of network resources and a reduced perception of latency by introducing header field compression and allowing multiple concurrent exchanges on the same connection. It also introduces unsolicited push of representations from servers to clients.

This specification is an alternative to, but does not obsolete, the HTTP/1.1 message syntax. HTTP's existing semantics remain unchanged.

Editorial Note (To be removed by RFC Editor)

Discussion of this draft takes place on the HTTPBIS working group mailing list (ietf-http-wg@w3.org), which is archived at [1].

Working Group information can be found at [2]; that specific to HTTP/2 are at [3].

The changes in this draft are summarized in Appendix B.

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 August 15, 2015.

Copyright Notice

Copyright (c) 2015 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.

Table of Contents

1. Introduction	4
2. HTTP/2 Protocol Overview	5
2.1. Document Organization	6
2.2. Conventions and Terminology	6
3. Starting HTTP/2	7
3.1. HTTP/2 Version Identification	8
3.2. Starting HTTP/2 for "http" URIs	9
3.2.1. HTTP2-Settings Header Field	10
3.3. Starting HTTP/2 for "https" URIs	11
3.4. Starting HTTP/2 with Prior Knowledge	11
3.5. HTTP/2 Connection Preface	11
4. HTTP Frames	12
4.1. Frame Format	13
4.2. Frame Size	14
4.3. Header Compression and Decompression	14
5. Streams and Multiplexing	15
5.1. Stream States	16
5.1.1. Stream Identifiers	21
5.1.2. Stream Concurrency	22
5.2. Flow Control	23
5.2.1. Flow Control Principles	23
5.2.2. Appropriate Use of Flow Control	24
5.3. Stream priority	25
5.3.1. Stream Dependencies	25
5.3.2. Dependency Weighting	26
5.3.3. Reprioritization	27
5.3.4. Prioritization State Management	27
5.3.5. Default Priorities	29
5.4. Error Handling	29

5.4.1.	Connection Error Handling	29
5.4.2.	Stream Error Handling	30
5.4.3.	Connection Termination	30
5.5.	Extending HTTP/2	30
6.	Frame Definitions	31
6.1.	DATA	31
6.2.	HEADERS	33
6.3.	PRIORITY	35
6.4.	RST_STREAM	36
6.5.	SETTINGS	37
6.5.1.	SETTINGS Format	38
6.5.2.	Defined SETTINGS Parameters	38
6.5.3.	Settings Synchronization	40
6.6.	PUSH_PROMISE	40
6.7.	PING	42
6.8.	GOAWAY	43
6.9.	WINDOW_UPDATE	46
6.9.1.	The Flow Control Window	47
6.9.2.	Initial Flow Control Window Size	48
6.9.3.	Reducing the Stream Window Size	49
6.10.	CONTINUATION	49
7.	Error Codes	50
8.	HTTP Message Exchanges	51
8.1.	HTTP Request/Response Exchange	51
8.1.1.	Upgrading From HTTP/2	53
8.1.2.	HTTP Header Fields	53
8.1.3.	Examples	57
8.1.4.	Request Reliability Mechanisms in HTTP/2	59
8.2.	Server Push	60
8.2.1.	Push Requests	61
8.2.2.	Push Responses	62
8.3.	The CONNECT Method	63
9.	Additional HTTP Requirements/Considerations	64
9.1.	Connection Management	64
9.1.1.	Connection Reuse	65
9.1.2.	The 421 (Misdirected Request) Status Code	66
9.2.	Use of TLS Features	66
9.2.1.	TLS 1.2 Features	67
9.2.2.	TLS 1.2 Cipher Suites	68
10.	Security Considerations	68
10.1.	Server Authority	68
10.2.	Cross-Protocol Attacks	68
10.3.	Intermediary Encapsulation Attacks	69
10.4.	Cacheability of Pushed Responses	69
10.5.	Denial of Service Considerations	70
10.5.1.	Limits on Header Block Size	71
10.5.2.	CONNECT Issues	71
10.6.	Use of Compression	72

10.7.	Use of Padding	72
10.8.	Privacy Considerations	73
11.	IANA Considerations	73
11.1.	Registration of HTTP/2 Identification Strings	74
11.2.	Frame Type Registry	74
11.3.	Settings Registry	75
11.4.	Error Code Registry	76
11.5.	HTTP2-Settings Header Field Registration	77
11.6.	PRI Method Registration	78
11.7.	The 421 (Misdirected Request) HTTP Status Code	78
12.	Acknowledgements	78
13.	References	79
13.1.	Normative References	79
13.2.	Informative References	80
13.3.	URIs	81
Appendix A.	TLS 1.2 Cipher Suite Black List	82
Appendix B.	Change Log	86
B.1.	Since draft-ietf-httpbis-http2-15	86
B.2.	Since draft-ietf-httpbis-http2-14	86
B.3.	Since draft-ietf-httpbis-http2-13	87
B.4.	Since draft-ietf-httpbis-http2-12	87
B.5.	Since draft-ietf-httpbis-http2-11	87
B.6.	Since draft-ietf-httpbis-http2-10	87
B.7.	Since draft-ietf-httpbis-http2-09	88
B.8.	Since draft-ietf-httpbis-http2-08	88
B.9.	Since draft-ietf-httpbis-http2-07	89
B.10.	Since draft-ietf-httpbis-http2-06	89
B.11.	Since draft-ietf-httpbis-http2-05	89
B.12.	Since draft-ietf-httpbis-http2-04	89
B.13.	Since draft-ietf-httpbis-http2-03	90
B.14.	Since draft-ietf-httpbis-http2-02	90
B.15.	Since draft-ietf-httpbis-http2-01	90
B.16.	Since draft-ietf-httpbis-http2-00	91
B.17.	Since draft-mbelshe-httpbis-spdyl-00	91

1. Introduction

The Hypertext Transfer Protocol (HTTP) is a wildly successful protocol. However, how HTTP/1.1 uses the underlying transport ([RFC7230], Section 6) has several characteristics that have a negative overall effect on application performance today.

In particular, HTTP/1.0 allowed only one request to be outstanding at a time on a given TCP connection. HTTP/1.1 added request pipelining, but this only partially addressed request concurrency and still suffers from head-of-line blocking. Therefore, HTTP/1.0 and HTTP/1.1 clients that need to make many requests use multiple connections to a server in order to achieve concurrency and thereby reduce latency.

Furthermore, HTTP header fields are often repetitive and verbose, causing unnecessary network traffic, as well as causing the initial TCP [TCP] congestion window to quickly fill. This can result in excessive latency when multiple requests are made on a new TCP connection.

HTTP/2 addresses these issues by defining an optimized mapping of HTTP's semantics to an underlying connection. Specifically, it allows interleaving of request and response messages on the same connection and uses an efficient coding for HTTP header fields. It also allows prioritization of requests, letting more important requests complete more quickly, further improving performance.

The resulting protocol is more friendly to the network, because fewer TCP connections can be used in comparison to HTTP/1.x. This means less competition with other flows, and longer-lived connections, which in turn leads to better utilization of available network capacity.

Finally, HTTP/2 also enables more efficient processing of messages through use of binary message framing.

2. HTTP/2 Protocol Overview

HTTP/2 provides an optimized transport for HTTP semantics. HTTP/2 supports all of the core features of HTTP/1.1, but aims to be more efficient in several ways.

The basic protocol unit in HTTP/2 is a frame (Section 4.1). Each frame type serves a different purpose. For example, HEADERS and DATA frames form the basis of HTTP requests and responses (Section 8.1); other frame types like SETTINGS, WINDOW_UPDATE, and PUSH_PROMISE are used in support of other HTTP/2 features.

Multiplexing of requests is achieved by having each HTTP request-response exchange associated with its own stream (Section 5). Streams are largely independent of each other, so a blocked or stalled request or response does not prevent progress on other streams.

Flow control and prioritization ensure that it is possible to efficiently use multiplexed streams. Flow control (Section 5.2) helps to ensure that only data that can be used by a receiver is transmitted. Prioritization (Section 5.3) ensures that limited resources can be directed to the most important streams first.

HTTP/2 adds a new interaction mode, whereby a server can push responses to a client (Section 8.2). Server push allows a server to

speculatively send data to a client that the server anticipates the client will need, trading off some network usage against a potential latency gain. The server does this by synthesizing a request, which it sends as a PUSH_PROMISE frame. The server is then able to send a response to the synthetic request on a separate stream.

Because HTTP header fields used in a connection can contain large amounts of redundant data, frames that contain them are compressed (Section 4.3). This has especially advantageous impact upon request sizes in the common case, allowing many requests to be compressed into one packet.

2.1. Document Organization

The HTTP/2 specification is split into four parts:

- o Starting HTTP/2 (Section 3) covers how an HTTP/2 connection is initiated.
- o The framing (Section 4) and streams (Section 5) layers describe the way HTTP/2 frames are structured and formed into multiplexed streams.
- o Frame (Section 6) and error (Section 7) definitions include details of the frame and error types used in HTTP/2.
- o HTTP mappings (Section 8) and additional requirements (Section 9) describe how HTTP semantics are expressed using frames and streams.

While some of the frame and stream layer concepts are isolated from HTTP, this specification does not define a completely generic framing layer. The framing and streams layers are tailored to the needs of the HTTP protocol and server push.

2.2. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

All numeric values are in network byte order. Values are unsigned unless otherwise indicated. Literal values are provided in decimal or hexadecimal as appropriate. Hexadecimal literals are prefixed with "0x" to distinguish them from decimal literals.

The following terms are used:

client: The endpoint that initiates an HTTP/2 connection. Clients send HTTP requests and receive HTTP responses.

connection: A transport-layer connection between two endpoints.

connection error: An error that affects the entire HTTP/2 connection.

endpoint: Either the client or server of the connection.

frame: The smallest unit of communication within an HTTP/2 connection, consisting of a header and a variable-length sequence of octets structured according to the frame type.

peer: An endpoint. When discussing a particular endpoint, "peer" refers to the endpoint that is remote to the primary subject of discussion.

receiver: An endpoint that is receiving frames.

sender: An endpoint that is transmitting frames.

server: The endpoint that accepts an HTTP/2 connection. Servers receive HTTP requests and serve HTTP responses.

stream: A bi-directional flow of frames within the HTTP/2 connection.

stream error: An error on the individual HTTP/2 stream.

Finally, the terms "gateway", "intermediary", "proxy", and "tunnel" are defined in Section 2.3 of [RFC7230]. Intermediaries act as both client and server at different times.

The term "payload body" is defined in Section 3.3 of [RFC7230].

3. Starting HTTP/2

An HTTP/2 connection is an application layer protocol running on top of a TCP connection ([TCP]). The client is the TCP connection initiator.

HTTP/2 uses the same "http" and "https" URI schemes used by HTTP/1.1. HTTP/2 shares the same default port numbers: 80 for "http" URIs and 443 for "https" URIs. As a result, implementations processing requests for target resource URIs like "http://example.org/foo" or "https://example.com/bar" are required to first discover whether the

upstream server (the immediate peer to which the client wishes to establish a connection) supports HTTP/2.

The means by which support for HTTP/2 is determined is different for "http" and "https" URIs. Discovery for "http" URIs is described in Section 3.2. Discovery for "https" URIs is described in Section 3.3.

3.1. HTTP/2 Version Identification

The protocol defined in this document has two identifiers.

- o The string "h2" identifies the protocol where HTTP/2 uses TLS [TLS12]. This identifier is used in the TLS application layer protocol negotiation extension (ALPN) [TLS-ALPN] field and in any place where HTTP/2 over TLS is identified.

The "h2" string is serialized into an ALPN protocol identifier as the two octet sequence: 0x68, 0x32.

- o The string "h2c" identifies the protocol where HTTP/2 is run over cleartext TCP. This identifier is used in the HTTP/1.1 Upgrade header field and in any place where HTTP/2 over TCP is identified.

The "h2c" string is reserved from the ALPN identifier space, but describes a protocol that does not use TLS.

Negotiating "h2" or "h2c" implies the use of the transport, security, framing and message semantics described in this document.

[[CREF1: RFC Editor's Note: please remove the remainder of this section prior to the publication of a final version of this document.]]

Only implementations of the final, published RFC can identify themselves as "h2" or "h2c". Until such an RFC exists, implementations MUST NOT identify themselves using these strings.

Implementations of draft versions of the protocol MUST add the string "-" and the corresponding draft number to the identifier. For example, draft-ietf-httpbis-http2-11 over TLS is identified using the string "h2-11".

Non-compatible experiments that are based on these draft versions MUST append the string "-" and an experiment name to the identifier. For example, an experimental implementation of packet mood-based encoding based on draft-ietf-httpbis-http2-09 might identify itself as "h2-09-emo". Note that any label MUST conform to the "token" syntax defined in Section 3.2.6 of [RFC7230]. Experimenters are

encouraged to coordinate their experiments on the ietf-http-wg@w3.org mailing list.

3.2. Starting HTTP/2 for "http" URIs

A client that makes a request for an "http" URI without prior knowledge about support for HTTP/2 on the next hop uses the HTTP Upgrade mechanism (Section 6.7 of [RFC7230]). The client does so by making an HTTP/1.1 request that includes an Upgrade header field with the "h2c" token. Such an HTTP/1.1 request MUST include exactly one HTTP2-Settings (Section 3.2.1) header field.

For example:

```
GET / HTTP/1.1
Host: server.example.com
Connection: Upgrade, HTTP2-Settings
Upgrade: h2c
HTTP2-Settings: <base64url encoding of HTTP/2 SETTINGS payload>
```

Requests that contain an payload body MUST be sent in their entirety before the client can send HTTP/2 frames. This means that a large request can block the use of the connection until it is completely sent.

If concurrency of an initial request with subsequent requests is important, an OPTIONS request can be used to perform the upgrade to HTTP/2, at the cost of an additional round-trip.

A server that does not support HTTP/2 can respond to the request as though the Upgrade header field were absent:

```
HTTP/1.1 200 OK
Content-Length: 243
Content-Type: text/html
```

...

A server MUST ignore an "h2" token in an Upgrade header field. Presence of a token with "h2" implies HTTP/2 over TLS, which is instead negotiated as described in Section 3.3.

A server that supports HTTP/2 accepts the upgrade with a 101 (Switching Protocols) response. After the empty line that terminates the 101 response, the server can begin sending HTTP/2 frames. These frames MUST include a response to the request that initiated the Upgrade.

For example:

```
HTTP/1.1 101 Switching Protocols
Connection: Upgrade
Upgrade: h2c
```

```
[ HTTP/2 connection ...
```

The first HTTP/2 frame sent by the server MUST be a SETTINGS frame (Section 6.5) as the server connection preface (Section 3.5). Upon receiving the 101 response, the client MUST send a connection preface (Section 3.5), which includes a SETTINGS frame.

The HTTP/1.1 request that is sent prior to upgrade is assigned a stream identifier of 1 (see Section 5.1.1) with default priority values (Section 5.3.5). Stream 1 is implicitly "half closed" from the client toward the server (see Section 5.1), since the request is completed as an HTTP/1.1 request. After commencing the HTTP/2 connection, stream 1 is used for the response.

3.2.1. HTTP2-Settings Header Field

A request that upgrades from HTTP/1.1 to HTTP/2 MUST include exactly one "HTTP2-Settings" header field. The "HTTP2-Settings" header field is a connection-specific header field that includes parameters that govern the HTTP/2 connection, provided in anticipation of the server accepting the request to upgrade.

```
HTTP2-Settings    = token68
```

A server MUST NOT upgrade the connection to HTTP/2 if this header field is not present, or if more than one is present. A server MUST NOT send this header field.

The content of the "HTTP2-Settings" header field is the payload of a SETTINGS frame (Section 6.5), encoded as a base64url string (that is, the URL- and filename-safe Base64 encoding described in Section 5 of [RFC4648], with any trailing '=' characters omitted). The ABNF [RFC5234] production for "token68" is defined in Section 2.1 of [RFC7235].

Since the upgrade is only intended to apply to the immediate connection, a client sending "HTTP2-Settings" MUST also send "HTTP2-Settings" as a connection option in the "Connection" header field to prevent it from being forwarded (see Section 6.1 of [RFC7230]).

A server decodes and interprets these values as it would any other SETTINGS frame. Explicit acknowledgement of these settings (Section 6.5.3) is not necessary, since a 101 response serves as implicit acknowledgment. Providing these values in the Upgrade request gives a client an opportunity to provide parameters prior to receiving any frames from the server.

3.3. Starting HTTP/2 for "https" URIs

A client that makes a request to an "https" URI uses TLS [TLS12] with the application layer protocol negotiation (ALPN) extension [TLS-ALPN].

HTTP/2 over TLS uses the "h2" protocol identifier. The "h2c" protocol identifier MUST NOT be sent by a client or selected by a server; the "h2c" protocol identifier describes a protocol that does not use TLS.

Once TLS negotiation is complete, both the client and the server MUST send a connection preface (Section 3.5).

3.4. Starting HTTP/2 with Prior Knowledge

A client can learn that a particular server supports HTTP/2 by other means. For example, [ALT-SVC] describes a mechanism for advertising this capability.

A client MUST send the connection preface (Section 3.5), and then MAY immediately send HTTP/2 frames to such a server; servers can identify these connections by the presence of the connection preface. This only affects the establishment of HTTP/2 connections over cleartext TCP; implementations that support HTTP/2 over TLS MUST use protocol negotiation in TLS [TLS-ALPN].

Likewise, the server MUST send a connection preface (Section 3.5).

Without additional information, prior support for HTTP/2 is not a strong signal that a given server will support HTTP/2 for future connections. For example, it is possible for server configurations to change, for configurations to differ between instances in clustered servers, or for network conditions to change.

3.5. HTTP/2 Connection Preface

In HTTP/2, each endpoint is required to send a connection preface as a final confirmation of the protocol in use, and to establish the initial settings for the HTTP/2 connection. The client and server each send a different connection preface.

The client connection preface starts with a sequence of 24 octets, which in hex notation are:

```
0x505249202a20485454502f322e300d0a0d0a534d0d0a0d0a
```

(the string "PRI * HTTP/2.0\r\n\r\nSM\r\n\r\n"). This sequence MUST be followed by a SETTINGS frame (Section 6.5), which MAY be empty. The client sends the client connection preface immediately upon receipt of a 101 Switching Protocols response (indicating a successful upgrade), or as the first application data octets of a TLS connection. If starting an HTTP/2 connection with prior knowledge of server support for the protocol, the client connection preface is sent upon connection establishment.

The client connection preface is selected so that a large proportion of HTTP/1.1 or HTTP/1.0 servers and intermediaries do not attempt to process further frames. Note that this does not address the concerns raised in [TALKING].

The server connection preface consists of a potentially empty SETTINGS frame (Section 6.5) that MUST be the first frame the server sends in the HTTP/2 connection.

The SETTINGS frames received from a peer as part of the connection preface MUST be acknowledged (see Section 6.5.3) after sending the connection preface.

To avoid unnecessary latency, clients are permitted to send additional frames to the server immediately after sending the client connection preface, without waiting to receive the server connection preface. It is important to note, however, that the server connection preface SETTINGS frame might include parameters that necessarily alter how a client is expected to communicate with the server. Upon receiving the SETTINGS frame, the client is expected to honor any parameters established. In some configurations, it is possible for the server to transmit SETTINGS before the client sends additional frames, providing an opportunity to avoid this issue.

Clients and servers MUST treat an invalid connection preface as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`. A GOAWAY frame (Section 6.8) MAY be omitted in this case, since an invalid preface indicates that the peer is not using HTTP/2.

4. HTTP Frames

Once the HTTP/2 connection is established, endpoints can begin exchanging frames.

4.1. Frame Format

All frames begin with a fixed 9-octet header followed by a variable-length payload.




Figure 1: Frame Layout

The fields of the frame header are defined as:

Length: The length of the frame payload expressed as an unsigned 24-bit integer. Values greater than 2^{14} (16,384) MUST NOT be sent unless the receiver has set a larger value for `SETTINGS_MAX_FRAME_SIZE`.

The 9 octets of the frame header are not included in this value.

Type: The 8-bit type of the frame. The frame type determines the format and semantics of the frame. Implementations MUST ignore and discard any frame that has a type that is unknown.

Flags: An 8-bit field reserved for frame-type specific boolean flags.

Flags are assigned semantics specific to the indicated frame type. Flags that have no defined semantics for a particular frame type MUST be ignored, and MUST be left unset (0x0) when sending.

R: A reserved 1-bit field. The semantics of this bit are undefined and the bit MUST remain unset (0x0) when sending and MUST be ignored when receiving.

Stream Identifier: A stream identifier (see Section 5.1.1) expressed as an unsigned 31-bit integer. The value 0x0 is reserved for frames that are associated with the connection as a whole as opposed to an individual stream.

The structure and content of the frame payload is dependent entirely on the frame type.

4.2. Frame Size

The size of a frame payload is limited by the maximum size that a receiver advertises in the `SETTINGS_MAX_FRAME_SIZE` setting. This setting can have any value between 2^{14} (16,384) and $2^{24}-1$ (16,777,215) octets, inclusive.

All implementations **MUST** be capable of receiving and minimally processing frames up to 2^{14} octets in length, plus the 9 octet frame header (Section 4.1). The size of the frame header is not included when describing frame sizes.

Note: Certain frame types, such as PING (Section 6.7), impose additional limits on the amount of payload data allowed.

An endpoint **MUST** send a `FRAME_SIZE_ERROR` error if a frame exceeds the size defined in `SETTINGS_MAX_FRAME_SIZE`, any limit defined for the frame type, or it is too small to contain mandatory frame data. A frame size error in a frame that could alter the state of the entire connection **MUST** be treated as a connection error (Section 5.4.1); this includes any frame carrying a header block (Section 4.3) (that is, `HEADERS`, `PUSH_PROMISE`, and `CONTINUATION`), `SETTINGS`, and any frame with a stream identifier of 0.

Endpoints are not obligated to use all available space in a frame. Responsiveness can be improved by using frames that are smaller than the permitted maximum size. Sending large frames can result in delays in sending time-sensitive frames (such as `RST_STREAM`, `WINDOW_UPDATE`, or `PRIORITY`) which if blocked by the transmission of a large frame, could affect performance.

4.3. Header Compression and Decompression

Just as in HTTP/1, a header field in HTTP/2 is a name with one or more associated values. They are used within HTTP request and response messages as well as server push operations (see Section 8.2).

Header lists are collections of zero or more header fields. When transmitted over a connection, a header list is serialized into a header block using HTTP Header Compression [`COMPRESSION`]. The serialized header block is then divided into one or more octet sequences, called header block fragments, and transmitted within the payload of `HEADERS` (Section 6.2), `PUSH_PROMISE` (Section 6.6) or `CONTINUATION` (Section 6.10) frames.

The Cookie header field [`COOKIE`] is treated specially by the HTTP mapping (see Section 8.1.2.5).

A receiving endpoint reassembles the header block by concatenating its fragments, then decompresses the block to reconstruct the header list.

A complete header block consists of either:

- o a single HEADERS or PUSH_PROMISE frame, with the END_HEADERS flag set, or
- o a HEADERS or PUSH_PROMISE frame with the END_HEADERS flag cleared and one or more CONTINUATION frames, where the last CONTINUATION frame has the END_HEADERS flag set.

Header compression is stateful. One compression context and one decompression context is used for the entire connection. A decoding error in a header block MUST be treated as a connection error (Section 5.4.1) of type `COMPRESSION_ERROR`.

Each header block is processed as a discrete unit. Header blocks MUST be transmitted as a contiguous sequence of frames, with no interleaved frames of any other type or from any other stream. The last frame in a sequence of HEADERS or CONTINUATION frames has the END_HEADERS flag set. The last frame in a sequence of PUSH_PROMISE or CONTINUATION frames has the END_HEADERS flag set. This allows a header block to be logically equivalent to a single frame.

Header block fragments can only be sent as the payload of HEADERS, PUSH_PROMISE or CONTINUATION frames, because these frames carry data that can modify the compression context maintained by a receiver. An endpoint receiving HEADERS, PUSH_PROMISE or CONTINUATION frames needs to reassemble header blocks and perform decompression even if the frames are to be discarded. A receiver MUST terminate the connection with a connection error (Section 5.4.1) of type `COMPRESSION_ERROR` if it does not decompress a header block.

5. Streams and Multiplexing


A "stream" is an independent, bi-directional sequence of frames exchanged between the client and server within an HTTP/2 connection. Streams have several important characteristics:

- o A single HTTP/2 connection can contain multiple concurrently open streams, with either endpoint interleaving frames from multiple streams.
- o Streams can be established and used unilaterally or shared by either the client or server.

- o Streams can be closed by either endpoint.
- o The order in which frames are sent on a stream is significant. Recipients process frames in the order they are received. In particular, the order of HEADERS, and DATA frames is semantically significant.
- o Streams are identified by an integer. Stream identifiers are assigned to streams by the endpoint initiating the stream.

5.1. Stream States

The lifecycle of a stream is shown in Figure 2.



send: endpoint sends this frame
 recv: endpoint receives this frame

H: HEADERS frame (with implied CONTINUATIONS)
 PP: PUSH_PROMISE frame (with implied CONTINUATIONS)
 ES: END_STREAM flag
 R: RST_STREAM frame

Figure 2: Stream States

Note that this diagram shows stream state transitions and the frames and flags that affect those transitions only. In this regard, CONTINUATION frames do not result in state transitions; they are effectively part of the HEADERS or PUSH_PROMISE that they follow. For the purpose of state transitions, the END_STREAM flag is processed as a separate event to the frame that bears it; a HEADERS frame with the END_STREAM flag set can cause two state transitions.

Both endpoints have a subjective view of the state of a stream that could be different when frames are in transit. Endpoints do not coordinate the creation of streams; they are created unilaterally by either endpoint. The negative consequences of a mismatch in states are limited to the "closed" state after sending RST_STREAM, where frames might be received for some time after closing.

Streams have the following states:

idle:

All streams start in the "idle" state.

The following transitions are valid from this state:

- * Sending or receiving a HEADERS frame causes the stream to become "open". The stream identifier is selected as described in Section 5.1.1. The same HEADERS frame can also cause a stream to immediately become "half closed".
- * Sending a PUSH_PROMISE frame on another stream reserves the idle stream that is identified for later use. The stream state for the reserved stream transitions to "reserved (local)".
- * Receiving a PUSH_PROMISE frame on another stream reserves an idle stream that is identified for later use. The stream state for the reserved stream transitions to "reserved (remote)".
- * Note that the PUSH_PROMISE frame is not sent on the idle stream, but references the newly reserved stream in the Promised Stream ID field.

Receiving any frame other than HEADERS or PRIORITY on a stream in this state MUST be treated as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

reserved (local):

A stream in the "reserved (local)" state is one that has been promised by sending a PUSH_PROMISE frame. A PUSH_PROMISE frame reserves an idle stream by associating the stream with an open stream that was initiated by the remote peer (see Section 8.2).

In this state, only the following transitions are possible:

- * The endpoint can send a HEADERS frame. This causes the stream to open in a "half closed (remote)" state.
- * Either endpoint can send a RST_STREAM frame to cause the stream to become "closed". This releases the stream reservation.

An endpoint MUST NOT send any type of frame other than HEADERS, RST_STREAM, or PRIORITY in this state.

A PRIORITY or WINDOW_UPDATE frame MAY be received in this state. Receiving any type of frame other than RST_STREAM, PRIORITY or WINDOW_UPDATE on a stream in this state MUST be treated as a connection error (Section 5.4.1) of type PROTOCOL_ERROR.

reserved (remote):

A stream in the "reserved (remote)" state has been reserved by a remote peer.

In this state, only the following transitions are possible:

- * Receiving a HEADERS frame causes the stream to transition to "half closed (local)".
- * Either endpoint can send a RST_STREAM frame to cause the stream to become "closed". This releases the stream reservation.

An endpoint MAY send a PRIORITY frame in this state to reprioritize the reserved stream. An endpoint MUST NOT send any type of frame other than RST_STREAM, WINDOW_UPDATE, or PRIORITY in this state.

Receiving any type of frame other than HEADERS, RST_STREAM or PRIORITY on a stream in this state MUST be treated as a connection error (Section 5.4.1) of type PROTOCOL_ERROR.

open:

A stream in the "open" state may be used by both peers to send frames of any type. In this state, sending peers observe advertised stream level flow control limits (Section 5.2).

From this state either endpoint can send a frame with an END_STREAM flag set, which causes the stream to transition into one of the "half closed" states: an endpoint sending an END_STREAM flag causes the stream state to become "half closed (local)"; an endpoint receiving an END_STREAM flag causes the stream state to become "half closed (remote)".

Either endpoint can send a RST_STREAM frame from this state, causing it to transition immediately to "closed".

half closed (local):

A stream that is in the "half closed (local)" state cannot be used for sending frames other than WINDOW_UPDATE, PRIORITY and RST_STREAM.

A stream transitions from this state to "closed" when a frame that contains an END_STREAM flag is received, or when either peer sends a RST_STREAM frame.

An endpoint can receive any type of frame in this state. Providing flow control credit using WINDOW_UPDATE frames is necessary to continue receiving flow controlled frames. A receiver can ignore WINDOW_UPDATE frames in this state, which might arrive for a short period after a frame bearing the END_STREAM flag is sent.

PRIORITY frames received in this state are used to reprioritize streams that depend on the identified stream.

half closed (remote):

A stream that is "half closed (remote)" is no longer being used by the peer to send frames. In this state, an endpoint is no longer obligated to maintain a receiver flow control window.

If an endpoint receives additional frames for a stream that is in this state, other than WINDOW_UPDATE, PRIORITY or RST_STREAM, it MUST respond with a stream error (Section 5.4.2) of type STREAM_CLOSED.

A stream that is "half closed (remote)" can be used by the endpoint to send frames of any type. In this state, the endpoint continues to observe advertised stream level flow control limits (Section 5.2).

A stream can transition from this state to "closed" by sending a frame that contains an END_STREAM flag, or when either peer sends a RST_STREAM frame.

closed:

The "closed" state is the terminal state.

An endpoint MUST NOT send frames other than PRIORITY on a closed stream. An endpoint that receives any frame other than PRIORITY after receiving a RST_STREAM MUST treat that as a stream error (Section 5.4.2) of type STREAM_CLOSED. Similarly, an endpoint that receives any frames after receiving a frame with the END_STREAM flag set MUST treat that as a connection error (Section 5.4.1) of type STREAM_CLOSED, unless the frame is permitted as described below.

WINDOW_UPDATE or RST_STREAM frames can be received in this state for a short period after a DATA or HEADERS frame containing an END_STREAM flag is sent. Until the remote peer receives and

processes RST_STREAM or the frame bearing the END_STREAM flag, it might send frames of these types. Endpoints MUST ignore WINDOW_UPDATE or RST_STREAM frames received in this state, though endpoints MAY choose to treat frames that arrive a significant time after sending END_STREAM as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

PRIORITY frames can be sent on closed streams to prioritize streams that are dependent on the closed stream. Endpoints SHOULD process PRIORITY frames, though they can be ignored if the stream has been removed from the dependency tree (see Section 5.3.4).

If this state is reached as a result of sending a RST_STREAM frame, the peer that receives the RST_STREAM might have already sent - or enqueued for sending - frames on the stream that cannot be withdrawn. An endpoint MUST ignore frames that it receives on closed streams after it has sent a RST_STREAM frame. An endpoint MAY choose to limit the period over which it ignores frames and treat frames that arrive after this time as being in error.

Flow controlled frames (i.e., DATA) received after sending RST_STREAM are counted toward the connection flow control window. Even though these frames might be ignored, because they are sent before the sender receives the RST_STREAM, the sender will consider the frames to count against the flow control window.

An endpoint might receive a PUSH_PROMISE frame after it sends RST_STREAM. PUSH_PROMISE causes a stream to become "reserved" even if the associated stream has been reset. Therefore, a RST_STREAM is needed to close an unwanted promised stream.

In the absence of more specific guidance elsewhere in this document, implementations SHOULD treat the receipt of a frame that is not expressly permitted in the description of a state as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`. Note that PRIORITY can be sent and received in any stream state. Frames of unknown types are ignored.

An example of the state transitions for an HTTP request/response exchange can be found in Section 8.1. An example of the state transitions for server push can be found in Section 8.2.1 and Section 8.2.2.

5.1.1. Stream Identifiers

Streams are identified with an unsigned 31-bit integer. Streams initiated by a client MUST use odd-numbered stream identifiers; those initiated by the server MUST use even-numbered stream identifiers. A

stream identifier of zero (0x0) is used for connection control messages; the stream identifier zero cannot be used to establish a new stream.

HTTP/1.1 requests that are upgraded to HTTP/2 (see Section 3.2) are responded to with a stream identifier of one (0x1). After the upgrade completes, stream 0x1 is "half closed (local)" to the client. Therefore, stream 0x1 cannot be selected as a new stream identifier by a client that upgrades from HTTP/1.1.

The identifier of a newly established stream MUST be numerically greater than all streams that the initiating endpoint has opened or reserved. This governs streams that are opened using a HEADERS frame and streams that are reserved using PUSH_PROMISE. An endpoint that receives an unexpected stream identifier MUST respond with a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

The first use of a new stream identifier implicitly closes all streams in the "idle" state that might have been initiated by that peer with a lower-valued stream identifier. For example, if a client sends a HEADERS frame on stream 7 without ever sending a frame on stream 5, then stream 5 transitions to the "closed" state when the first frame for stream 7 is sent or received.

Stream identifiers cannot be reused. Long-lived connections can result in an endpoint exhausting the available range of stream identifiers. A client that is unable to establish a new stream identifier can establish a new connection for new streams. A server that is unable to establish a new stream identifier can send a GOAWAY frame so that the client is forced to open a new connection for new streams.

5.1.2. Stream Concurrency

A peer can limit the number of concurrently active streams using the `SETTINGS_MAX_CONCURRENT_STREAMS` parameter (see Section 6.5.2) within a SETTINGS frame. The maximum concurrent streams setting is specific to each endpoint and applies only to the peer that receives the setting. That is, clients specify the maximum number of concurrent streams the server can initiate, and servers specify the maximum number of concurrent streams the client can initiate.

Streams that are in the "open" state, or either of the "half closed" states count toward the maximum number of streams that an endpoint is permitted to open. Streams in any of these three states count toward the limit advertised in the `SETTINGS_MAX_CONCURRENT_STREAMS` setting. Streams in either of the "reserved" states do not count toward the stream limit.

Endpoints **MUST NOT** exceed the limit set by their peer. An endpoint that receives a HEADERS frame that causes their advertised concurrent stream limit to be exceeded **MUST** treat this as a stream error (Section 5.4.2) of type `PROTOCOL_ERROR` or `REFUSED_STREAM`. The choice of error code determines whether the endpoint wishes to enable automatic retry, see Section 8.1.4) for details.

An endpoint that wishes to reduce the value of `SETTINGS_MAX_CONCURRENT_STREAMS` to a value that is below the current number of open streams can either close streams that exceed the new value or allow streams to complete.

5.2. Flow Control

Using streams for multiplexing introduces contention over use of the TCP connection, resulting in blocked streams. A flow control scheme ensures that streams on the same connection do not destructively interfere with each other. Flow control is used for both individual streams and for the connection as a whole.

HTTP/2 provides for flow control through use of the `WINDOW_UPDATE` frame (Section 6.9).

5.2.1. Flow Control Principles

HTTP/2 stream flow control aims to allow a variety of flow control algorithms to be used without requiring protocol changes. Flow control in HTTP/2 has the following characteristics:

1. Flow control is specific to a connection. Both types of flow control are between the endpoints of a single hop, and not over the entire end-to-end path.
2. Flow control is based on window update frames. Receivers advertise how many octets they are prepared to receive on a stream and for the entire connection. This is a credit-based scheme.
3. Flow control is directional with overall control provided by the receiver. A receiver **MAY** choose to set any window size that it desires for each stream and for the entire connection. A sender **MUST** respect flow control limits imposed by a receiver. Clients, servers and intermediaries all independently advertise their flow control window as a receiver and abide by the flow control limits set by their peer when sending.
4. The initial value for the flow control window is 65,535 octets for both new streams and the overall connection.

5. The frame type determines whether flow control applies to a frame. Of the frames specified in this document, only DATA frames are subject to flow control; all other frame types do not consume space in the advertised flow control window. This ensures that important control frames are not blocked by flow control.
6. Flow control cannot be disabled.
7. HTTP/2 defines only the format and semantics of the WINDOW_UPDATE frame (Section 6.9). This document does not stipulate how a receiver decides when to send this frame or the value that it sends, nor does it specify how a sender chooses to send packets. Implementations are able to select any algorithm that suits their needs.

Implementations are also responsible for managing how requests and responses are sent based on priority; choosing how to avoid head of line blocking for requests; and managing the creation of new streams. Algorithm choices for these could interact with any flow control algorithm.

5.2.2. Appropriate Use of Flow Control

Flow control is defined to protect endpoints that are operating under resource constraints. For example, a proxy needs to share memory between many connections, and also might have a slow upstream connection and a fast downstream one. Flow control addresses cases where the receiver is unable to process data on one stream, yet wants to continue to process other streams in the same connection.

Deployments that do not require this capability can advertise a flow control window of the maximum size ($2^{31}-1$), and by maintaining this window by sending a WINDOW_UPDATE frame when any data is received. This effectively disables flow control for that receiver. Conversely, a sender is always subject to the flow control window advertised by the receiver.

Deployments with constrained resources (for example, memory) can employ flow control to limit the amount of memory a peer can consume. Note, however, that this can lead to suboptimal use of available network resources if flow control is enabled without knowledge of the bandwidth-delay product (see [RFC7323]).

Even with full awareness of the current bandwidth-delay product, implementation of flow control can be difficult. When using flow control, the receiver **MUST** read from the TCP receive buffer in a

timely fashion. Failure to do so could lead to a deadlock when critical frames, such as WINDOW_UPDATE, are not read and acted upon.

5.3. Stream priority

A client can assign a priority for a new stream by including prioritization information in the HEADERS frame (Section 6.2) that opens the stream. At any other time, the PRIORITY frame (Section 6.3) can be used to change the priority of a stream.

The purpose of prioritization is to allow an endpoint to express how it would prefer its peer allocate resources when managing concurrent streams. Most importantly, priority can be used to select streams for transmitting frames when there is limited capacity for sending.

Streams can be prioritized by marking them as dependent on the completion of other streams (Section 5.3.1). Each dependency is assigned a relative weight, a number that is used to determine the relative proportion of available resources that are assigned to streams dependent on the same stream.

Explicitly setting the priority for a stream is input to a prioritization process. It does not guarantee any particular processing or transmission order for the stream relative to any other stream. An endpoint cannot force a peer to process concurrent streams in a particular order using priority. Expressing priority is therefore only ever a suggestion.

Prioritization information can be omitted from messages. Defaults are used prior to any explicit values being provided (Section 5.3.5).

5.3.1. Stream Dependencies

Each stream can be given an explicit dependency on another stream. Including a dependency expresses a preference to allocate resources to the identified stream rather than to the dependent stream.

A stream that is not dependent on any other stream is given a stream dependency of 0x0. In other words, the non-existent stream 0 forms the root of the tree.

A stream that depends on another stream is a dependent stream. The stream upon which a stream is dependent is a parent stream. A dependency on a stream that is not currently in the tree - such as a stream in the "idle" state - results in that stream being given a default priority (Section 5.3.5).

When assigning a dependency on another stream, the stream is added as a new dependency of the parent stream. Dependent streams that share the same parent are not ordered with respect to each other. For example, if streams B and C are dependent on stream A, and if stream D is created with a dependency on stream A, this results in a dependency order of A followed by B, C, and D in any order.




Figure 3: Example of Default Dependency Creation

An exclusive flag allows for the insertion of a new level of dependencies. The exclusive flag causes the stream to become the sole dependency of its parent stream, causing other dependencies to become dependent on the exclusive stream. In the previous example, if stream D is created with an exclusive dependency on stream A, this results in D becoming the dependency parent of B and C.



Figure 4: Example of Exclusive Dependency Creation

Inside the dependency tree, a dependent stream **SHOULD** only be allocated resources if all of the streams that it depends on (the chain of parent streams up to 0x0) are either closed, or it is not possible to make progress on them.

A stream cannot depend on itself. An endpoint **MUST** treat this as a stream error (Section 5.4.2) of type `PROTOCOL_ERROR`.

5.3.2. Dependency Weighting

All dependent streams are allocated an integer weight between 1 and 256 (inclusive).

Streams with the same parent **SHOULD** be allocated resources proportionally based on their weight. Thus, if stream B depends on stream A with weight 4, and C depends on stream A with weight 12, and if no progress can be made on A, stream B ideally receives one third of the resources allocated to stream C.

5.3.3. Reprioritization

Stream priorities are changed using the PRIORITY frame. Setting a dependency causes a stream to become dependent on the identified parent stream.

Dependent streams move with their parent stream if the parent is reprioritized. Setting a dependency with the exclusive flag for a reprioritized stream moves all the dependencies of the new parent stream to become dependent on the reprioritized stream.

If a stream is made dependent on one of its own dependencies, the formerly dependent stream is first moved to be dependent on the reprioritized stream's previous parent. The moved dependency retains its weight.

For example, consider an original dependency tree where B and C depend on A, D and E depend on C, and F depends on D. If A is made dependent on D, then D takes the place of A. All other dependency relationships stay the same, except for F, which becomes dependent on A if the reprioritization is exclusive.




Figure 5: Example of Dependency Reordering

5.3.4. Prioritization State Management

When a stream is removed from the dependency tree, its dependencies can be moved to become dependent on the parent of the closed stream. The weights of new dependencies are recalculated by distributing the weight of the dependency of the closed stream proportionally based on the weights of its dependencies.

Streams that are removed from the dependency tree cause some prioritization information to be lost. Resources are shared between streams with the same parent stream, which means that if a stream in that set closes or becomes blocked, any spare capacity allocated to a stream is distributed to the immediate neighbors of the stream.

However, if the common dependency is removed from the tree, those streams share resources with streams at the next highest level.

For example, assume streams A and B share a parent, and streams C and D both depend on stream A. Prior to the removal of stream A, if streams A and D are unable to proceed, then stream C receives all the resources dedicated to stream A. If stream A is removed from the tree, the weight of stream A is divided between streams C and D. If stream D is still unable to proceed, this results in stream C receiving a reduced proportion of resources. For equal starting weights, C receives one third, rather than one half, of available resources.

It is possible for a stream to become closed while prioritization information that creates a dependency on that stream is in transit. If a stream identified in a dependency has no associated priority information, then the dependent stream is instead assigned a default priority (Section 5.3.5). This potentially creates suboptimal prioritization, since the stream could be given a priority that is different to what is intended.

To avoid these problems, an endpoint SHOULD retain stream prioritization state for a period after streams become closed. The longer state is retained, the lower the chance that streams are assigned incorrect or default priority values.

Similarly, streams that are in the "idle" state can be assigned priority or become a parent of other streams. This allows for the creation of a grouping node in the dependency tree, which enables more flexible expressions of priority. Idle streams begin with a default priority (Section 5.3.5).

The retention of priority information for streams that are not counted toward the limit set by `SETTINGS_MAX_CONCURRENT_STREAMS` could create a large state burden for an endpoint. Therefore the amount of prioritization state that is retained MAY be limited.

The amount of additional state an endpoint maintains for prioritization could be dependent on load; under high load, prioritization state can be discarded to limit resource commitments. In extreme cases, an endpoint could even discard prioritization state for active or reserved streams. If a limit is applied, endpoints SHOULD maintain state for at least as many streams as allowed by their setting for `SETTINGS_MAX_CONCURRENT_STREAMS`. Implementations SHOULD also attempt to retain state for streams that are in active use in the priority tree.

An endpoint receiving a PRIORITY frame that changes the priority of a closed stream SHOULD alter the dependencies of the streams that depend on it, if it has retained enough state to do so.

5.3.5. Default Priorities

All streams are initially assigned a non-exclusive dependency on stream 0x0. Pushed streams (Section 8.2) initially depend on their associated stream. In both cases, streams are assigned a default weight of 16.

5.4. Error Handling

HTTP/2 framing permits two classes of error:

- o An error condition that renders the entire connection unusable is a connection error.
- o An error in an individual stream is a stream error.

A list of error codes is included in Section 7.

5.4.1. Connection Error Handling

A connection error is any error which prevents further processing of the framing layer, or which corrupts any connection state.

An endpoint that encounters a connection error SHOULD first send a GOAWAY frame (Section 6.8) with the stream identifier of the last stream that it successfully received from its peer. The GOAWAY frame includes an error code that indicates why the connection is terminating. After sending the GOAWAY frame for an error condition, the endpoint MUST close the TCP connection.

It is possible that the GOAWAY will not be reliably received by the receiving endpoint (see [RFC7230], Section 6.6). In the event of a connection error, GOAWAY only provides a best effort attempt to communicate with the peer about why the connection is being terminated.

An endpoint can end a connection at any time. In particular, an endpoint MAY choose to treat a stream error as a connection error. Endpoints SHOULD send a GOAWAY frame when ending a connection, providing that circumstances permit it.

5.4.2. Stream Error Handling

A stream error is an error related to a specific stream that does not affect processing of other streams.

An endpoint that detects a stream error sends a RST_STREAM frame (Section 6.4) that contains the stream identifier of the stream where the error occurred. The RST_STREAM frame includes an error code that indicates the type of error.

A RST_STREAM is the last frame that an endpoint can send on a stream. The peer that sends the RST_STREAM frame MUST be prepared to receive any frames that were sent or enqueued for sending by the remote peer. These frames can be ignored, except where they modify connection state (such as the state maintained for header compression (Section 4.3), or flow control).

Normally, an endpoint SHOULD NOT send more than one RST_STREAM frame for any stream. However, an endpoint MAY send additional RST_STREAM frames if it receives frames on a closed stream after more than a round-trip time. This behavior is permitted to deal with misbehaving implementations.

An endpoint MUST NOT send a RST_STREAM in response to a RST_STREAM frame, to avoid looping.

5.4.3. Connection Termination

If the TCP connection is closed or reset while streams remain in open or half closed states, then the affected streams cannot be automatically retried (see Section 8.1.4 for details).

5.5. Extending HTTP/2

HTTP/2 permits extension of the protocol. Protocol extensions can be used to provide additional services or alter any aspect of the protocol, within the limitations described in this section. Extensions are effective only within the scope of a single HTTP/2 connection.

This applies to the protocol elements defined in this document. This does not affect the existing options for extending HTTP, such as defining new methods, status codes, or header fields.

Extensions are permitted to use new frame types (Section 4.1), new settings (Section 6.5.2), or new error codes (Section 7). Registries are established for managing these extension points: frame types

(Section 11.2), settings (Section 11.3) and error codes (Section 11.4).

Implementations MUST ignore unknown or unsupported values in all extensible protocol elements. Implementations MUST discard frames that have unknown or unsupported types. This means that any of these extension points can be safely used by extensions without prior arrangement or negotiation. However, extension frames that appear in the middle of a header block (Section 4.3) are not permitted; these MUST be treated as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

Extensions that could change the semantics of existing protocol components MUST be negotiated before being used. For example, an extension that changes the layout of the HEADERS frame cannot be used until the peer has given a positive signal that this is acceptable. In this case, it could also be necessary to coordinate when the revised layout comes into effect. Note that treating any frame other than DATA frames as flow controlled is such a change in semantics, and can only be done through negotiation.

This document doesn't mandate a specific method for negotiating the use of an extension, but notes that a setting (Section 6.5.2) could be used for that purpose. If both peers set a value that indicates willingness to use the extension, then the extension can be used. If a setting is used for extension negotiation, the initial value MUST be defined in such a fashion that the extension is initially disabled.

6. Frame Definitions

This specification defines a number of frame types, each identified by a unique 8-bit type code. Each frame type serves a distinct purpose either in the establishment and management of the connection as a whole, or of individual streams.

The transmission of specific frame types can alter the state of a connection. If endpoints fail to maintain a synchronized view of the connection state, successful communication within the connection will no longer be possible. Therefore, it is important that endpoints have a shared comprehension of how the state is affected by the use any given frame.

6.1. DATA

DATA frames (type=0x0) convey arbitrary, variable-length sequences of octets associated with a stream. One or more DATA frames are used, for instance, to carry HTTP request or response payloads.

DATA frames MAY also contain padding. Padding can be added to DATA frames to obscure the size of messages. Padding is a security feature; see Section 10.7.




Figure 6: DATA Frame Payload

The DATA frame contains the following fields:

Pad Length: An 8-bit field containing the length of the frame padding in units of octets. This field is conditional and is only present if the PADDED flag is set.

Data: Application data. The amount of data is the remainder of the frame payload after subtracting the length of the other fields that are present.

Padding: Padding octets that contain no application semantic value. Padding octets MUST be set to zero when sending. A receiver is not obligated to verify padding, but MAY treat non-zero padding as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

The DATA frame defines the following flags:

END_STREAM (0x1): Bit 0 being set indicates that this frame is the last that the endpoint will send for the identified stream. Setting this flag causes the stream to enter one of the "half closed" states or the "closed" state (Section 5.1).

PADDED (0x8): Bit 3 being set indicates that the Pad Length field and any padding that it describes is present.

DATA frames MUST be associated with a stream. If a DATA frame is received whose stream identifier field is 0x0, the recipient MUST respond with a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

DATA frames are subject to flow control and can only be sent when a stream is in the "open" or "half closed (remote)" states. The entire DATA frame payload is included in flow control, including Pad Length and Padding fields if present. If a DATA frame is received whose

stream is not in "open" or "half closed (local)" state, the recipient MUST respond with a stream error (Section 5.4.2) of type `STREAM_CLOSED`.

The total number of padding octets is determined by the value of the Pad Length field. If the length of the padding is the length of the frame payload or greater, the recipient MUST treat this as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

Note: A frame can be increased in size by one octet by including a Pad Length field with a value of zero.

6.2. HEADERS

The HEADERS frame (type=0x1) is used to open a stream (Section 5.1), and additionally carries a header block fragment. HEADERS frames can be sent on a stream in the "open" or "half closed (remote)" states.

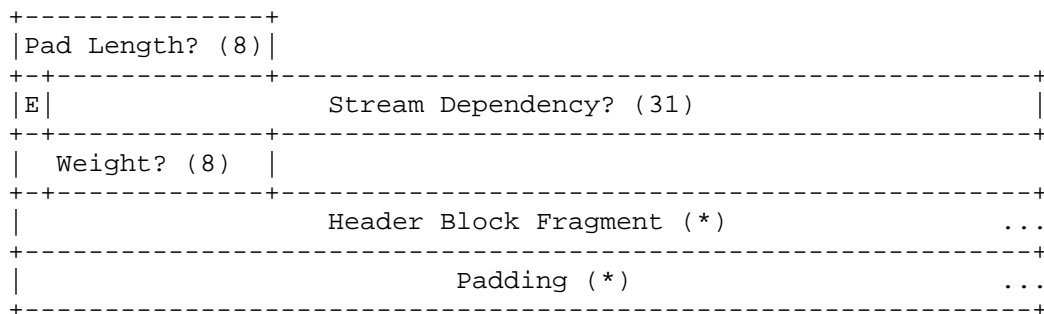


Figure 7: HEADERS Frame Payload

The HEADERS frame payload has the following fields:

Pad Length: An 8-bit field containing the length of the frame padding in units of octets. This field is only present if the `PADDED` flag is set.

E: A single bit flag indicates that the stream dependency is exclusive, see Section 5.3. This field is only present if the `PRIORITY` flag is set.

Stream Dependency: A 31-bit stream identifier for the stream that this stream depends on, see Section 5.3. This field is only present if the `PRIORITY` flag is set.

Weight: An unsigned 8-bit integer representing a priority weight for the stream, see Section 5.3. Add one to the value to obtain a

weight between 1 and 256. This field is only present if the `PRIORITY` flag is set.

Header Block Fragment: A header block fragment (Section 4.3).

Padding: Padding octets.

The HEADERS frame defines the following flags:

`END_STREAM` (0x1): Bit 0 being set indicates that the header block (Section 4.3) is the last that the endpoint will send for the identified stream.

A HEADERS frame carries the `END_STREAM` flag that signals the end of a stream. However, a HEADERS frame with the `END_STREAM` flag set can be followed by `CONTINUATION` frames on the same stream. Logically, the `CONTINUATION` frames are part of the HEADERS frame.

`END_HEADERS` (0x4): Bit 2 being set indicates that this frame contains an entire header block (Section 4.3) and is not followed by any `CONTINUATION` frames.

A HEADERS frame without the `END_HEADERS` flag set **MUST** be followed by a `CONTINUATION` frame for the same stream. A receiver **MUST** treat the receipt of any other type of frame or a frame on a different stream as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

`PADDED` (0x8): Bit 3 being set indicates that the Pad Length field and any padding that it describes is present.

`PRIORITY` (0x20): Bit 5 being set indicates that the Exclusive Flag (E), Stream Dependency, and Weight fields are present; see Section 5.3.

The payload of a HEADERS frame contains a header block fragment (Section 4.3). A header block that does not fit within a HEADERS frame is continued in a `CONTINUATION` frame (Section 6.10).

HEADERS frames **MUST** be associated with a stream. If a HEADERS frame is received whose stream identifier field is 0x0, the recipient **MUST** respond with a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

The HEADERS frame changes the connection state as described in Section 4.3.

The HEADERS frame can include padding. Padding fields and flags are identical to those defined for DATA frames (Section 6.1).

Prioritization information in a HEADERS frame is logically equivalent to a separate PRIORITY frame, but inclusion in HEADERS avoids the potential for churn in stream prioritization when new streams are created. Prioritization fields in HEADERS frames subsequent to the first on a stream reprioritize the stream (Section 5.3.3).

6.3. PRIORITY

The PRIORITY frame (type=0x2) specifies the sender-advised priority of a stream (Section 5.3). It can be sent at any time for any stream, including idle or closed streams.




Figure 8: PRIORITY Frame Payload

The payload of a PRIORITY frame contains the following fields:

E: A single bit flag indicates that the stream dependency is exclusive, see Section 5.3.

Stream Dependency: A 31-bit stream identifier for the stream that this stream depends on, see Section 5.3.

Weight: An unsigned 8-bit integer representing a priority weight for the stream, see Section 5.3. Add one to the value to obtain a weight between 1 and 256.

The PRIORITY frame does not define any flags.

The PRIORITY frame always identifies a stream. If a PRIORITY frame is received with a stream identifier of 0x0, the recipient **MUST** respond with a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

The PRIORITY frame can be sent on a stream in any state, though it cannot be sent between consecutive frames that comprise a single header block (Section 4.3). Note that this frame could arrive after processing or frame sending has completed, which would cause it to have no effect on the identified stream. For a stream that is in the "half closed (remote)" or "closed" - state, this frame can only

affect processing of the identified stream and its dependent streams and not frame transmission on that stream.

The PRIORITY frame can be sent for a stream in the "idle" or "closed" states. This allows for the reprioritization of a group of dependent streams by altering the priority of an unused or closed parent stream.

A PRIORITY frame with a length other than 5 octets MUST be treated as a stream error (Section 5.4.2) of type FRAME_SIZE_ERROR.

6.4. RST_STREAM

The RST_STREAM frame (type=0x3) allows for immediate termination of a stream. RST_STREAM is sent to request cancellation of a stream, or to indicate that an error condition has occurred.




Figure 9: RST_STREAM Frame Payload

The RST_STREAM frame contains a single unsigned, 32-bit integer identifying the error code (Section 7). The error code indicates why the stream is being terminated.

The RST_STREAM frame does not define any flags.

The RST_STREAM frame fully terminates the referenced stream and causes it to enter the closed state. After receiving a RST_STREAM on a stream, the receiver MUST NOT send additional frames for that stream, with the exception of PRIORITY. However, after sending the RST_STREAM, the sending endpoint MUST be prepared to receive and process additional frames sent on the stream that might have been sent by the peer prior to the arrival of the RST_STREAM.

RST_STREAM frames MUST be associated with a stream. If a RST_STREAM frame is received with a stream identifier of 0x0, the recipient MUST treat this as a connection error (Section 5.4.1) of type PROTOCOL_ERROR.

RST_STREAM frames MUST NOT be sent for a stream in the "idle" state. If a RST_STREAM frame identifying an idle stream is received, the recipient MUST treat this as a connection error (Section 5.4.1) of type PROTOCOL_ERROR.

A RST_STREAM frame with a length other than 4 octets MUST be treated as a connection error (Section 5.4.1) of type FRAME_SIZE_ERROR.

6.5. SETTINGS

The SETTINGS frame (type=0x4) conveys configuration parameters that affect how endpoints communicate, such as preferences and constraints on peer behavior. The SETTINGS frame is also used to acknowledge the receipt of those parameters. Individually, a SETTINGS parameter can also be referred to as a "setting".

SETTINGS parameters are not negotiated; they describe characteristics of the sending peer, which are used by the receiving peer. Different values for the same parameter can be advertised by each peer. For example, a client might set a high initial flow control window, whereas a server might set a lower value to conserve resources.

A SETTINGS frame MUST be sent by both endpoints at the start of a connection, and MAY be sent at any other time by either endpoint over the lifetime of the connection. Implementations MUST support all of the parameters defined by this specification.

Each parameter in a SETTINGS frame replaces any existing value for that parameter. Parameters are processed in the order in which they appear, and a receiver of a SETTINGS frame does not need to maintain any state other than the current value of its parameters. Therefore, the value of a SETTINGS parameter is the last value that is seen by a receiver.

SETTINGS parameters are acknowledged by the receiving peer. To enable this, the SETTINGS frame defines the following flag:

ACK (0x1): Bit 0 being set indicates that this frame acknowledges receipt and application of the peer's SETTINGS frame. When this bit is set, the payload of the SETTINGS frame MUST be empty. Receipt of a SETTINGS frame with the ACK flag set and a length field value other than 0 MUST be treated as a connection error (Section 5.4.1) of type FRAME_SIZE_ERROR. For more info, see Settings Synchronization (Section 6.5.3).

SETTINGS frames always apply to a connection, never a single stream. The stream identifier for a SETTINGS frame MUST be zero (0x0). If an endpoint receives a SETTINGS frame whose stream identifier field is anything other than 0x0, the endpoint MUST respond with a connection error (Section 5.4.1) of type PROTOCOL_ERROR.

The SETTINGS frame affects connection state. A badly formed or incomplete SETTINGS frame MUST be treated as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

A SETTINGS frame with a length other than a multiple of 6 octets MUST be treated as a connection error (Section 5.4.1) of type `FRAME_SIZE_ERROR`.

6.5.1. SETTINGS Format

The payload of a SETTINGS frame consists of zero or more parameters, each consisting of an unsigned 16-bit setting identifier and an unsigned 32-bit value.




Figure 10: Setting Format

6.5.2. Defined SETTINGS Parameters

The following parameters are defined:

SETTINGS_HEADER_TABLE_SIZE (0x1): Allows the sender to inform the remote endpoint of the maximum size of the header compression table used to decode header blocks, in octets. The encoder can select any size equal to or less than this value by using signaling specific to the header compression format inside a header block, see [COMPRESSION]. The initial value is 4,096 octets.

SETTINGS_ENABLE_PUSH (0x2): This setting can be use to disable server push (Section 8.2). An endpoint MUST NOT send a `PUSH_PROMISE` frame if it receives this parameter set to a value of 0. An endpoint that has both set this parameter to 0 and had it acknowledged MUST treat the receipt of a `PUSH_PROMISE` frame as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

The initial value is 1, which indicates that server push is permitted. Any value other than 0 or 1 MUST be treated as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

SETTINGS_MAX_CONCURRENT_STREAMS (0x3): Indicates the maximum number of concurrent streams that the sender will allow. This limit is directional: it applies to the number of streams that the sender

permits the receiver to create. Initially there is no limit to this value. It is recommended that this value be no smaller than 100, so as to not unnecessarily limit parallelism.

A value of 0 for `SETTINGS_MAX_CONCURRENT_STREAMS` SHOULD NOT be treated as special by endpoints. A zero value does prevent the creation of new streams, however this can also happen for any limit that is exhausted with active streams. Servers SHOULD only set a zero value for short durations; if a server does not wish to accept requests, closing the connection is more appropriate.

`SETTINGS_INITIAL_WINDOW_SIZE` (0x4): Indicates the sender's initial window size (in octets) for stream level flow control. The initial value is $2^{16}-1$ (65,535) octets.

This setting affects the window size of all streams, see Section 6.9.2.

Values above the maximum flow control window size of $2^{31}-1$ MUST be treated as a connection error (Section 5.4.1) of type `FLOW_CONTROL_ERROR`.

`SETTINGS_MAX_FRAME_SIZE` (0x5): Indicates the size of the largest frame payload that the sender is willing to receive, in octets.

The initial value is 2^{14} (16,384) octets. The value advertised by an endpoint MUST be between this initial value and the maximum allowed frame size ($2^{24}-1$ or 16,777,215 octets), inclusive. Values outside this range MUST be treated as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

`SETTINGS_MAX_HEADER_LIST_SIZE` (0x6): This advisory setting informs a peer of the maximum size of header list that the sender is prepared to accept, in octets. The value is based on the uncompressed size of header fields, including the length of the name and value in octets plus an overhead of 32 octets for each header field.

For any given request, a lower limit than what is advertised MAY be enforced. The initial value of this setting is unlimited.

An endpoint that receives a `SETTINGS` frame with any unknown or unsupported identifier MUST ignore that setting.

6.5.3. Settings Synchronization

Most values in SETTINGS benefit from or require an understanding of when the peer has received and applied the changed parameter values. In order to provide such synchronization timepoints, the recipient of a SETTINGS frame in which the ACK flag is not set MUST apply the updated parameters as soon as possible upon receipt.

The values in the SETTINGS frame MUST be processed in the order they appear, with no other frame processing between values. Unsupported parameters MUST be ignored. Once all values have been processed, the recipient MUST immediately emit a SETTINGS frame with the ACK flag set. Upon receiving a SETTINGS frame with the ACK flag set, the sender of the altered parameters can rely on the setting having been applied.

If the sender of a SETTINGS frame does not receive an acknowledgement within a reasonable amount of time, it MAY issue a connection error (Section 5.4.1) of type SETTINGS_TIMEOUT.

6.6. PUSH_PROMISE

The PUSH_PROMISE frame (type=0x5) is used to notify the peer endpoint in advance of streams the sender intends to initiate. The PUSH_PROMISE frame includes the unsigned 31-bit identifier of the stream the endpoint plans to create along with a set of headers that provide additional context for the stream. Section 8.2 contains a thorough description of the use of PUSH_PROMISE frames.




Figure 11: PUSH_PROMISE Payload Format

The PUSH_PROMISE frame payload has the following fields:

Pad Length: An 8-bit field containing the length of the frame padding in units of octets. This field is only present if the PADDED flag is set.

R: A single reserved bit.

Promised Stream ID: An unsigned 31-bit integer that identifies the stream that is reserved by the PUSH_PROMISE. The promised stream identifier **MUST** be a valid choice for the next stream sent by the sender (see new stream identifier (Section 5.1.1)).

Header Block Fragment: A header block fragment (Section 4.3) containing request header fields.

Padding: Padding octets.

The PUSH_PROMISE frame defines the following flags:

END_HEADERS (0x4): Bit 2 being set indicates that this frame contains an entire header block (Section 4.3) and is not followed by any CONTINUATION frames.

A PUSH_PROMISE frame without the END_HEADERS flag set **MUST** be followed by a CONTINUATION frame for the same stream. A receiver **MUST** treat the receipt of any other type of frame or a frame on a different stream as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

PADDED (0x8): Bit 3 being set indicates that the Pad Length field and any padding that it describes is present.

PUSH_PROMISE frames **MUST** be associated with a peer-initiated stream that is in either the "open" or "half closed (remote)" state. The stream identifier of a PUSH_PROMISE frame indicates the stream it is associated with. If the stream identifier field specifies the value 0x0, a recipient **MUST** respond with a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

Promised streams are not required to be used in the order they are promised. The PUSH_PROMISE only reserves stream identifiers for later use.

PUSH_PROMISE **MUST NOT** be sent if the `SETTINGS_ENABLE_PUSH` setting of the peer endpoint is set to 0. An endpoint that has set this setting and has received acknowledgement **MUST** treat the receipt of a PUSH_PROMISE frame as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

Recipients of PUSH_PROMISE frames can choose to reject promised streams by returning a `RST_STREAM` referencing the promised stream identifier back to the sender of the PUSH_PROMISE.

A PUSH_PROMISE frame modifies the connection state in two ways. The inclusion of a header block (Section 4.3) potentially modifies the

state maintained for header compression. PUSH_PROMISE also reserves a stream for later use, causing the promised stream to enter the "reserved" state. A sender MUST NOT send a PUSH_PROMISE on a stream unless that stream is either "open" or "half closed (remote)"; the sender MUST ensure that the promised stream is a valid choice for a new stream identifier (Section 5.1.1) (that is, the promised stream MUST be in the "idle" state).

Since `PUSH_PROMISE` reserves a stream, ignoring a `PUSH_PROMISE` frame causes the stream state to become indeterminate. A receiver **MUST** treat the receipt of a `PUSH_PROMISE` on a stream that is neither "open" nor "half closed (local)" as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`. However, an endpoint that has sent `RST_STREAM` on the associated stream **MUST** handle `PUSH_PROMISE` frames that might have been created before the `RST_STREAM` frame is received and processed.

A receiver MUST treat the receipt of a PUSH_PROMISE that promises an illegal stream identifier (Section 5.1.1) (that is, an identifier for a stream that is not currently in the "idle" state) as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

The `PUSH_PROMISE` frame can include padding. Padding fields and flags are identical to those defined for `DATA` frames (Section 6.1).

6.7. PING

The PING frame (type=0x6) is a mechanism for measuring a minimal round trip time from the sender, as well as determining whether an idle connection is still functional. PING frames can be sent from any endpoint.




Figure 12: PING Payload Format

In addition to the frame header, PING frames MUST contain 8 octets of data in the payload. A sender can include any value it chooses and use those octets in any fashion.

Receivers of a PING frame that does not include an ACK flag MUST send a PING frame with the ACK flag set in response, with an identical payload. PING responses SHOULD be given higher priority than any other frame.

The PING frame defines the following flags:

ACK (0x1): Bit 0 being set indicates that this PING frame is a PING response. An endpoint **MUST** set this flag in PING responses. An endpoint **MUST NOT** respond to PING frames containing this flag.

PING frames are not associated with any individual stream. If a PING frame is received with a stream identifier field value other than 0x0, the recipient **MUST** respond with a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

Receipt of a PING frame with a length field value other than 8 **MUST** be treated as a connection error (Section 5.4.1) of type `FRAME_SIZE_ERROR`.

6.8. GOAWAY

The GOAWAY frame (type=0x7) informs the remote peer to stop creating streams on this connection. GOAWAY can be sent by either the client or the server. Once sent, the sender will ignore frames sent on any new streams with identifiers higher than the included last stream identifier. Receivers of a GOAWAY frame **MUST NOT** open additional streams on the connection, although a new connection can be established for new streams.

The purpose of this frame is to allow an endpoint to gracefully stop accepting new streams, while still finishing processing of previously established streams. This enables administrative actions, like server maintenance.

There is an inherent race condition between an endpoint starting new streams and the remote sending a GOAWAY frame. To deal with this case, the GOAWAY contains the stream identifier of the last peer-initiated stream which was or might be processed on the sending endpoint in this connection. For instance, if the server sends a GOAWAY frame, the identified stream is the highest numbered stream initiated by the client.

If the receiver of the GOAWAY has sent data on streams with a higher stream identifier than what is indicated in the GOAWAY frame, those streams are not or will not be processed. The receiver of the GOAWAY frame can treat the streams as though they had never been created at all, thereby allowing those streams to be retried later on a new connection.

Endpoints **SHOULD** always send a GOAWAY frame before closing a connection so that the remote peer can know whether a stream has been partially processed or not. For example, if an HTTP client sends a

POST at the same time that a server closes a connection, the client cannot know if the server started to process that POST request if the server does not send a GOAWAY frame to indicate what streams it might have acted on.

An endpoint might choose to close a connection without sending GOAWAY for misbehaving peers.




Figure 13: GOAWAY Payload Format

The GOAWAY frame does not define any flags.

The GOAWAY frame applies to the connection, not a specific stream. An endpoint **MUST** treat a GOAWAY frame with a stream identifier other than 0x0 as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

The last stream identifier in the GOAWAY frame contains the highest numbered stream identifier for which the sender of the GOAWAY frame might have taken some action on, or might yet take action on. All streams up to and including the identified stream might have been processed in some way. The last stream identifier can be set to 0 if no streams were processed.

Note: In this context, "processed" means that some data from the stream was passed to some higher layer of software that might have taken some action as a result.

If a connection terminates without a GOAWAY frame, the last stream identifier is effectively the highest possible stream identifier.

On streams with lower or equal numbered identifiers that were not closed completely prior to the connection being closed, re-attempting requests, transactions, or any protocol activity is not possible, with the exception of idempotent actions like HTTP GET, PUT, or DELETE. Any protocol activity that uses higher numbered streams can be safely retried using a new connection.

Activity on streams numbered lower or equal to the last stream identifier might still complete successfully. The sender of a GOAWAY

frame might gracefully shut down a connection by sending a GOAWAY frame, maintaining the connection in an open state until all in-progress streams complete.

An endpoint MAY send multiple GOAWAY frames if circumstances change. For instance, an endpoint that sends GOAWAY with NO_ERROR during graceful shutdown could subsequently encounter a condition that requires immediate termination of the connection. The last stream identifier from the last GOAWAY frame received indicates which streams could have been acted upon. Endpoints MUST NOT increase the value they send in the last stream identifier, since the peers might already have retried unprocessed requests on another connection.

A client that is unable to retry requests loses all requests that are in flight when the server closes the connection. This is especially true for intermediaries that might not be serving clients using HTTP/2. A server that is attempting to gracefully shut down a connection SHOULD send an initial GOAWAY frame with the last stream identifier set to $2^{31}-1$ and a NO_ERROR code. This signals to the client that a shutdown is imminent and that no further requests can be initiated. After waiting at least one round trip time, the server can send another GOAWAY frame with an updated last stream identifier. This ensures that a connection can be cleanly shut down without losing requests.

After sending a GOAWAY frame, the sender can discard frames for streams with identifiers higher than the identified last stream. However, any frames that alter connection state cannot be completely ignored. For instance, HEADERS, PUSH_PROMISE and CONTINUATION frames MUST be minimally processed to ensure the state maintained for header compression is consistent (see Section 4.3); similarly DATA frames MUST be counted toward the connection flow control window. Failure to process these frames can cause flow control or header compression state to become unsynchronized.

The GOAWAY frame also contains a 32-bit error code (Section 7) that contains the reason for closing the connection.

Endpoints MAY append opaque data to the payload of any GOAWAY frame. Additional debug data is intended for diagnostic purposes only and carries no semantic value. Debug information could contain security- or privacy-sensitive data. Logged or otherwise persistently stored debug data MUST have adequate safeguards to prevent unauthorized access.

6.9. WINDOW_UPDATE

The WINDOW_UPDATE frame (type=0x8) is used to implement flow control; see Section 5.2 for an overview.

Flow control operates at two levels: on each individual stream and on the entire connection.

Both types of flow control are hop-by-hop; that is, only between the two endpoints. Intermediaries do not forward WINDOW_UPDATE frames between dependent connections. However, throttling of data transfer by any receiver can indirectly cause the propagation of flow control information toward the original sender.

Flow control only applies to frames that are identified as being subject to flow control. Of the frame types defined in this document, this includes only DATA frames. Frames that are exempt from flow control MUST be accepted and processed, unless the receiver is unable to assign resources to handling the frame. A receiver MAY respond with a stream error (Section 5.4.2) or connection error (Section 5.4.1) of type FLOW_CONTROL_ERROR if it is unable to accept a frame.




Figure 14: WINDOW_UPDATE Payload Format

The payload of a WINDOW_UPDATE frame is one reserved bit, plus an unsigned 31-bit integer indicating the number of octets that the sender can transmit in addition to the existing flow control window. The legal range for the increment to the flow control window is 1 to $2^{31}-1$ (2,147,483,647) octets.

The WINDOW_UPDATE frame does not define any flags.

The WINDOW_UPDATE frame can be specific to a stream or to the entire connection. In the former case, the frame's stream identifier indicates the affected stream; in the latter, the value "0" indicates that the entire connection is the subject of the frame.

A receiver MUST treat the receipt of a WINDOW_UPDATE frame with a flow control window increment of 0 as a stream error (Section 5.4.2) of type PROTOCOL_ERROR; errors on the connection flow control window MUST be treated as a connection error (Section 5.4.1).

WINDOW_UPDATE can be sent by a peer that has sent a frame bearing the END_STREAM flag. This means that a receiver could receive a WINDOW_UPDATE frame on a "half closed (remote)" or "closed" stream. A receiver MUST NOT treat this as an error, see Section 5.1.

A receiver that receives a flow controlled frame MUST always account for its contribution against the connection flow control window, unless the receiver treats this as a connection error (Section 5.4.1). This is necessary even if the frame is in error. Since the sender counts the frame toward the flow control window, if the receiver does not, the flow control window at sender and receiver can become different.

A WINDOW_UPDATE frame with a length other than 4 octets MUST be treated as a connection error (Section 5.4.1) of type FRAME_SIZE_ERROR.

6.9.1. The Flow Control Window

Flow control in HTTP/2 is implemented using a window kept by each sender on every stream. The flow control window is a simple integer value that indicates how many octets of data the sender is permitted to transmit; as such, its size is a measure of the buffering capacity of the receiver.

Two flow control windows are applicable: the stream flow control window and the connection flow control window. The sender MUST NOT send a flow controlled frame with a length that exceeds the space available in either of the flow control windows advertised by the receiver. Frames with zero length with the END_STREAM flag set (that is, an empty DATA frame) MAY be sent if there is no available space in either flow control window.

For flow control calculations, the 9 octet frame header is not counted.

After sending a flow controlled frame, the sender reduces the space available in both windows by the length of the transmitted frame.

The receiver of a frame sends a WINDOW_UPDATE frame as it consumes data and frees up space in flow control windows. Separate WINDOW_UPDATE frames are sent for the stream and connection level flow control windows.

A sender that receives a WINDOW_UPDATE frame updates the corresponding window by the amount specified in the frame.

A sender MUST NOT allow a flow control window to exceed $2^{31}-1$ octets. If a sender receives a WINDOW_UPDATE that causes a flow control window to exceed this maximum it MUST terminate either the stream or the connection, as appropriate. For streams, the sender sends a RST_STREAM with the error code of FLOW_CONTROL_ERROR code; for the connection, a GOAWAY frame with a FLOW_CONTROL_ERROR code.

Flow controlled frames from the sender and WINDOW_UPDATE frames from the receiver are completely asynchronous with respect to each other. This property allows a receiver to aggressively update the window size kept by the sender to prevent streams from stalling.

6.9.2. Initial Flow Control Window Size

When an HTTP/2 connection is first established, new streams are created with an initial flow control window size of 65,535 octets. The connection flow control window is 65,535 octets. Both endpoints can adjust the initial window size for new streams by including a value for SETTINGS_INITIAL_WINDOW_SIZE in the SETTINGS frame that forms part of the connection preface. The connection flow control window can only be changed using WINDOW_UPDATE frames.

Prior to receiving a SETTINGS frame that sets a value for SETTINGS_INITIAL_WINDOW_SIZE, an endpoint can only use the default initial window size when sending flow controlled frames. Similarly, the connection flow control window is set to the default initial window size until a WINDOW_UPDATE frame is received.

A SETTINGS frame can alter the initial flow control window size for all streams in the "open" or "half closed (remote)" state. When the value of SETTINGS_INITIAL_WINDOW_SIZE changes, a receiver MUST adjust the size of all stream flow control windows that it maintains by the difference between the new value and the old value.

A change to SETTINGS_INITIAL_WINDOW_SIZE can cause the available space in a flow control window to become negative. A sender MUST track the negative flow control window, and MUST NOT send new flow controlled frames until it receives WINDOW_UPDATE frames that cause the flow control window to become positive.

For example, if the client sends 60KB immediately on connection establishment, and the server sets the initial window size to be 16KB, the client will recalculate the available flow control window to be -44KB on receipt of the SETTINGS frame. The client retains a negative flow control window until WINDOW_UPDATE frames restore the window to being positive, after which the client can resume sending.

A SETTINGS frame cannot alter the connection flow control window.

An endpoint **MUST** treat a change to `SETTINGS_INITIAL_WINDOW_SIZE` that causes any flow control window to exceed the maximum size as a connection error (Section 5.4.1) of type `FLOW_CONTROL_ERROR`.

6.9.3. Reducing the Stream Window Size

A receiver that wishes to use a smaller flow control window than the current size can send a new `SETTINGS` frame. However, the receiver **MUST** be prepared to receive data that exceeds this window size, since the sender might send data that exceeds the lower limit prior to processing the `SETTINGS` frame.

After sending a `SETTINGS` frame that reduces the initial flow control window size, a receiver **MAY** continue to process streams that exceed flow control limits. Allowing streams to continue does not allow the receiver to immediately reduce the space it reserves for flow control windows. Progress on these streams can also stall, since `WINDOW_UPDATE` frames are needed to allow the sender to resume sending. The receiver **MAY** instead send a `RST_STREAM` with `FLOW_CONTROL_ERROR` error code for the affected streams.

6.10. CONTINUATION

The `CONTINUATION` frame (type=0x9) is used to continue a sequence of header block fragments (Section 4.3). Any number of `CONTINUATION` frames can be sent, as long as the preceding frame is on the same stream and is a `HEADERS`, `PUSH_PROMISE` or `CONTINUATION` frame without the `END_HEADERS` flag set.

```
+-----+
|               Header Block Fragment (*)               ...
+-----+
```

Figure 15: `CONTINUATION` Frame Payload

The `CONTINUATION` frame payload contains a header block fragment (Section 4.3).

The `CONTINUATION` frame defines the following flag:

`END_HEADERS` (0x4): Bit 2 being set indicates that this frame ends a header block (Section 4.3).

If the `END_HEADERS` bit is not set, this frame **MUST** be followed by another `CONTINUATION` frame. A receiver **MUST** treat the receipt of any other type of frame or a frame on a different stream as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

The CONTINUATION frame changes the connection state as defined in Section 4.3.

CONTINUATION frames MUST be associated with a stream. If a CONTINUATION frame is received whose stream identifier field is 0x0, the recipient MUST respond with a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

A CONTINUATION frame MUST be preceded by a `HEADERS`, `PUSH_PROMISE` or CONTINUATION frame without the `END_HEADERS` flag set. A recipient that observes violation of this rule MUST respond with a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

7. Error Codes

Error codes are 32-bit fields that are used in `RST_STREAM` and `GOAWAY` frames to convey the reasons for the stream or connection error.

Error codes share a common code space. Some error codes apply only to either streams or the entire connection and have no defined semantics in the other context.

The following error codes are defined:

`NO_ERROR` (0x0): The associated condition is not as a result of an error. For example, a `GOAWAY` might include this code to indicate graceful shutdown of a connection.

`PROTOCOL_ERROR` (0x1): The endpoint detected an unspecific protocol error. This error is for use when a more specific error code is not available.

`INTERNAL_ERROR` (0x2): The endpoint encountered an unexpected internal error.

`FLOW_CONTROL_ERROR` (0x3): The endpoint detected that its peer violated the flow control protocol.

`SETTINGS_TIMEOUT` (0x4): The endpoint sent a `SETTINGS` frame, but did not receive a response in a timely manner. See `Settings Synchronization` (Section 6.5.3).

`STREAM_CLOSED` (0x5): The endpoint received a frame after a stream was half closed.

`FRAME_SIZE_ERROR` (0x6): The endpoint received a frame with an invalid size.

REFUSED_STREAM (0x7): The endpoint refuses the stream prior to performing any application processing, see Section 8.1.4 for details.

CANCEL (0x8): Used by the endpoint to indicate that the stream is no longer needed.

COMPRESSION_ERROR (0x9): The endpoint is unable to maintain the header compression context for the connection.

CONNECT_ERROR (0xa): The connection established in response to a CONNECT request (Section 8.3) was reset or abnormally closed.

ENHANCE_YOUR_CALM (0xb): The endpoint detected that its peer is exhibiting a behavior that might be generating excessive load.

INADEQUATE_SECURITY (0xc): The underlying transport has properties that do not meet minimum security requirements (see Section 9.2).

HTTP_1_1_REQUIRED (0xd): The endpoint requires that HTTP/1.1 be used instead of HTTP/2.

Unknown or unsupported error codes MUST NOT trigger any special behavior. These MAY be treated by an implementation as being equivalent to INTERNAL_ERROR.

8. HTTP Message Exchanges

HTTP/2 is intended to be as compatible as possible with current uses of HTTP. This means that, from the application perspective, the features of the protocol are largely unchanged. To achieve this, all request and response semantics are preserved, although the syntax of conveying those semantics has changed.

Thus, the specification and requirements of HTTP/1.1 Semantics and Content [RFC7231], Conditional Requests [RFC7232], Range Requests [RFC7233], Caching [RFC7234] and Authentication [RFC7235] are applicable to HTTP/2. Selected portions of HTTP/1.1 Message Syntax and Routing [RFC7230], such as the HTTP and HTTPS URI schemes, are also applicable in HTTP/2, but the expression of those semantics for this protocol are defined in the sections below.

8.1. HTTP Request/Response Exchange

A client sends an HTTP request on a new stream, using a previously unused stream identifier (Section 5.1.1). A server sends an HTTP response on the same stream as the request.

An HTTP message (request or response) consists of:

1. for a response only, zero or more HEADERS frames (each followed by zero or more CONTINUATION frames) containing the message headers of informational (1xx) HTTP responses (see [RFC7230], Section 3.2 and [RFC7231], Section 6.2), and
2. one HEADERS frame (followed by zero or more CONTINUATION frames) containing the message headers (see [RFC7230], Section 3.2), and
3. zero or more DATA frames containing the payload body (see [RFC7230], Section 3.3), and
4. optionally, one HEADERS frame, followed by zero or more CONTINUATION frames containing the trailer-part, if present (see [RFC7230], Section 4.1.2).

The last frame in the sequence bears an END_STREAM flag, noting that a HEADERS frame bearing the END_STREAM flag can be followed by CONTINUATION frames that carry any remaining portions of the header block.

Other frames (from any stream) MUST NOT occur between either HEADERS frame and any CONTINUATION frames that might follow.

HTTP/2 uses DATA frames to carry message payloads. The "chunked" transfer encoding defined in Section 4.1 of [RFC7230] MUST NOT be used in HTTP/2.

Trailing header fields are carried in a header block that also terminates the stream. Such a header block is a sequence starting with a HEADERS frame, followed by zero or more CONTINUATION frames, where the HEADERS frame bears an END_STREAM flag. Header blocks after the first that do not terminate the stream are not part of an HTTP request or response.

A HEADERS frame (and associated CONTINUATION frames) can only appear at the start or end of a stream. An endpoint that receives a HEADERS frame without the END_STREAM flag set after receiving a final (non-informational) status code MUST treat the corresponding request or response as malformed (Section 8.1.2.6).

An HTTP request/response exchange fully consumes a single stream. A request starts with the HEADERS frame that puts the stream into an "open" state. The request ends with a frame bearing END_STREAM, which causes the stream to become "half closed (local)" for the client and "half closed (remote)" for the server. A response starts

with a HEADERS frame and ends with a frame bearing END_STREAM, which places the stream in the "closed" state.

An HTTP response is complete after the server sends - or the client receives - a frame with the END_STREAM flag set (including any CONTINUATION frames needed to complete a header block). A server can send a complete response prior to the client sending an entire request if the response does not depend on any portion of the request that has not been sent and received. When this is true, a server MAY request that the client abort transmission of a request without error by sending a RST_STREAM with an error code of NO_ERROR after sending a complete response (i.e., a frame with the END_STREAM flag). Clients MUST NOT discard responses as a result of receiving such a RST_STREAM, though clients can always discard responses at their discretion for other reasons.

8.1.1.1. Upgrading From HTTP/2

HTTP/2 removes support for the 101 (Switching Protocols) informational status code ([RFC7231], Section 6.2.2).

The semantics of 101 (Switching Protocols) aren't applicable to a multiplexed protocol. Alternative protocols are able to use the same mechanisms that HTTP/2 uses to negotiate their use (see Section 3).

8.1.1.2. HTTP Header Fields

HTTP header fields carry information as a series of key-value pairs. For a listing of registered HTTP headers, see the Message Header Field Registry maintained at [4].

Just as in HTTP/1.x, header field names are strings of ASCII characters that are compared in a case-insensitive fashion. However, header field names MUST be converted to lowercase prior to their encoding in HTTP/2. A request or response containing uppercase header field names MUST be treated as malformed (Section 8.1.2.6).

8.1.2.1. Pseudo-Header Fields

While HTTP/1.x used the message start-line (see [RFC7230], Section 3.1) to convey the target URI and method of the request, and the status code for the response, HTTP/2 uses special pseudo-header fields beginning with ':' character (ASCII 0x3a) for this purpose.

Pseudo-header fields are not HTTP header fields. Endpoints MUST NOT generate pseudo-header fields other than those defined in this document.

Pseudo-header fields are only valid in the context in which they are defined. Pseudo-header fields defined for requests MUST NOT appear in responses; pseudo-header fields defined for responses MUST NOT appear in requests. Pseudo-header fields MUST NOT appear in trailers. Endpoints MUST treat a request or response that contains undefined or invalid pseudo-header fields as malformed (Section 8.1.2.6).

All pseudo-header fields MUST appear in the header block before regular header fields. Any request or response that contains a pseudo-header field that appears in a header block after a regular header field MUST be treated as malformed (Section 8.1.2.6).

8.1.2.2. Connection-Specific Header Fields

HTTP/2 does not use the "Connection" header field to indicate connection-specific header fields; in this protocol, connection-specific metadata is conveyed by other means. An endpoint MUST NOT generate an HTTP/2 message containing connection-specific header fields; any message containing connection-specific header fields MUST be treated as malformed (Section 8.1.2.6).

The only exception to this is the TE header field, which MAY be present in an HTTP/2 request; when it is, it MUST NOT contain any value other than "trailers".

This means that an intermediary transforming an HTTP/1.x message to HTTP/2 will need to remove any header fields nominated by the Connection header field, along with the Connection header field itself. Such intermediaries SHOULD also remove other connection-specific header fields, such as Keep-Alive, Proxy-Connection, Transfer-Encoding and Upgrade, even if they are not nominated by Connection.

Note: HTTP/2 purposefully does not support upgrade to another protocol. The handshake methods described in Section 3 are believed sufficient to negotiate the use of alternative protocols.

8.1.2.3. Request Pseudo-Header Fields

The following pseudo-header fields are defined for HTTP/2 requests:

- o The ":method" pseudo-header field includes the HTTP method ([RFC7231], Section 4).
- o The ":scheme" pseudo-header field includes the scheme portion of the target URI ([RFC3986], Section 3.1).

":scheme" is not restricted to "http" and "https" schemes URIs. A proxy or gateway can translate requests for non-HTTP schemes, enabling the use of HTTP to interact with non-HTTP services.

- o The ":authority" pseudo-header field includes the authority portion of the target URI ([RFC3986], Section 3.2). The authority MUST NOT include the deprecated "userinfo" subcomponent for "http" or "https" schemes URIs.

To ensure that the HTTP/1.1 request line can be reproduced accurately, this pseudo-header field MUST be omitted when translating from an HTTP/1.1 request that has a request target in origin or asterisk form (see [RFC7230], Section 5.3). Clients that generate HTTP/2 requests directly SHOULD use the ":authority" pseudo-header field instead of the "Host" header field. An intermediary that converts an HTTP/2 request to HTTP/1.1 MUST create a "Host" header field if one is not present in a request by copying the value of the ":authority" pseudo-header field.

- o The ":path" pseudo-header field includes the path and query parts of the target URI (the "path-absolute" production from [RFC3986] and optionally a '?' character followed by the "query" production, see [RFC3986], Section 3.3 and [RFC3986], Section 3.4). A request in asterisk form includes the value '*' for the ":path" pseudo-header field.

This pseudo-header field MUST NOT be empty for "http" or "https" URIs; "http" or "https" URIs that do not contain a path component MUST include a value of '/'. The exception to this rule is an OPTIONS request for an "http" or "https" URI that does not include a path component; these MUST include a ":path" pseudo-header field with a value of '*' (see [RFC7230], Section 5.3.4).

All HTTP/2 requests MUST include exactly one valid value for the ":method", ":scheme", and ":path" pseudo-header fields, unless it is a CONNECT request (Section 8.3). An HTTP request that omits mandatory pseudo-header fields is malformed (Section 8.1.2.6).

HTTP/2 does not define a way to carry the version identifier that is included in the HTTP/1.1 request line.

8.1.2.4. Response Pseudo-Header Fields

For HTTP/2 responses, a single ":status" pseudo-header field is defined that carries the HTTP status code field (see [RFC7231], Section 6). This pseudo-header field MUST be included in all responses, otherwise the response is malformed (Section 8.1.2.6).

HTTP/2 does not define a way to carry the version or reason phrase that is included in an HTTP/1.1 status line.

8.1.2.5. Compressing the Cookie Header Field

The Cookie header field [COOKIE] uses a semi-colon (";") to delimit cookie-pairs (or "crumbs"). This header field doesn't follow the list construction rules in HTTP (see [RFC7230], Section 3.2.2), which prevents cookie-pairs from being separated into different name-value pairs. This can significantly reduce compression efficiency as individual cookie-pairs are updated.

To allow for better compression efficiency, the Cookie header field MAY be split into separate header fields, each with one or more cookie-pairs. If there are multiple Cookie header fields after decompression, these MUST be concatenated into a single octet string using the two octet delimiter of 0x3B, 0x20 (the ASCII string "; ") before being passed into a non-HTTP/2 context, such as an HTTP/1.1 connection, or a generic HTTP server application.

Therefore, the following two lists of Cookie header fields are semantically equivalent.

```
cookie: a=b; c=d; e=f
```

```
cookie: a=b  
cookie: c=d  
cookie: e=f
```

8.1.2.6. Malformed Requests and Responses

A malformed request or response is one that is an otherwise valid sequence of HTTP/2 frames, but is otherwise invalid due to the presence of extraneous frames, prohibited header fields, the absence of mandatory header fields, or the inclusion of uppercase header field names.

A request or response that includes a payload body can include a "content-length" header field. A request or response is also malformed if the value of a "content-length" header field does not equal the sum of the DATA frame payload lengths that form the body. A response that is defined to have no payload, as described in [RFC7230], Section 3.3.2, can have a non-zero "content-length" header field, even though no content is included in DATA frames.

Intermediaries that process HTTP requests or responses (i.e., any intermediary not acting as a tunnel) MUST NOT forward a malformed request or response. Malformed requests or responses that are

detected MUST be treated as a stream error (Section 5.4.2) of type `PROTOCOL_ERROR`.

For malformed requests, a server MAY send an HTTP response prior to closing or resetting the stream. Clients MUST NOT accept a malformed response. Note that these requirements are intended to protect against several types of common attacks against HTTP; they are deliberately strict, because being permissive can expose implementations to these vulnerabilities.

8.1.3. Examples

This section shows HTTP/1.1 requests and responses, with illustrations of equivalent HTTP/2 requests and responses.

An HTTP GET request includes request header fields and no payload body and is therefore transmitted as a single HEADERS frame, followed by zero or more CONTINUATION frames containing the serialized block of request header fields. The HEADERS frame in the following has both the `END_HEADERS` and `END_STREAM` flags set; no CONTINUATION frames are sent:

```
GET /resource HTTP/1.1          HEADERS
Host: example.org               ==>  + END_STREAM
Accept: image/jpeg              + END_HEADERS
                                :method = GET
                                :scheme = https
                                :path = /resource
                                host = example.org
                                accept = image/jpeg
```

Similarly, a response that includes only response header fields is transmitted as a HEADERS frame (again, followed by zero or more CONTINUATION frames) containing the serialized block of response header fields.

```
HTTP/1.1 304 Not Modified        HEADERS
ETag: "xyzzy"                   ==>  + END_STREAM
Expires: Thu, 23 Jan ...         + END_HEADERS
                                :status = 304
                                etag = "xyzzy"
                                expires = Thu, 23 Jan ...
```

An HTTP POST request that includes request header fields and payload data is transmitted as one HEADERS frame, followed by zero or more CONTINUATION frames containing the request header fields, followed by one or more DATA frames, with the last CONTINUATION (or HEADERS)

frame having the END_HEADERS flag set and the final DATA frame having the END_STREAM flag set:

```
POST /resource HTTP/1.1      HEADERS
Host: example.org            ==>  - END_STREAM
Content-Type: image/jpeg      - END_HEADERS
Content-Length: 123           :method = POST
                               :path = /resource
                               :scheme = https
{binary data}

CONTINUATION
+ END_HEADERS
  content-type = image/jpeg
  host = example.org
  content-length = 123

DATA
+ END_STREAM
{binary data}
```

Note that data contributing to any given header field could be spread between header block fragments. The allocation of header fields to frames in this example is illustrative only.

A response that includes header fields and payload data is transmitted as a HEADERS frame, followed by zero or more CONTINUATION frames, followed by one or more DATA frames, with the last DATA frame in the sequence having the END_STREAM flag set:

```
HTTP/1.1 200 OK              HEADERS
Content-Type: image/jpeg      ==>  - END_STREAM
Content-Length: 123           + END_HEADERS
                               :status = 200
                               content-type = image/jpeg
                               content-length = 123
{binary data}

DATA
+ END_STREAM
{binary data}
```

An informational response using a 1xx status code other than 101 is transmitted as a HEADERS frame, followed by zero or more CONTINUATION frames.

Trailing header fields are sent as a header block after both the request or response header block and all the DATA frames have been sent. The HEADERS frame starting the trailers header block has the END_STREAM flag set.

The following example includes both a 100 (Continue) status code, which is sent in response to a request containing a "100-continue" token in the Expect header field, and trailing header fields:

HTTP/1.1 100 Continue		HEADERS
Extension-Field: bar	==>	- END_STREAM
		+ END_HEADERS
		:status = 100
		extension-field = bar
HTTP/1.1 200 OK		HEADERS
Content-Type: image/jpeg	==>	- END_STREAM
Transfer-Encoding: chunked		+ END_HEADERS
Trailer: Foo		:status = 200
		content-length = 123
123		content-type = image/jpeg
{binary data}		trailer = Foo
0		
Foo: bar		DATA
		- END_STREAM
		{binary data}
		HEADERS
		+ END_STREAM
		+ END_HEADERS
		foo = bar

8.1.4. Request Reliability Mechanisms in HTTP/2

In HTTP/1.1, an HTTP client is unable to retry a non-idempotent request when an error occurs, because there is no means to determine the nature of the error. It is possible that some server processing occurred prior to the error, which could result in undesirable effects if the request were reattempted.

HTTP/2 provides two mechanisms for providing a guarantee to a client that a request has not been processed:

- o The GOAWAY frame indicates the highest stream number that might have been processed. Requests on streams with higher numbers are therefore guaranteed to be safe to retry.
- o The REFUSED_STREAM error code can be included in a RST_STREAM frame to indicate that the stream is being closed prior to any processing having occurred. Any request that was sent on the reset stream can be safely retried.

Requests that have not been processed have not failed; clients MAY automatically retry them, even those with non-idempotent methods.

A server MUST NOT indicate that a stream has not been processed unless it can guarantee that fact. If frames that are on a stream are passed to the application layer for any stream, then `REFUSED_STREAM` MUST NOT be used for that stream, and a `GOAWAY` frame MUST include a stream identifier that is greater than or equal to the given stream identifier.

In addition to these mechanisms, the `PING` frame provides a way for a client to easily test a connection. Connections that remain idle can become broken as some middleboxes (for instance, network address translators, or load balancers) silently discard connection bindings. The `PING` frame allows a client to safely test whether a connection is still active without sending a request.

8.2. Server Push

HTTP/2 allows a server to pre-emptively send (or "push") responses (along with corresponding "promised" requests) to a client in association with a previous client-initiated request. This can be useful when the server knows the client will need to have those responses available in order to fully process the response to the original request.

A client can request that server push be disabled, though this is negotiated for each hop independently. The `SETTINGS_ENABLE_PUSH` setting can be set to 0 to indicate that server push is disabled.

Promised requests MUST be cacheable (see [RFC7231], Section 4.2.3), MUST be safe (see [RFC7231], Section 4.2.1) and MUST NOT include a request body. Clients that receive a promised request that is not cacheable, is not known to be safe or that indicates the presence of a request body MUST reset the promised stream with a stream error (Section 5.4.2) of type `PROTOCOL_ERROR`. Note this could result in the promised stream being reset if the client does not recognize a newly defined method as being safe.

Pushed responses that are cacheable (see [RFC7234], Section 3) can be stored by the client, if it implements an HTTP cache. Pushed responses are considered successfully validated on the origin server (e.g., if the "no-cache" cache response directive [RFC7234], Section 5.2.2 is present) while the stream identified by the promised stream ID is still open.

Pushed responses that are not cacheable MUST NOT be stored by any HTTP cache. They MAY be made available to the application separately.

The server MUST include a value in the ":authority" header field for which the server is authoritative (see Section 10.1). A client MUST treat a PUSH_PROMISE for which the server is not authoritative as a stream error (Section 5.4.2) of type `PROTOCOL_ERROR`.

An intermediary can receive pushes from the server and choose not to forward them on to the client. In other words, how to make use of the pushed information is up to that intermediary. Equally, the intermediary might choose to make additional pushes to the client, without any action taken by the server.

A client cannot push. Thus, servers MUST treat the receipt of a PUSH_PROMISE frame as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`. Clients MUST reject any attempt to change the `SETTINGS_ENABLE_PUSH` setting to a value other than 0 by treating the message as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`.

8.2.1. Push Requests

Server push is semantically equivalent to a server responding to a request; however, in this case that request is also sent by the server, as a PUSH_PROMISE frame.

The PUSH_PROMISE frame includes a header block that contains a complete set of request header fields that the server attributes to the request. It is not possible to push a response to a request that includes a request body.

Pushed responses are always associated with an explicit request from the client. The PUSH_PROMISE frames sent by the server are sent on that explicit request's stream. The PUSH_PROMISE frame also includes a promised stream identifier, chosen from the stream identifiers available to the server (see Section 5.1.1).

The header fields in PUSH_PROMISE and any subsequent CONTINUATION frames MUST be a valid and complete set of request header fields (Section 8.1.2.3). The server MUST include a method in the ":method" header field that is safe and cacheable. If a client receives a PUSH_PROMISE that does not include a complete and valid set of header fields, or the ":method" header field identifies a method that is not safe, it MUST respond with a stream error (Section 5.4.2) of type `PROTOCOL_ERROR`.

The server SHOULD send PUSH_PROMISE (Section 6.6) frames prior to sending any frames that reference the promised responses. This avoids a race where clients issue requests prior to receiving any PUSH_PROMISE frames.

For example, if the server receives a request for a document containing embedded links to multiple image files, and the server chooses to push those additional images to the client, sending push promises before the DATA frames that contain the image links ensures that the client is able to see the promises before discovering embedded links. Similarly, if the server pushes responses referenced by the header block (for instance, in Link header fields), sending the push promises before sending the header block ensures that clients do not request them.

PUSH_PROMISE frames MUST NOT be sent by the client.

PUSH_PROMISE frames can be sent by the server in response to any client-initiated stream, but the stream MUST be in either the "open" or "half closed (remote)" state with respect to the server. PUSH_PROMISE frames are interspersed with the frames that comprise a response, though they cannot be interspersed with HEADERS and CONTINUATION frames that comprise a single header block.

Sending a PUSH_PROMISE frame creates a new stream and puts the stream into the "reserved (local)" state for the server and the "reserved (remote)" state for the client.

8.2.2. Push Responses

After sending the PUSH_PROMISE frame, the server can begin delivering the pushed response as a response (Section 8.1.2.4) on a server-initiated stream that uses the promised stream identifier. The server uses this stream to transmit an HTTP response, using the same sequence of frames as defined in Section 8.1. This stream becomes "half closed" to the client (Section 5.1) after the initial HEADERS frame is sent.

Once a client receives a PUSH_PROMISE frame and chooses to accept the pushed response, the client SHOULD NOT issue any requests for the promised response until after the promised stream has closed.

If the client determines, for any reason, that it does not wish to receive the pushed response from the server, or if the server takes too long to begin sending the promised response, the client can send an RST_STREAM frame, using either the CANCEL or REFUSED_STREAM codes, and referencing the pushed stream's identifier.

A client can use the `SETTINGS_MAX_CONCURRENT_STREAMS` setting to limit the number of responses that can be concurrently pushed by a server. Advertising a `SETTINGS_MAX_CONCURRENT_STREAMS` value of zero disables server push by preventing the server from creating the necessary streams. This does not prohibit a server from sending `PUSH_PROMISE` frames; clients need to reset any promised streams that are not wanted.

Clients receiving a pushed response MUST validate that either the server is authoritative (see Section 10.1), or the proxy that provided the pushed response is configured for the corresponding request. For example, a server that offers a certificate for only the "example.com" DNS-ID or Common Name is not permitted to push a response for "https://www.example.org/doc".

The response for a `PUSH_PROMISE` stream begins with a `HEADERS` frame, which immediately puts the stream into the "half closed (remote)" state for the server and "half closed (local)" state for the client, and ends with a frame bearing `END_STREAM`, which places the stream in the "closed" state.

Note: The client never sends a frame with the `END_STREAM` flag for a server push.

8.3. The CONNECT Method

In HTTP/1.x, the pseudo-method `CONNECT` ([RFC7231], Section 4.3.6) is used to convert an HTTP connection into a tunnel to a remote host. `CONNECT` is primarily used with HTTP proxies to establish a TLS session with an origin server for the purposes of interacting with "https" resources.

In HTTP/2, the `CONNECT` method is used to establish a tunnel over a single HTTP/2 stream to a remote host, for similar purposes. The HTTP header field mapping works as defined in Request Header Fields (Section 8.1.2.3), with a few differences. Specifically:

- o The `:method` header field is set to `"CONNECT"`.
- o The `:scheme` and `:path` header fields MUST be omitted.
- o The `:authority` header field contains the host and port to connect to (equivalent to the authority-form of the request-target of `CONNECT` requests, see [RFC7230], Section 5.3).

A `CONNECT` request that does not conform to these restrictions is malformed (Section 8.1.2.6).

A proxy that supports CONNECT establishes a TCP connection [TCP] to the server identified in the ":authority" header field. Once this connection is successfully established, the proxy sends a HEADERS frame containing a 2xx series status code to the client, as defined in [RFC7231], Section 4.3.6.

After the initial HEADERS frame sent by each peer, all subsequent DATA frames correspond to data sent on the TCP connection. The payload of any DATA frames sent by the client is transmitted by the proxy to the TCP server; data received from the TCP server is assembled into DATA frames by the proxy. Frame types other than DATA or stream management frames (RST_STREAM, WINDOW_UPDATE, and PRIORITY) MUST NOT be sent on a connected stream, and MUST be treated as a stream error (Section 5.4.2) if received.

The TCP connection can be closed by either peer. The END_STREAM flag on a DATA frame is treated as being equivalent to the TCP FIN bit. A client is expected to send a DATA frame with the END_STREAM flag set after receiving a frame bearing the END_STREAM flag. A proxy that receives a DATA frame with the END_STREAM flag set sends the attached data with the FIN bit set on the last TCP segment. A proxy that receives a TCP segment with the FIN bit set sends a DATA frame with the END_STREAM flag set. Note that the final TCP segment or DATA frame could be empty.

A TCP connection error is signaled with RST_STREAM. A proxy treats any error in the TCP connection, which includes receiving a TCP segment with the RST bit set, as a stream error (Section 5.4.2) of type CONNECT_ERROR. Correspondingly, a proxy MUST send a TCP segment with the RST bit set if it detects an error with the stream or the HTTP/2 connection.

9. Additional HTTP Requirements/Considerations

This section outlines attributes of the HTTP protocol that improve interoperability, reduce exposure to known security vulnerabilities, or reduce the potential for implementation variation.

9.1. Connection Management

HTTP/2 connections are persistent. For best performance, it is expected clients will not close connections until it is determined that no further communication with a server is necessary (for example, when a user navigates away from a particular web page), or until the server closes the connection.

Clients SHOULD NOT open more than one HTTP/2 connection to a given host and port pair, where host is derived from a URI, a selected alternative service [ALT-SVC], or a configured proxy.

A client can create additional connections as replacements, either to replace connections that are near to exhausting the available stream identifier space (Section 5.1.1), to refresh the keying material for a TLS connection, or to replace connections that have encountered errors (Section 5.4.1).

A client MAY open multiple connections to the same IP address and TCP port using different Server Name Indication [TLS-EXT] values or to provide different TLS client certificates, but SHOULD avoid creating multiple connections with the same configuration.

Servers are encouraged to maintain open connections for as long as possible, but are permitted to terminate idle connections if necessary. When either endpoint chooses to close the transport-layer TCP connection, the terminating endpoint SHOULD first send a GOAWAY (Section 6.8) frame so that both endpoints can reliably determine whether previously sent frames have been processed and gracefully complete or terminate any necessary remaining tasks.

9.1.1.1. Connection Reuse

Connections that are made to an origin server, either directly or through a tunnel created using the CONNECT method (Section 8.3) MAY be reused for requests with multiple different URI authority components. A connection can be reused as long as the origin server is authoritative (Section 10.1). For TCP connections without TLS, this depends on the host having resolved to the same IP address.

For "https" resources, connection reuse additionally depends on having a certificate that is valid for the host in the URI. The certificate presented by the server MUST satisfy any checks that the client would perform when forming a new TLS connection for the host in the URI.

An origin server might offer a certificate with multiple "subjectAltName" attributes, or names with wildcards, one of which is valid for the authority in the URI. For example, a certificate with a "subjectAltName" of "*.example.com" might permit the use of the same connection for requests to URIs starting with "https://a.example.com/" and "https://b.example.com/".

In some deployments, reusing a connection for multiple origins can result in requests being directed to the wrong origin server. For example, TLS termination might be performed by a middlebox that uses

the TLS Server Name Indication (SNI) [TLS-EXT] extension to select an origin server. This means that it is possible for clients to send confidential information to servers that might not be the intended target for the request, even though the server is otherwise authoritative.

A server that does not wish clients to reuse connections can indicate that it is not authoritative for a request by sending a 421 (Misdirected Request) status code in response to the request (see Section 9.1.2).

A client that is configured to use a proxy over HTTP/2 directs requests to that proxy through a single connection. That is, all requests sent via a proxy reuse the connection to the proxy.

9.1.2. The 421 (Misdirected Request) Status Code

The 421 (Misdirected Request) status code indicates that the request was directed at a server that is not able to produce a response. This can be sent by a server that is not configured to produce responses for the combination of scheme and authority that are included in the request URI.

Clients receiving a 421 (Misdirected Request) response from a server MAY retry the request - whether the request method is idempotent or not - over a different connection. This is possible if a connection is reused (Section 9.1.1) or if an alternative service is selected ([ALT-SVC]).

This status code MUST NOT be generated by proxies.

A 421 response is cacheable by default; i.e., unless otherwise indicated by the method definition or explicit cache controls (see Section 4.2.2 of [RFC7234]).

9.2. Use of TLS Features

Implementations of HTTP/2 MUST use TLS [TLS12] version 1.2 or higher for HTTP/2 over TLS. The general TLS usage guidance in [TLSBCP] SHOULD be followed, with some additional restrictions that are specific to HTTP/2.

The TLS implementation MUST support the Server Name Indication (SNI) [TLS-EXT] extension to TLS. HTTP/2 clients MUST indicate the target domain name when negotiating TLS.

Deployments of HTTP/2 that negotiate TLS 1.3 or higher need only support and use the SNI extension; deployments of TLS 1.2 are subject

to the requirements in the following sections. Implementations are encouraged to provide defaults that comply, but it is recognized that deployments are ultimately responsible for compliance.

9.2.1. TLS 1.2 Features

This section describes restrictions on the TLS 1.2 feature set that can be used with HTTP/2. Due to deployment limitations, it might not be possible to fail TLS negotiation when these restrictions are not met. An endpoint MAY immediately terminate an HTTP/2 connection that does not meet these TLS requirements with a connection error (Section 5.4.1) of type `INADEQUATE_SECURITY`.

A deployment of HTTP/2 over TLS 1.2 MUST disable compression. TLS compression can lead to the exposure of information that would not otherwise be revealed [RFC3749]. Generic compression is unnecessary since HTTP/2 provides compression features that are more aware of context and therefore likely to be more appropriate for use for performance, security or other reasons.

A deployment of HTTP/2 over TLS 1.2 MUST disable renegotiation. An endpoint MUST treat a TLS renegotiation as a connection error (Section 5.4.1) of type `PROTOCOL_ERROR`. Note that disabling renegotiation can result in long-lived connections becoming unusable due to limits on the number of messages the underlying cipher suite can encipher.

An endpoint MAY use renegotiation to provide confidentiality protection for client credentials offered in the handshake, but any renegotiation MUST occur prior to sending the connection preface. A server SHOULD request a client certificate if it sees a renegotiation request immediately after establishing a connection.

This effectively prevents the use of renegotiation in response to a request for a specific protected resource. A future specification might provide a way to support this use case. Alternatively, a server might use an error (Section 5.4) of type `HTTP_1_1_REQUIRED` to request the client use a protocol which supports renegotiation.

Implementations MUST support ephemeral key exchange sizes of at least 2048 bits for cipher suites that use ephemeral finite field Diffie-Hellman (DHE) [TLS12] and 224 bits for cipher suites that use ephemeral elliptic curve Diffie-Hellman (ECDHE) [RFC4492]. Clients MUST accept DHE sizes of up to 4096 bits. Endpoints MAY treat negotiation of key sizes smaller than the lower limits as a connection error (Section 5.4.1) of type `INADEQUATE_SECURITY`.

9.2.2. TLS 1.2 Cipher Suites

A deployment of HTTP/2 over TLS 1.2 SHOULD NOT use any of the cipher suites that are listed in the cipher suite black list (Appendix A).

Endpoints MAY choose to generate a connection error (Section 5.4.1) of type INADEQUATE_SECURITY if one of the cipher suites from the black list are negotiated. A deployment that chooses to use a black-listed cipher suite risks triggering a connection error unless the set of potential peers is known to accept that cipher suite.

Implementations MUST NOT generate this error in reaction to the negotiation of a cipher suite that is not on the black list. Consequently, when clients offer a cipher suite that is not on the black list, they have to be prepared to use that cipher suite with HTTP/2.

The black list includes the cipher suite that TLS 1.2 makes mandatory, which means that TLS 1.2 deployments could have non-intersecting sets of permitted cipher suites. To avoid this problem causing TLS handshake failures, deployments of HTTP/2 that use TLS 1.2 MUST support TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 [TLS-ECDHE] with the P256 elliptic curve [FIPS186].

Note that clients might advertise support of cipher suites that are on the black list in order to allow for connection to servers that do not support HTTP/2. This allows servers to select HTTP/1.1 with a cipher suite that is on the HTTP/2 black list. However, this can result in HTTP/2 being negotiated with a black-listed cipher suite if the application protocol and cipher suite are independently selected.

10. Security Considerations

10.1. Server Authority

HTTP/2 relies on the HTTP/1.1 definition of authority for determining whether a server is authoritative in providing a given response, see [RFC7230], Section 9.1. This relies on local name resolution for the "http" URI scheme, and the authenticated server identity for the "https" scheme (see [RFC2818], Section 3).

10.2. Cross-Protocol Attacks

In a cross-protocol attack, an attacker causes a client to initiate a transaction in one protocol toward a server that understands a different protocol. An attacker might be able to cause the transaction to appear as valid transaction in the second protocol.

In combination with the capabilities of the web context, this can be used to interact with poorly protected servers in private networks.

Completing a TLS handshake with an ALPN identifier for HTTP/2 can be considered sufficient protection against cross protocol attacks. ALPN provides a positive indication that a server is willing to proceed with HTTP/2, which prevents attacks on other TLS-based protocols.

The encryption in TLS makes it difficult for attackers to control the data which could be used in a cross-protocol attack on a cleartext protocol.

The cleartext version of HTTP/2 has minimal protection against cross-protocol attacks. The connection preface (Section 3.5) contains a string that is designed to confuse HTTP/1.1 servers, but no special protection is offered for other protocols. A server that is willing to ignore parts of an HTTP/1.1 request containing an Upgrade header field in addition to the client connection preface could be exposed to a cross-protocol attack.

10.3. Intermediary Encapsulation Attacks

The HTTP/2 header field encoding allows the expression of names that are not valid field names in the Internet Message Syntax used by HTTP/1.1. Requests or responses containing invalid header field names MUST be treated as malformed (Section 8.1.2.6). An intermediary therefore cannot translate an HTTP/2 request or response containing an invalid field name into an HTTP/1.1 message.

Similarly, HTTP/2 allows header field values that are not valid. While most of the values that can be encoded will not alter header field parsing, carriage return (CR, ASCII 0xd), line feed (LF, ASCII 0xa), and the zero character (NUL, ASCII 0x0) might be exploited by an attacker if they are translated verbatim. Any request or response that contains a character not permitted in a header field value MUST be treated as malformed (Section 8.1.2.6). Valid characters are defined by the "field-content" ABNF rule in Section 3.2 of [RFC7230].

10.4. Cacheability of Pushed Responses

Pushed responses do not have an explicit request from the client; the request is provided by the server in the PUSH_PROMISE frame.

Caching responses that are pushed is possible based on the guidance provided by the origin server in the Cache-Control header field. However, this can cause issues if a single server hosts more than one

tenant. For example, a server might offer multiple users each a small portion of its URI space.

Where multiple tenants share space on the same server, that server MUST ensure that tenants are not able to push representations of resources that they do not have authority over. Failure to enforce this would allow a tenant to provide a representation that would be served out of cache, overriding the actual representation that the authoritative tenant provides.

Pushed responses for which an origin server is not authoritative (see Section 10.1) MUST NOT be used or cached.

10.5. Denial of Service Considerations

An HTTP/2 connection can demand a greater commitment of resources to operate than a HTTP/1.1 connection. The use of header compression and flow control depend on a commitment of resources for storing a greater amount of state. Settings for these features ensure that memory commitments for these features are strictly bounded.

The number of PUSH_PROMISE frames is not constrained in the same fashion. A client that accepts server push SHOULD limit the number of streams it allows to be in the "reserved (remote)" state. Excessive number of server push streams can be treated as a stream error (Section 5.4.2) of type ENHANCE_YOUR_CALM.

Processing capacity cannot be guarded as effectively as state capacity.

The SETTINGS frame can be abused to cause a peer to expend additional processing time. This might be done by pointlessly changing SETTINGS parameters, setting multiple undefined parameters, or changing the same setting multiple times in the same frame. WINDOW_UPDATE or PRIORITY frames can be abused to cause an unnecessary waste of resources.

Large numbers of small or empty frames can be abused to cause a peer to expend time processing frame headers. Note however that some uses are entirely legitimate, such as the sending of an empty DATA or CONTINUATION frame at the end of a stream.

Header compression also offers some opportunities to waste processing resources; see Section 7 of [COMPRESSION] for more details on potential abuses.

Limits in SETTINGS parameters cannot be reduced instantaneously, which leaves an endpoint exposed to behavior from a peer that could

exceed the new limits. In particular, immediately after establishing a connection, limits set by a server are not known to clients and could be exceeded without being an obvious protocol violation.

All these features - i.e., SETTINGS changes, small frames, header compression - have legitimate uses. These features become a burden only when they are used unnecessarily or to excess.

An endpoint that doesn't monitor this behavior exposes itself to a risk of denial of service attack. Implementations SHOULD track the use of these features and set limits on their use. An endpoint MAY treat activity that is suspicious as a connection error (Section 5.4.1) of type `ENHANCE_YOUR_CALM`.

10.5.1. Limits on Header Block Size

A large header block (Section 4.3) can cause an implementation to commit a large amount of state. Header fields that are critical for routing can appear toward the end of a header block, which prevents streaming of header fields to their ultimate destination. This ordering and other reasons, such as ensuring cache correctness, means that an endpoint might need to buffer the entire header block. Since there is no hard limit to the size of a header block, some endpoints could be forced to commit a large amount of available memory for header fields.

An endpoint can use the `SETTINGS_MAX_HEADER_LIST_SIZE` to advise peers of limits that might apply on the size of header blocks. This setting is only advisory, so endpoints MAY choose to send header blocks that exceed this limit and risk having the request or response being treated as malformed. This setting is specific to a connection, so any request or response could encounter a hop with a lower, unknown limit. An intermediary can attempt to avoid this problem by passing on values presented by different peers, but they are not obligated to do so.

A server that receives a larger header block than it is willing to handle can send an HTTP 431 (Request Header Fields Too Large) status code [RFC6585]. A client can discard responses that it cannot process. The header block MUST be processed to ensure a consistent connection state, unless the connection is closed.

10.5.2. CONNECT Issues

The CONNECT method can be used to create disproportionate load on an proxy, since stream creation is relatively inexpensive when compared to the creation and maintenance of a TCP connection. A proxy might also maintain some resources for a TCP connection beyond the closing

of the stream that carries the CONNECT request, since the outgoing TCP connection remains in the TIME_WAIT state. A proxy therefore cannot rely on SETTINGS_MAX_CONCURRENT_STREAMS alone to limit the resources consumed by CONNECT requests.

10.6. Use of Compression

Compression can allow an attacker to recover secret data when it is compressed in the same context as data under attacker control. HTTP/2 enables compression of header fields (Section 4.3); the following concerns also apply to the use of HTTP compressed content-codings ([RFC7231], Section 3.1.2.1).

There are demonstrable attacks on compression that exploit the characteristics of the web (e.g., [BREACH]). The attacker induces multiple requests containing varying plaintext, observing the length of the resulting ciphertext in each, which reveals a shorter length when a guess about the secret is correct.

Implementations communicating on a secure channel MUST NOT compress content that includes both confidential and attacker-controlled data unless separate compression dictionaries are used for each source of data. Compression MUST NOT be used if the source of data cannot be reliably determined. Generic stream compression, such as that provided by TLS MUST NOT be used with HTTP/2 (see Section 9.2).

Further considerations regarding the compression of header fields are described in [COMPRESSION].

10.7. Use of Padding

Padding within HTTP/2 is not intended as a replacement for general purpose padding, such as might be provided by TLS [TLS12]. Redundant padding could even be counterproductive. Correct application can depend on having specific knowledge of the data that is being padded.

To mitigate attacks that rely on compression, disabling or limiting compression might be preferable to padding as a countermeasure.

Padding can be used to obscure the exact size of frame content, and is provided to mitigate specific attacks within HTTP. For example, attacks where compressed content includes both attacker-controlled plaintext and secret data (see for example, [BREACH]).

Use of padding can result in less protection than might seem immediately obvious. At best, padding only makes it more difficult for an attacker to infer length information by increasing the number of frames an attacker has to observe. Incorrectly implemented

padding schemes can be easily defeated. In particular, randomized padding with a predictable distribution provides very little protection; similarly, padding payloads to a fixed size exposes information as payload sizes cross the fixed size boundary, which could be possible if an attacker can control plaintext.

Intermediaries SHOULD retain padding for DATA frames, but MAY drop padding for HEADERS and PUSH_PROMISE frames. A valid reason for an intermediary to change the amount of padding of frames is to improve the protections that padding provides.

10.8. Privacy Considerations

Several characteristics of HTTP/2 provide an observer an opportunity to correlate actions of a single client or server over time. This includes the value of settings, the manner in which flow control windows are managed, the way priorities are allocated to streams, timing of reactions to stimulus, and handling of any features that are controlled by settings.

As far as this creates observable differences in behavior, they could be used as a basis for fingerprinting a specific client, as defined in Section 1.8 of [HTML5].

HTTP/2's preference for using a single TCP connection allows correlation of a user's activity on a site. If connections are reused for different origins, this allows tracking across those origins.

Because the PING and SETTINGS frames solicit immediate responses, they can be used by an endpoint to measure latency to their peer. This might have privacy implications in certain scenarios.

11. IANA Considerations

A string for identifying HTTP/2 is entered into the "Application Layer Protocol Negotiation (ALPN) Protocol IDs" registry established in [TLS-ALPN].

This document establishes a registry for frame types, settings, and error codes. These new registries are entered into a new "Hypertext Transfer Protocol (HTTP) 2 Parameters" section.

This document registers the "HTTP2-Settings" header field for use in HTTP; and the 421 (Misdirected Request) status code.

This document registers the "PRI" method for use in HTTP, to avoid collisions with the connection preface (Section 3.5).

11.1. Registration of HTTP/2 Identification Strings

This document creates two registrations for the identification of HTTP/2 in the "Application Layer Protocol Negotiation (ALPN) Protocol IDs" registry established in [TLS-ALPN].

The "h2" string identifies HTTP/2 when used over TLS:

Protocol: HTTP/2 over TLS

Identification Sequence: 0x68 0x32 ("h2")

Specification: This document

The "h2c" string identifies HTTP/2 when used over cleartext TCP:

Protocol: HTTP/2 over TCP

Identification Sequence: 0x68 0x32 0x63 ("h2c")

Specification: This document

11.2. Frame Type Registry

This document establishes a registry for HTTP/2 frame type codes. The "HTTP/2 Frame Type" registry manages an 8-bit space. The "HTTP/2 Frame Type" registry operates under either of the "IETF Review" or "IESG Approval" policies [RFC5226] for values between 0x00 and 0xef, with values between 0xf0 and 0xff being reserved for experimental use.

New entries in this registry require the following information:

Frame Type: A name or label for the frame type.

Code: The 8-bit code assigned to the frame type.

Specification: A reference to a specification that includes a description of the frame layout, its semantics, and flags that the frame type uses, including any parts of the frame that are conditionally present based on the value of flags.

The entries in the following table are registered by this document.

Frame Type	Code	Section
DATA	0x0	Section 6.1
HEADERS	0x1	Section 6.2
PRIORITY	0x2	Section 6.3
RST_STREAM	0x3	Section 6.4
SETTINGS	0x4	Section 6.5
PUSH_PROMISE	0x5	Section 6.6
PING	0x6	Section 6.7
GOAWAY	0x7	Section 6.8
WINDOW_UPDATE	0x8	Section 6.9
CONTINUATION	0x9	Section 6.10

11.3. Settings Registry

This document establishes a registry for HTTP/2 settings. The "HTTP/2 Settings" registry manages a 16-bit space. The "HTTP/2 Settings" registry operates under the "Expert Review" policy [RFC5226] for values in the range from 0x0000 to 0xffff, with values between and 0xf000 and 0xffff being reserved for experimental use.

New registrations are advised to provide the following information:

Name: A symbolic name for the setting. Specifying a setting name is optional.

Code: The 16-bit code assigned to the setting.

Initial Value: An initial value for the setting.

Specification: An optional reference to a specification that describes the use of the setting.

An initial set of setting registrations can be found in Section 6.5.2.

Name	Code	Initial Value	Specification
HEADER_TABLE_SIZE	0x1	4096	Section 6.5.2
ENABLE_PUSH	0x2	1	Section 6.5.2
MAX_CONCURRENT_STREAMS	0x3	(infinite)	Section 6.5.2
INITIAL_WINDOW_SIZE	0x4	65535	Section 6.5.2
MAX_FRAME_SIZE	0x5	16384	Section 6.5.2
MAX_HEADER_LIST_SIZE	0x6	(infinite)	Section 6.5.2

11.4. Error Code Registry

This document establishes a registry for HTTP/2 error codes. The "HTTP/2 Error Code" registry manages a 32-bit space. The "HTTP/2 Error Code" registry operates under the "Expert Review" policy [RFC5226].

Registrations for error codes are required to include a description of the error code. An expert reviewer is advised to examine new registrations for possible duplication with existing error codes. Use of existing registrations is to be encouraged, but not mandated.

New registrations are advised to provide the following information:

Name: A name for the error code. Specifying an error code name is optional.

Code: The 32-bit error code value.

Description: A brief description of the error code semantics, longer if no detailed specification is provided.

Specification: An optional reference for a specification that defines the error code.

The entries in the following table are registered by this document.

Name	Code	Description	Specification
NO_ERROR	0x0	Graceful shutdown	Section 7
PROTOCOL_ERROR	0x1	Protocol error detected	Section 7
INTERNAL_ERROR	0x2	Implementation fault	Section 7
FLOW_CONTROL_ERROR	0x3	Flow control limits exceeded	Section 7
SETTINGS_TIMEOUT	0x4	Settings not acknowledged	Section 7
STREAM_CLOSED	0x5	Frame received for closed stream	Section 7
FRAME_SIZE_ERROR	0x6	Frame size incorrect	Section 7
REFUSED_STREAM	0x7	Stream not processed	Section 7
CANCEL	0x8	Stream cancelled	Section 7
COMPRESSION_ERROR	0x9	Compression state not updated	Section 7
CONNECT_ERROR	0xa	TCP connection error for CONNECT method	Section 7
ENHANCE_YOUR_CALM	0xb	Processing capacity exceeded	Section 7
INADEQUATE_SECURITY	0xc	Negotiated TLS parameters not acceptable	Section 7
HTTP_1_1_REQUIRED	0xd	Use HTTP/1.1 for the request	Section 7

11.5. HTTP2-Settings Header Field Registration

This section registers the "HTTP2-Settings" header field in the Permanent Message Header Field Registry [BCP90].

Header field name: HTTP2-Settings

Applicable protocol: http

Status: standard

Author/Change controller: IETF

Specification document(s): Section 3.2.1 of this document

Related information: This header field is only used by an HTTP/2 client for Upgrade-based negotiation.

11.6. PRI Method Registration

This section registers the "PRI" method in the HTTP Method Registry ([RFC7231], Section 8.1).

Method Name: PRI

Safe Yes

Idempotent Yes

Specification document(s) Section 3.5 of this document

Related information: This method is never used by an actual client. This method will appear to be used when an HTTP/1.1 server or intermediary attempts to parse an HTTP/2 connection preface.

11.7. The 421 (Misdirected Request) HTTP Status Code

This document registers the 421 (Misdirected Request) HTTP Status code in the Hypertext Transfer Protocol (HTTP) Status Code Registry ([RFC7231], Section 8.2).

Status Code: 421

Short Description: Misdirected Request

Specification: Section 9.1.2 of this document

12. Acknowledgements

This document includes substantial input from the following individuals:

- o Adam Langley, Wan-Teh Chang, Jim Morrison, Mark Nottingham, Alyssa Wilk, Costin Manolache, William Chan, Vitaliy Lvin, Joe Chan, Adam Barth, Ryan Hamilton, Gavin Peters, Kent Alstad, Kevin Lindsay, Paul Amer, Fan Yang, Jonathan Leighton (SPDY contributors).
- o Gabriel Montenegro and Willy Tarreau (Upgrade mechanism).
- o William Chan, Salvatore Loreto, Osama Mazahir, Gabriel Montenegro, Jitu Padhye, Roberto Peon, Rob Trace (Flow control).
- o Mike Bishop (Extensibility).
- o Mark Nottingham, Julian Reschke, James Snell, Jeff Pinner, Mike Bishop, Herve Ruellan (Substantial editorial contributions).

- o Kari Hurtta, Tatsuhiro Tsujikawa, Greg Wilkins, Poul-Henning Kamp, Jonathan Thackray.
- o Alexey Melnikov was an editor of this document during 2013.
- o A substantial proportion of Martin's contribution was supported by Microsoft during his employment there.
- o The Japanese HTTP/2 community provided an invaluable contribution, including a number of implementations, plus numerous technical and editorial contributions.

13. References

13.1. Normative References

[COMPRESSION]

Ruellan, H. and R. Peon, "HPACK - Header Compression for HTTP/2", draft-ietf-httpbis-header-compression-11 (work in progress), February 2015.

[COOKIE] Barth, A., "HTTP State Management Mechanism", RFC 6265, April 2011.

[FIPS186] NIST, "Digital Signature Standard (DSS)", FIPS PUB 186-4, July 2013, <<http://dx.doi.org/10.6028/NIST.FIPS.186-4>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

[RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, October 2006.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.

- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, June 2014.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, June 2014.
- [RFC7232] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests", RFC 7232, June 2014.
- [RFC7233] Fielding, R., Ed., Lafon, Y., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Range Requests", RFC 7233, June 2014.
- [RFC7234] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Caching", RFC 7234, June 2014.
- [RFC7235] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Authentication", RFC 7235, June 2014.
- [TCP] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [TLS-ALPN] Friedl, S., Popov, A., Langley, A., and E. Stephan, "Transport Layer Security (TLS) Application-Layer Protocol Negotiation Extension", RFC 7301, July 2014.
- [TLS-ECDHE] Rescorla, E., "TLS Elliptic Curve Cipher Suites with SHA-256/384 and AES Galois Counter Mode (GCM)", RFC 5289, August 2008.
- [TLS-EXT] Eastlake, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", RFC 6066, January 2011.
- [TLS12] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

13.2. Informative References

- [ALT-SVC] Nottingham, M., McManus, P., and J. Reschke, "HTTP Alternative Services", draft-ietf-httpbis-alt-svc-06 (work in progress), February 2015.

- [BCP90] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", BCP 90, RFC 3864, September 2004.
- [BREACH] Gluck, Y., Harris, N., and A. Prado, "BREACH: Reviving the CRIME Attack", July 2013, <<http://breachattack.com/resources/BREACH%20-%20SSL,%20gone%20in%2030%20seconds.pdf>>.
- [HTML5] Hickson, I., Berjon, R., Faulkner, S., Leithead, T., Doyle Navara, E., O'Connor, E., and S. Pfeiffer, "HTML5", W3C Recommendation REC-html5-20141028, October 2014, <<http://www.w3.org/TR/2014/REC-html5-20141028/>>.
- Latest version available at [5].
- [RFC3749] Hollenbeck, S., "Transport Layer Security Protocol Compression Methods", RFC 3749, May 2004.
- [RFC4492] Blake-Wilson, S., Bolyard, N., Gupta, V., Hawk, C., and B. Moeller, "Elliptic Curve Cryptography (ECC) Cipher Suites for Transport Layer Security (TLS)", RFC 4492, May 2006.
- [RFC6585] Nottingham, M. and R. Fielding, "Additional HTTP Status Codes", RFC 6585, April 2012.
- [RFC7323] Borman, D., Braden, B., Jacobson, V., and R. Scheffenegger, "TCP Extensions for High Performance", RFC 7323, September 2014.
- [TALKING] Huang, L-S., Chen, E., Barth, A., Rescorla, E., and C. Jackson, "Talking to Yourself for Fun and Profit", 2011, <<http://w2spconf.com/2011/papers/websocket.pdf>>.
- [TLSBCP] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of TLS and DTLS", draft-ietf-uta-tls-bcp-08 (work in progress), December 2014.

13.3. URIs

- [1] <https://www.iana.org/assignments/message-headers>
- [2] <https://groups.google.com/forum/?fromgroups#!topic/spdy-dev/cfUef2gL3iU>
- [3] <https://tools.ietf.org/html/draft-montenegro-httpbis-http2-fc-principles-01>

Appendix A. TLS 1.2 Cipher Suite Black List

An HTTP/2 implementation MAY treat the negotiation of any of the following cipher suites with TLS 1.2 as a connection error (Section 5.4.1) of type INADEQUATE_SECURITY: TLS_NULL_WITH_NULL_NULL, TLS_RSA_WITH_NULL_MD5, TLS_RSA_WITH_NULL_SHA, TLS_RSA_EXPORT_WITH_RC4_40_MD5, TLS_RSA_WITH_RC4_128_MD5, TLS_RSA_WITH_RC4_128_SHA, TLS_RSA_EXPORT_WITH_RC2_CBC_40_MD5, TLS_RSA_WITH_IDEA_CBC_SHA, TLS_RSA_EXPORT_WITH_DES40_CBC_SHA, TLS_RSA_WITH_DES_CBC_SHA, TLS_RSA_WITH_3DES_EDE_CBC_SHA, TLS_DH_DSS_EXPORT_WITH_DES40_CBC_SHA, TLS_DH_DSS_WITH_DES_CBC_SHA, TLS_DH_DSS_WITH_3DES_EDE_CBC_SHA, TLS_DH_RSA_EXPORT_WITH_DES40_CBC_SHA, TLS_DH_RSA_WITH_DES_CBC_SHA, TLS_DH_RSA_WITH_3DES_EDE_CBC_SHA, TLS_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA, TLS_DHE_DSS_WITH_DES_CBC_SHA, TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA, TLS_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA, TLS_DHE_RSA_WITH_DES_CBC_SHA, TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA, TLS_DH_anon_EXPORT_WITH_RC4_40_MD5, TLS_DH_anon_WITH_RC4_128_MD5, TLS_DH_anon_EXPORT_WITH_DES40_CBC_SHA, TLS_DH_anon_WITH_DES_CBC_SHA, TLS_DH_anon_WITH_3DES_EDE_CBC_SHA, TLS_KRB5_WITH_DES_CBC_SHA, TLS_KRB5_WITH_3DES_EDE_CBC_SHA, TLS_KRB5_WITH_RC4_128_SHA, TLS_KRB5_WITH_IDEA_CBC_SHA, TLS_KRB5_WITH_DES_CBC_MD5, TLS_KRB5_WITH_3DES_EDE_CBC_MD5, TLS_KRB5_WITH_RC4_128_MD5, TLS_KRB5_WITH_IDEA_CBC_MD5, TLS_KRB5_EXPORT_WITH_DES_CBC_40_SHA, TLS_KRB5_EXPORT_WITH_RC2_CBC_40_SHA, TLS_KRB5_EXPORT_WITH_RC4_40_SHA, TLS_KRB5_EXPORT_WITH_DES_CBC_40_MD5, TLS_KRB5_EXPORT_WITH_RC2_CBC_40_MD5, TLS_KRB5_EXPORT_WITH_RC4_40_MD5, TLS_PSK_WITH_NULL_SHA, TLS_DHE_PSK_WITH_NULL_SHA, TLS_RSA_PSK_WITH_NULL_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_DH_DSS_WITH_AES_128_CBC_SHA, TLS_DH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_DH_anon_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_DH_DSS_WITH_AES_256_CBC_SHA, TLS_DH_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DH_anon_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_NULL_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_DH_DSS_WITH_AES_128_CBC_SHA256, TLS_DH_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_CAMELLIA_128_CBC_SHA, TLS_DH_DSS_WITH_CAMELLIA_128_CBC_SHA, TLS_DH_RSA_WITH_CAMELLIA_128_CBC_SHA, TLS_DHE_DSS_WITH_CAMELLIA_128_CBC_SHA, TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA, TLS_DH_anon_WITH_CAMELLIA_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DH_DSS_WITH_AES_256_CBC_SHA256,

TLS_DH_RSA_WITH_AES_256_CBC_SHA256,
TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,
TLS_DH_anon_WITH_AES_128_CBC_SHA256,
TLS_DH_anon_WITH_AES_256_CBC_SHA256,
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA,
TLS_DH_DSS_WITH_CAMELLIA_256_CBC_SHA,
TLS_DH_RSA_WITH_CAMELLIA_256_CBC_SHA,
TLS_DHE_DSS_WITH_CAMELLIA_256_CBC_SHA,
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA,
TLS_DH_anon_WITH_CAMELLIA_256_CBC_SHA, TLS_PSK_WITH_RC4_128_SHA,
TLS_PSK_WITH_3DES_EDE_CBC_SHA, TLS_PSK_WITH_AES_128_CBC_SHA,
TLS_PSK_WITH_AES_256_CBC_SHA, TLS_DHE_PSK_WITH_RC4_128_SHA,
TLS_DHE_PSK_WITH_3DES_EDE_CBC_SHA, TLS_DHE_PSK_WITH_AES_128_CBC_SHA,
TLS_DHE_PSK_WITH_AES_256_CBC_SHA, TLS_RSA_PSK_WITH_RC4_128_SHA,
TLS_RSA_PSK_WITH_3DES_EDE_CBC_SHA, TLS_RSA_PSK_WITH_AES_128_CBC_SHA,
TLS_RSA_PSK_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_SEED_CBC_SHA,
TLS_DH_DSS_WITH_SEED_CBC_SHA, TLS_DH_RSA_WITH_SEED_CBC_SHA,
TLS_DHE_DSS_WITH_SEED_CBC_SHA, TLS_DHE_RSA_WITH_SEED_CBC_SHA,
TLS_DH_anon_WITH_SEED_CBC_SHA, TLS_RSA_WITH_AES_128_GCM_SHA256,
TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_DH_RSA_WITH_AES_128_GCM_SHA256,
TLS_DH_RSA_WITH_AES_256_GCM_SHA384,
TLS_DH_DSS_WITH_AES_128_GCM_SHA256,
TLS_DH_DSS_WITH_AES_256_GCM_SHA384,
TLS_DH_anon_WITH_AES_128_GCM_SHA256,
TLS_DH_anon_WITH_AES_256_GCM_SHA384, TLS_PSK_WITH_AES_128_GCM_SHA256,
TLS_PSK_WITH_AES_256_GCM_SHA384, TLS_RSA_PSK_WITH_AES_128_GCM_SHA256,
TLS_RSA_PSK_WITH_AES_256_GCM_SHA384, TLS_PSK_WITH_AES_128_CBC_SHA256,
TLS_PSK_WITH_AES_256_CBC_SHA384, TLS_PSK_WITH_NULL_SHA256,
TLS_PSK_WITH_NULL_SHA384, TLS_DHE_PSK_WITH_AES_128_CBC_SHA256,
TLS_DHE_PSK_WITH_AES_256_CBC_SHA384, TLS_DHE_PSK_WITH_NULL_SHA256,
TLS_DHE_PSK_WITH_NULL_SHA384, TLS_RSA_PSK_WITH_AES_128_CBC_SHA256,
TLS_RSA_PSK_WITH_AES_256_CBC_SHA384, TLS_RSA_PSK_WITH_NULL_SHA256,
TLS_RSA_PSK_WITH_NULL_SHA384, TLS_RSA_WITH_CAMELLIA_128_CBC_SHA256,
TLS_DH_DSS_WITH_CAMELLIA_128_CBC_SHA256,
TLS_DH_RSA_WITH_CAMELLIA_128_CBC_SHA256,
TLS_DHE_DSS_WITH_CAMELLIA_128_CBC_SHA256,
TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA256,
TLS_DH_anon_WITH_CAMELLIA_128_CBC_SHA256,
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA256,
TLS_DH_DSS_WITH_CAMELLIA_256_CBC_SHA256,
TLS_DH_RSA_WITH_CAMELLIA_256_CBC_SHA256,
TLS_DHE_DSS_WITH_CAMELLIA_256_CBC_SHA256,
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA256,
TLS_DH_anon_WITH_CAMELLIA_256_CBC_SHA256,
TLS_EMPTY_RENEGOTIATION_INFO_SCSV, TLS_ECDH_ECDSA_WITH_NULL_SHA,
TLS_ECDH_ECDSA_WITH_RC4_128_SHA,
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA,

TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA,
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_NULL_SHA,
TLS_ECDHE_ECDSA_WITH_RC4_128_SHA,
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA,
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_NULL_SHA,
TLS_ECDH_RSA_WITH_RC4_128_SHA, TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA,
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA,
TLS_ECDHE_RSA_WITH_NULL_SHA, TLS_ECDHE_RSA_WITH_RC4_128_SHA,
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA,
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_anon_WITH_NULL_SHA,
TLS_ECDH_anon_WITH_RC4_128_SHA, TLS_ECDH_anon_WITH_3DES_EDE_CBC_SHA,
TLS_ECDH_anon_WITH_AES_128_CBC_SHA,
TLS_ECDH_anon_WITH_AES_256_CBC_SHA,
TLS_SRP_SHA_WITH_3DES_EDE_CBC_SHA,
TLS_SRP_SHA_RSA_WITH_3DES_EDE_CBC_SHA,
TLS_SRP_SHA_DSS_WITH_3DES_EDE_CBC_SHA,
TLS_SRP_SHA_WITH_AES_128_CBC_SHA,
TLS_SRP_SHA_RSA_WITH_AES_128_CBC_SHA,
TLS_SRP_SHA_DSS_WITH_AES_128_CBC_SHA,
TLS_SRP_SHA_WITH_AES_256_CBC_SHA,
TLS_SRP_SHA_RSA_WITH_AES_256_CBC_SHA,
TLS_SRP_SHA_DSS_WITH_AES_256_CBC_SHA,
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256,
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384,
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256,
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384,
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256,
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384,
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256,
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_PSK_WITH_RC4_128_SHA,
TLS_ECDHE_PSK_WITH_3DES_EDE_CBC_SHA,
TLS_ECDHE_PSK_WITH_AES_128_CBC_SHA,
TLS_ECDHE_PSK_WITH_AES_256_CBC_SHA,
TLS_ECDHE_PSK_WITH_AES_128_CBC_SHA256,
TLS_ECDHE_PSK_WITH_AES_256_CBC_SHA384, TLS_ECDHE_PSK_WITH_NULL_SHA,
TLS_ECDHE_PSK_WITH_NULL_SHA256, TLS_ECDHE_PSK_WITH_NULL_SHA384,
TLS_RSA_WITH_ARIA_128_CBC_SHA256, TLS_RSA_WITH_ARIA_256_CBC_SHA384,
TLS_DH_DSS_WITH_ARIA_128_CBC_SHA256,
TLS_DH_DSS_WITH_ARIA_256_CBC_SHA384,
TLS_DH_RSA_WITH_ARIA_128_CBC_SHA256,
TLS_DH_RSA_WITH_ARIA_256_CBC_SHA384,
TLS_DHE_DSS_WITH_ARIA_128_CBC_SHA256,

TLS_DHE_DSS_WITH_ARIA_256_CBC_SHA384,
TLS_DHE_RSA_WITH_ARIA_128_CBC_SHA256,
TLS_DHE_RSA_WITH_ARIA_256_CBC_SHA384,
TLS_DH_anon_WITH_ARIA_128_CBC_SHA256,
TLS_DH_anon_WITH_ARIA_256_CBC_SHA384,
TLS_ECDHE_ECDSA_WITH_ARIA_128_CBC_SHA256,
TLS_ECDHE_ECDSA_WITH_ARIA_256_CBC_SHA384,
TLS_ECDH_ECDSA_WITH_ARIA_128_CBC_SHA256,
TLS_ECDH_ECDSA_WITH_ARIA_256_CBC_SHA384,
TLS_ECDHE_RSA_WITH_ARIA_128_CBC_SHA256,
TLS_ECDHE_RSA_WITH_ARIA_256_CBC_SHA384,
TLS_ECDH_RSA_WITH_ARIA_128_CBC_SHA256,
TLS_ECDH_RSA_WITH_ARIA_256_CBC_SHA384,
TLS_RSA_WITH_ARIA_128_GCM_SHA256, TLS_RSA_WITH_ARIA_256_GCM_SHA384,
TLS_DH_RSA_WITH_ARIA_128_GCM_SHA256,
TLS_DH_RSA_WITH_ARIA_256_GCM_SHA384,
TLS_DH_DSS_WITH_ARIA_128_GCM_SHA256,
TLS_DH_DSS_WITH_ARIA_256_GCM_SHA384,
TLS_DH_anon_WITH_ARIA_128_GCM_SHA256,
TLS_DH_anon_WITH_ARIA_256_GCM_SHA384,
TLS_ECDH_ECDSA_WITH_ARIA_128_GCM_SHA256,
TLS_ECDH_ECDSA_WITH_ARIA_256_GCM_SHA384,
TLS_ECDH_RSA_WITH_ARIA_128_GCM_SHA256,
TLS_ECDH_RSA_WITH_ARIA_256_GCM_SHA384,
TLS_PSK_WITH_ARIA_128_CBC_SHA256, TLS_PSK_WITH_ARIA_256_CBC_SHA384,
TLS_DHE_PSK_WITH_ARIA_128_CBC_SHA256,
TLS_DHE_PSK_WITH_ARIA_256_CBC_SHA384,
TLS_RSA_PSK_WITH_ARIA_128_CBC_SHA256,
TLS_RSA_PSK_WITH_ARIA_256_CBC_SHA384,
TLS_PSK_WITH_ARIA_128_GCM_SHA256, TLS_PSK_WITH_ARIA_256_GCM_SHA384,
TLS_RSA_PSK_WITH_ARIA_128_GCM_SHA256,
TLS_RSA_PSK_WITH_ARIA_256_GCM_SHA384,
TLS_ECDHE_PSK_WITH_ARIA_128_CBC_SHA256,
TLS_ECDHE_PSK_WITH_ARIA_256_CBC_SHA384,
TLS_ECDHE_ECDSA_WITH_CAMELLIA_128_CBC_SHA256,
TLS_ECDHE_ECDSA_WITH_CAMELLIA_256_CBC_SHA384,
TLS_ECDH_ECDSA_WITH_CAMELLIA_128_CBC_SHA256,
TLS_ECDH_ECDSA_WITH_CAMELLIA_256_CBC_SHA384,
TLS_ECDHE_RSA_WITH_CAMELLIA_128_CBC_SHA256,
TLS_ECDHE_RSA_WITH_CAMELLIA_256_CBC_SHA384,
TLS_ECDH_RSA_WITH_CAMELLIA_128_CBC_SHA256,
TLS_ECDH_RSA_WITH_CAMELLIA_256_CBC_SHA384,
TLS_RSA_WITH_CAMELLIA_128_GCM_SHA256,
TLS_RSA_WITH_CAMELLIA_256_GCM_SHA384,
TLS_DH_RSA_WITH_CAMELLIA_128_GCM_SHA256,
TLS_DH_RSA_WITH_CAMELLIA_256_GCM_SHA384,
TLS_DH_DSS_WITH_CAMELLIA_128_GCM_SHA256,
TLS_DH_DSS_WITH_CAMELLIA_256_GCM_SHA384,

TLS_DH_anon_WITH_CAMELLIA_128_GCM_SHA256,
TLS_DH_anon_WITH_CAMELLIA_256_GCM_SHA384,
TLS_ECDH_ECDSA_WITH_CAMELLIA_128_GCM_SHA256,
TLS_ECDH_ECDSA_WITH_CAMELLIA_256_GCM_SHA384,
TLS_ECDH_RSA_WITH_CAMELLIA_128_GCM_SHA256,
TLS_ECDH_RSA_WITH_CAMELLIA_256_GCM_SHA384,
TLS_PSK_WITH_CAMELLIA_128_GCM_SHA256,
TLS_PSK_WITH_CAMELLIA_256_GCM_SHA384,
TLS_RSA_PSK_WITH_CAMELLIA_128_GCM_SHA256,
TLS_RSA_PSK_WITH_CAMELLIA_256_GCM_SHA384,
TLS_PSK_WITH_CAMELLIA_128_CBC_SHA256,
TLS_PSK_WITH_CAMELLIA_256_CBC_SHA384,
TLS_DHE_PSK_WITH_CAMELLIA_128_CBC_SHA256,
TLS_DHE_PSK_WITH_CAMELLIA_256_CBC_SHA384,
TLS_RSA_PSK_WITH_CAMELLIA_128_CBC_SHA256,
TLS_RSA_PSK_WITH_CAMELLIA_256_CBC_SHA384,
TLS_ECDHE_PSK_WITH_CAMELLIA_128_CBC_SHA256,
TLS_ECDHE_PSK_WITH_CAMELLIA_256_CBC_SHA384, TLS_RSA_WITH_AES_128_CCM,
TLS_RSA_WITH_AES_256_CCM, TLS_RSA_WITH_AES_128_CCM_8,
TLS_RSA_WITH_AES_256_CCM_8, TLS_PSK_WITH_AES_128_CCM,
TLS_PSK_WITH_AES_256_CCM, TLS_PSK_WITH_AES_128_CCM_8,
TLS_PSK_WITH_AES_256_CCM_8.

Note: This list was assembled from the set of registered TLS cipher suites at the time of writing. This list includes those cipher suites that do not offer an ephemeral key exchange and those that are based on the TLS null, stream or block cipher type (as defined in Section 6.2.3 of [TLS12]). Additional cipher suites with these properties could be defined; these would not be explicitly prohibited.

Appendix B. Change Log

This section is to be removed by RFC Editor before publication.

B.1. Since draft-ietf-httpbis-http2-15

Enabled the sending of PRIORITY for any stream state.

Added a cipher suite blacklist and made several changes to the TLS usage section.

B.2. Since draft-ietf-httpbis-http2-14

Renamed Not Authoritative status code to Misdirected Request.

Added HTTP_1_1_REQUIRED error code.

B.3. Since draft-ietf-httpbis-http2-13

Pseudo-header fields are now required to appear strictly before regular ones.

Restored 1xx series status codes, except 101.

Changed frame length field 24-bits. Expanded frame header to 9 octets. Added a setting to limit the damage.

Added a setting to advise peers of header set size limits.

Removed segments.

Made non-semantic-bearing HEADERS frames illegal in the HTTP mapping.

B.4. Since draft-ietf-httpbis-http2-12

Restored extensibility options.

Restricting TLS cipher suites to AEAD only.

Removing Content-Encoding requirements.

Permitting the use of PRIORITY after stream close.

Removed ALTSVC frame.

Removed BLOCKED frame.

Reducing the maximum padding size to 256 octets; removing padding from CONTINUATION frames.

Removed per-frame GZIP compression.

B.5. Since draft-ietf-httpbis-http2-11

Added BLOCKED frame (at risk).

Simplified priority scheme.

Added DATA per-frame GZIP compression.

B.6. Since draft-ietf-httpbis-http2-10

Changed "connection header" to "connection preface" to avoid confusion.

Added dependency-based stream prioritization.

Added "h2c" identifier to distinguish between cleartext and secured HTTP/2.

Adding missing padding to PUSH_PROMISE.

Integrate ALTSVC frame and supporting text.

Dropping requirement on "deflate" Content-Encoding.

Improving security considerations around use of compression.

B.7. Since draft-ietf-httpbis-http2-09

Adding padding for data frames.

Renumbering frame types, error codes, and settings.

Adding INADEQUATE_SECURITY error code.

Updating TLS usage requirements to 1.2; forbidding TLS compression.

Removing extensibility for frames and settings.

Changing setting identifier size.

Removing the ability to disable flow control.

Changing the protocol identification token to "h2".

Changing the use of :authority to make it optional and to allow userinfo in non-HTTP cases.

Allowing split on 0x0 for Cookie.

Reserved PRI method in HTTP/1.1 to avoid possible future collisions.

B.8. Since draft-ietf-httpbis-http2-08

Added cookie crumbling for more efficient header compression.

Added header field ordering with the value-concatenation mechanism.

B.9. Since draft-ietf-httpbis-http2-07

Marked draft for implementation.

B.10. Since draft-ietf-httpbis-http2-06

Adding definition for CONNECT method.

Constraining the use of push to safe, cacheable methods with no request body.

Changing from :host to :authority to remove any potential confusion.

Adding setting for header compression table size.

Adding settings acknowledgement.

Removing unnecessary and potentially problematic flags from CONTINUATION.

Added denial of service considerations.

B.11. Since draft-ietf-httpbis-http2-05

Marking the draft ready for implementation.

Renumbering END_PUSH_PROMISE flag.

Editorial clarifications and changes.

B.12. Since draft-ietf-httpbis-http2-04

Added CONTINUATION frame for HEADERS and PUSH_PROMISE.

PUSH_PROMISE is no longer implicitly prohibited if SETTINGS_MAX_CONCURRENT_STREAMS is zero.

Push expanded to allow all safe methods without a request body.

Clarified the use of HTTP header fields in requests and responses. Prohibited HTTP/1.1 hop-by-hop header fields.

Requiring that intermediaries not forward requests with missing or illegal routing :-headers.

Clarified requirements around handling different frames after stream close, stream reset and GOAWAY.

Added more specific prohibitions for sending of different frame types in various stream states.

Making the last received setting value the effective value.

Clarified requirements on TLS version, extension and ciphers.

B.13. Since draft-ietf-httpbis-http2-03

Committed major restructuring atrocities.

Added reference to first header compression draft.

Added more formal description of frame lifecycle.

Moved END_STREAM (renamed from FINAL) back to HEADERS/DATA.

Removed HEADERS+PRIORITY, added optional priority to HEADERS frame.

Added PRIORITY frame.

B.14. Since draft-ietf-httpbis-http2-02

Added continuations to frames carrying header blocks.

Replaced use of "session" with "connection" to avoid confusion with other HTTP stateful concepts, like cookies.

Removed "message".

Switched to TLS ALPN from NPN.

Editorial changes.

B.15. Since draft-ietf-httpbis-http2-01

Added IANA considerations section for frame types, error codes and settings.

Removed data frame compression.

Added PUSH_PROMISE.

Added globally applicable flags to framing.

Removed zlib-based header compression mechanism.

Updated references.

Clarified stream identifier reuse.

Removed CREDENTIALS frame and associated mechanisms.

Added advice against naive implementation of flow control.

Added session header section.

Restructured frame header. Removed distinction between data and control frames.

Altered flow control properties to include session-level limits.

Added note on cacheability of pushed resources and multiple tenant servers.

Changed protocol label form based on discussions.

B.16. Since draft-ietf-httpbis-http2-00

Changed title throughout.

Removed section on Incompatibilities with SPDY draft#2.

Changed INTERNAL_ERROR on GOAWAY to have a value of 2 [6].

Replaced abstract and introduction.

Added section on starting HTTP/2.0, including upgrade mechanism.

Removed unused references.

Added flow control principles (Section 5.2.1) based on [7].

B.17. Since draft-mbelshe-httpbis-spdy-00

Adopted as base for draft-ietf-httpbis-http2.

Updated authors/editors list.

Added status note.

Authors' Addresses

Mike Belshe
Twist

EMail: mbelshe@chromium.org

Roberto Peon
Google, Inc

EMail: fenix@google.com

Martin Thomson (editor)
Mozilla
331 E Evelyn Street
Mountain View, CA 94041
US

EMail: martin.thomson@gmail.com

HTTP Working Group
Internet-Draft
Intended status: Experimental
Expires: September 18, 2017

M. Nottingham

M. Thomson
Mozilla
March 17, 2017

Opportunistic Security for HTTP/2
draft-ietf-httpbis-http2-encryption-11

Abstract

This document describes how "http" URIs can be accessed using Transport Layer Security (TLS) and HTTP/2 to mitigate pervasive monitoring attacks. This mechanism not a replacement for "https" URIs; it is vulnerable to active attacks.

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 September 18, 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.

Table of Contents

1. Introduction	2
1.1. Goals and Non-Goals	2
1.2. Notational Conventions	3
2. Using HTTP URIs over TLS	3
2.1. Alternative Server Opt-In	4
2.2. Interaction with "https" URIs	5
2.3. The "http-opportunistic" well-known URI	5
3. IANA Considerations	6
4. Security Considerations	6
4.1. Security Indicators	6
4.2. Downgrade Attacks	6
4.3. Privacy Considerations	6
4.4. Confusion Regarding Request Scheme	7
4.5. Server Controls	7
5. References	7
5.1. Normative References	7
5.2. Informative References	8
Appendix A. Acknowledgements	9
Authors' Addresses	9

1. Introduction

This document describes a use of HTTP Alternative Services [RFC7838] to decouple the URI scheme from the use and configuration of underlying encryption. It allows an "http" URI to be accessed using HTTP/2 [RFC7230] and Transport Layer Security (TLS) [RFC5246] with Opportunistic Security [RFC7435].

This document describes a usage model whereby sites can serve "http" URIs over TLS, thereby avoiding the problem of serving Mixed Content (described in [W3C.CR-mixed-content-20160802]) while still providing protection against passive attacks.

Opportunistic Security does not provide the same guarantees as using TLS with "https" URIs, because it is vulnerable to active attacks, and does not change the security context of the connection. Normally, users will not be able to tell that it is in use (i.e., there will be no "lock icon").

1.1. Goals and Non-Goals

The immediate goal is to make the use of HTTP more robust in the face of pervasive passive monitoring [RFC7258].

A secondary (but significant) goal is to provide for ease of implementation, deployment and operation. This mechanism is expected

to have a minimal impact upon performance, and require a trivial administrative effort to configure.

Preventing active attacks (such as a Man-in-the-Middle) is a non-goal for this specification. Furthermore, this specification is not intended to replace or offer an alternative to "https", since "https" both prevents active attacks and invokes a more stringent security model in most clients.

1.2. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Using HTTP URIs over TLS

An origin server that supports the resolution of "http" URIs can indicate support for this specification by providing an alternative service advertisement [RFC7838] for a protocol identifier that uses TLS, such as "h2" [RFC7540]. Such a protocol MUST include an explicit indication of the scheme of the resource. This excludes HTTP/1.1; HTTP/1.1 clients are forbidden from including the absolute form of a URI in requests to origin servers (see Section 5.3.1 of [RFC7230]).

A client that receives such an advertisement MAY make future requests intended for the associated origin [RFC6454] to the identified service (as specified by [RFC7838]), provided that the alternative service opts in as described in Section 2.1.

A client that places the importance of protection against passive attacks over performance might choose to withhold requests until an encrypted connection is available. However, if such a connection cannot be successfully established, the client can resume its use of the cleartext connection.

A client can also explicitly probe for an alternative service advertisement by sending a request that bears little or no sensitive information, such as one with the OPTIONS method. Likewise, clients with existing alternative services information could make such a request before they expire, in order minimize the delays that might be incurred.

Client certificates are not meaningful for URLs with the "http" scheme, and therefore clients creating new TLS connections to alternative services for the purposes of this specification MUST NOT present them. A server that also provides "https" resources on the

same port can request a certificate during the TLS handshake, but it MUST NOT abort the handshake if the client does not provide one.

2.1. Alternative Server Opt-In

It is possible that the server might become confused about whether requests' URLs have a "http" or "https" scheme, for various reasons; see Section 4.4. To ensure that the alternative service has opted into serving "http" URLs over TLS, clients are required to perform additional checks before directing "http" requests to it.

Clients MUST NOT send "http" requests over a secured connection, unless the chosen alternative service presents a certificate that is valid for the origin as defined in [RFC2818]. Using an authenticated alternative service establishes "reasonable assurances" for the purposes of [RFC7838]. In addition to authenticating the server, the client MUST have obtained a valid http-opportunistic response for an origin (as per Section 2.3) using the authenticated connection. An exception to the latter restriction is made for requests for the "http-opportunistic" well-known URI.

For example, assuming the following request is made over a TLS connection that is successfully authenticated for those origins, the following request/response pair would allow requests for the origins "http://www.example.com" or "http://example.com" to be sent using a secured connection:

HEADERS

```
+ END_STREAM
+ END_HEADERS
:method = GET
:scheme = http
:authority = example.com
:path = /.well-known/http-opportunistic
```

HEADERS

```
:status = 200
content-type = application/json
```

DATA

```
+ END_STREAM
[ "http://www.example.com", "http://example.com" ]
```

Though this document describes multiple origins, this is only for operational convenience. Only a request made to an origin (over an authenticated connection) can be used to acquire this resource for that origin. Thus in the example, the request to "http://example.com" cannot be assumed to also provide an http-opportunistic response for "http://www.example.com".

2.2. Interaction with "https" URIs

Clients MUST NOT send "http" requests and "https" requests on the same connection. Similarly, clients MUST NOT send "http" requests for multiple origins on the same connection.

2.3. The "http-opportunistic" well-known URI

This specification defines the "http-opportunistic" well-known URI [RFC5785]. A client is said to have a valid http-opportunistic response for a given origin when:

- o The client has requested the well-known URI from the origin over an authenticated connection and a 200 (OK) response was provided, and
- o That response is fresh [RFC7234] (potentially through revalidation [RFC7232]), and
- o That response has the media type "application/json", and
- o That response's payload, when parsed as JSON [RFC7159], contains an array as the root, and
- o The array contains a string that is a case-insensitive character-for-character match for the origin in question, serialised into Unicode as per Section 6.1 of [RFC6454].

A client MAY treat an "http-opportunistic" resource as invalid if values it contains are not strings.

This document does not define semantics for "http-opportunistic" resources on an "https" origin, nor does it define semantics if the resource includes "https" origins.

Allowing clients to cache the http-opportunistic resource means that all alternative services need to be able to respond to requests for "http" resources. A client is permitted to use an alternative service without acquiring the http-opportunistic resource from that service.

A client MUST NOT use any cached copies of an http-opportunistic resource that was acquired (or revalidated) over an unauthenticated connection. To avoid potential errors, a client can request or revalidate the http-opportunistic resource before using any connection to an alternative service.

Clients that use cached http-opportunistic responses MUST ensure that their cache is cleared of any responses that were acquired over an unauthenticated connection. Revalidating an unauthenticated response using an authenticated connection does not ensure the integrity of the response.

3. IANA Considerations

This specification registers a Well-Known URI [RFC5785]:

- o URI Suffix: http-opportunistic
- o Change Controller: IETF
- o Specification Document(s): Section 2.3 of [this specification]
- o Related Information:

4. Security Considerations

4.1. Security Indicators

User Agents MUST NOT provide any special security indicators when an "http" resource is acquired using TLS. In particular, indicators that might suggest the same level of security as "https" MUST NOT be used (e.g., a "lock device").

4.2. Downgrade Attacks

A downgrade attack against the negotiation for TLS is possible.

For example, because the "Alt-Svc" header field [RFC7838] likely appears in an unauthenticated and unencrypted channel, it is subject to downgrade by network attackers. In its simplest form, an attacker that wants the connection to remain in the clear need only strip the "Alt-Svc" header field from responses.

4.3. Privacy Considerations

Cached alternative services can be used to track clients over time; e.g., using a user-specific hostname. Clearing the cache reduces the ability of servers to track clients; therefore clients MUST clear cached alternative service information when clearing other origin-based state (i.e., cookies).

4.4. Confusion Regarding Request Scheme

HTTP implementations and applications sometimes use ambient signals to determine if a request is for an "https" resource; for example, they might look for TLS on the stack, or a server port number of 443.

This might be due to expected limitations in the protocol (the most common HTTP/1.1 request form does not carry an explicit indication of the URI scheme and the resource might have been developed assuming HTTP/1.1), or it may be because how the server and application are implemented (often, they are two separate entities, with a variety of possible interfaces between them).

Any security decisions based upon this information could be misled by the deployment of this specification, because it violates the assumption that the use of TLS (or port 443) means that the client is accessing a HTTPS URI, and operating in the security context implied by HTTPS.

Therefore, server implementers and administrators need to carefully examine the use of such signals before deploying this specification.

4.5. Server Controls

This specification requires that a server send both an Alternative Service advertisement and host content in a well-known location to send HTTP requests over TLS. Servers SHOULD take suitable measures to ensure that the content of the well-known resource remains under their control. Likewise, because the Alt-Svc header field is used to describe policies across an entire origin, servers SHOULD NOT permit user content to set or modify the value of this header.

5. References

5.1. 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>>.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, DOI 10.17487/RFC2818, May 2000, <<http://www.rfc-editor.org/info/rfc2818>>.

- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.
- [RFC5785] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", RFC 5785, DOI 10.17487/RFC5785, April 2010, <<http://www.rfc-editor.org/info/rfc5785>>.
- [RFC6454] Barth, A., "The Web Origin Concept", RFC 6454, DOI 10.17487/RFC6454, December 2011, <<http://www.rfc-editor.org/info/rfc6454>>.
- [RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <<http://www.rfc-editor.org/info/rfc7159>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", RFC 7540, DOI 10.17487/RFC7540, May 2015, <<http://www.rfc-editor.org/info/rfc7540>>.
- [RFC7838] Nottingham, M., McManus, P., and J. Reschke, "HTTP Alternative Services", RFC 7838, DOI 10.17487/RFC7838, April 2016, <<http://www.rfc-editor.org/info/rfc7838>>.

5.2. Informative References

- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <<http://www.rfc-editor.org/info/rfc7258>>.

- [RFC7435] Dukhovni, V., "Opportunistic Security: Some Protection Most of the Time", RFC 7435, DOI 10.17487/RFC7435, December 2014, <<http://www.rfc-editor.org/info/rfc7435>>.
- [RFC7469] Evans, C., Palmer, C., and R. Sleevi, "Public Key Pinning Extension for HTTP", RFC 7469, DOI 10.17487/RFC7469, April 2015, <<http://www.rfc-editor.org/info/rfc7469>>.
- [W3C.CR-mixed-content-20160802] West, M., "Mixed Content", World Wide Web Consortium CR CR-mixed-content-20160802, August 2016, <<https://www.w3.org/TR/2016/CR-mixed-content-20160802>>.

Appendix A. Acknowledgements

Mike Bishop contributed significant text to this document.

Thanks to Patrick McManus, Stefan Eissing, Eliot Lear, Stephen Farrell, Guy Podjarny, Stephen Ludin, Erik Nygren, Paul Hoffman, Adam Langley, Eric Rescorla, Julian Reschke, KariHurtta, and Richard Barnes for their feedback and suggestions.

Authors' Addresses

Mark Nottingham

Email: mnot@mnot.net

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

Martin Thomson
Mozilla

Email: martin.thomson@gmail.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 4, 2015

S. Loreto, Ed.
J. Mattsson
R. Skog
H. Spaak
Ericsson
G. Bourg
D. Druta
M. Hafeez
AT&T
July 3, 2014

Explicitly Authenticated Proxy in HTTP/2.0
draft-loreto-httpbis-explicitly-auth-proxy-01

Abstract

This document proposes the definition of an Explicitly Authenticated Proxy as intermediary of normally unprotected "http" URI scheme requests and responses of HTTP2 traffic.

An Explicitly Authenticated Proxy is a message forwarding agent that is selected, with explicit user's consent, and configured by the user agent to receive exclusively "http" URI scheme requests and attempt to satisfy those requests on behalf of the user agent. A client is connected to an Explicitly Authenticated Proxy through an authenticated TLS secured connection.

This document describes a method for a user agent to automatically discover and authenticate, and for an user to provide consent for an Explicitly Authenticated Proxy. This enables proxied communication to be encrypted and authenticated, explicitly acknowledged by the user agent and visible to the server end point.

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 January 4, 2015.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1. Introduction	3
1.1. Goals and non Goals	4
1.2. Explicitly Authenticated Proxy	4
2. Terminology	5
3. Establishing proxy connection	5
3.1. TLS Handshake with Proxy certificate	5
4. Connection to a mobile network	6
4.1. proxy discovery in a mobile network	7
5. Explicit Proxy behaviour	7
5.1. Explicitly Authenticated Forward Proxy towards HTTP2 origin server	7
5.2. Explicitly Authenticated Forward Proxy towards HTTP/1.1 Origin Server	9
5.3. Explicitly Authenticated Forward Proxy and https URIs . .	10
6. User Consent	11
6.1. Expected behaviour if the user opts out/revokes consent .	11
7. Signalling the presence of a Proxy in between	12
8. Security Considerations	12
9. Acknowledgments	14
10. References	14
10.1. Normative References	14
10.2. Informative References	15
10.3. URIs	15
Appendix A. Proxy certificate	15
Authors' Addresses	16

1. Introduction

HTTP/1.1 and earlier allowed for the use of proxies and gateways to satisfy requests through a chain of connections. This has made possible a Web ecosystem of various kinds of proxies and gateways: cache servers, security gateways, web accelerators, content filters, and many others. In some cases their presence is explicit (configured proxies), and in other they are completely transparent to the end user (interception proxies, and gateways such as reverse proxies).

The success and the presence of the proxies and gateways is also a problem for the evolution of the HTTP as their behaviour on protocol extensions, and especially on alternative wire formats of the protocol, is not predictable. This unpredictable behaviour can lead to difficulties to deploy new versions of the protocol before the intermediaries are themselves updated. As an example, see the difficulties in deploying the WebSocket Protocol [RFC6455] in clear. It can also lead to potentially problematic trust models where proxies are accessing traffic content without the user being aware. Relying on establishing an HTTPS tunnel has then become the popular way to bypass the intermediate proxies as it provides reliable deployment model for web protocols. The encrypted tunnel obfuscates the data from all intermediaries and provides integrity validation.

HTTPS tunnels, while speeding up the deployment, make it difficult for a forward proxy and other gateways to be used to enable caching, enhance anonymity for a user agent, or enhance security by scanning content for virus and malware. HTTPS tunnels also remove the possibility to enhance delivery performance based on the knowledge of the network status, and this become an important limitation especially with HTTP2 when multiple streams are multiplexed on top of the same TCP connection.

Several drafts analysing the role and the requirements for proxy have been submitted:

1. [I-D.nottingham-http-proxy-problem] discusses the use and configuration of proxies in HTTP, pointing out problems in the currently deployed Web infrastructure along the way
2. [I-D.vidya-httpbis-explicit-proxy-ps] describes the issues with HTTP proxies for TLS protected traffic and motivates the need for explicit proxying capability in HTTP. It also presents the goals that such a solution would need to satisfy and some example solution directions.

3. [I-D.rpeon-httpbis-exproxy] describes a method for connecting to a proxy via a secure channel, allowing, disallowing, and detecting any transforms that the proxy may perform, and allowing the proxy to connect via secure channel to another site on the user's behalf.

Use cases in form of stories for proxies are also listed in the wiki Proxy-User-Stories [1] and analysed in a matrix form in Trusted Proxy Use Case Analysis and Alternatives [2].

This draft explicitly narrows down the general discussion to the role of Proxy as intermediary of "http" scheme URIs of HTTP2 traffic.

1.1. Goals and non Goals

The primary goal is to define an intermediary to 'http' traffic, that is TLS connected to the browser, operates with the knowledge and explicit consent of the user.

Non goal is to define an intermediary for 'https' URI. However the intermediary's expected behaviour for this case is listed for completeness.

1.2. Explicitly Authenticated Proxy

An "Explicitly Authenticated", as defined in this document, is an HTTP Proxy (see section 2.3 [I-D.ietf-httpbis-pl-messaging]) that is certificate authenticated, user acknowledged and connected to over a TLS encrypted (and possibly integrity protected) connection. An Explicitly Authenticated Proxy is configured by the user agent to exclusively receive "http" URI scheme requests and attempt to satisfy those requests on behalf of the user agent.

The presence of a configured Explicitly Authenticated Proxy MUST NOT change the user agent behaviour for the "https" URI scheme requests.

To distinguish between an HTTP2 connection meant to transport "https" URIs resources and an HTTP2 connection meant to transport "http" URIs resource, this document defines the ALPN [I-D.ietf-tls-appplayerprotoneg] identifier "h2c" to signal that HTTP2 transports "http" URI requests and resources over TLS.

This document describes a method for an user agent to automatically discover and then for an user to accept or reject (i.e. to provide consent for) an Explicitly Authenticated Proxy to be securely involved when a request to an "http" URI resource is made.

Section 3 defines a solution based on sending a proxy certificate in the TLS handshake.

Section 5 describes the role of the Explicitly Authenticated Proxy in helping the user to fetch "http" URIs resource when the user has provided consent to the Explicitly Authenticated Proxy to be involved.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

This document defines the following terms:

Explicit proxy: an intercepting proxy (see section 2.3 [I-D.ietf-httpbis-pl-messaging]) that communicates its presence to the user agent and destination server..

Explicitly Authenticated Proxy: an HTTP Proxy that is certificate authenticated, user acknowledged and connected to over a TLS encrypted (and possibly integrity protected) connection. An Explicitly Authenticated Proxy is configured by the user agent to exclusively receive "http" URI scheme requests and attempt to satisfy those requests on behalf of the user agent. The presence of a configured Explicitly Authenticated Proxy MUST NOT change the user agent behaviour for the "https" URI scheme requests.

3. Establishing proxy connection

An Explicitly Authenticated Proxy indicates its presence, identity and willingness to serve the user agent by intercepting TLS ClientHello message containing "h2c" value (a new ALPN protocol type assigned for this purpose) in the ALPN [I-D.ietf-tls-applayerprotoneg] negotiation extension field. It answers the TLS initiation with a TLS ServerHello message containing the Proxy certificate Appendix A .

3.1. TLS Handshake with Proxy certificate

When a (TLS and HTTP) user agent receives a Server Certificate message, it checks whether the certificate contains an Extended Key Usage extension and if so whether the "proxyAuthentication" key purpose id is included. If it is included, the user agent concludes that the certificate belongs to a proxy. The user agent then SHOULD ensure user consent.

If the user provides consent, the user agent continues the TLS handshake with the proxy.

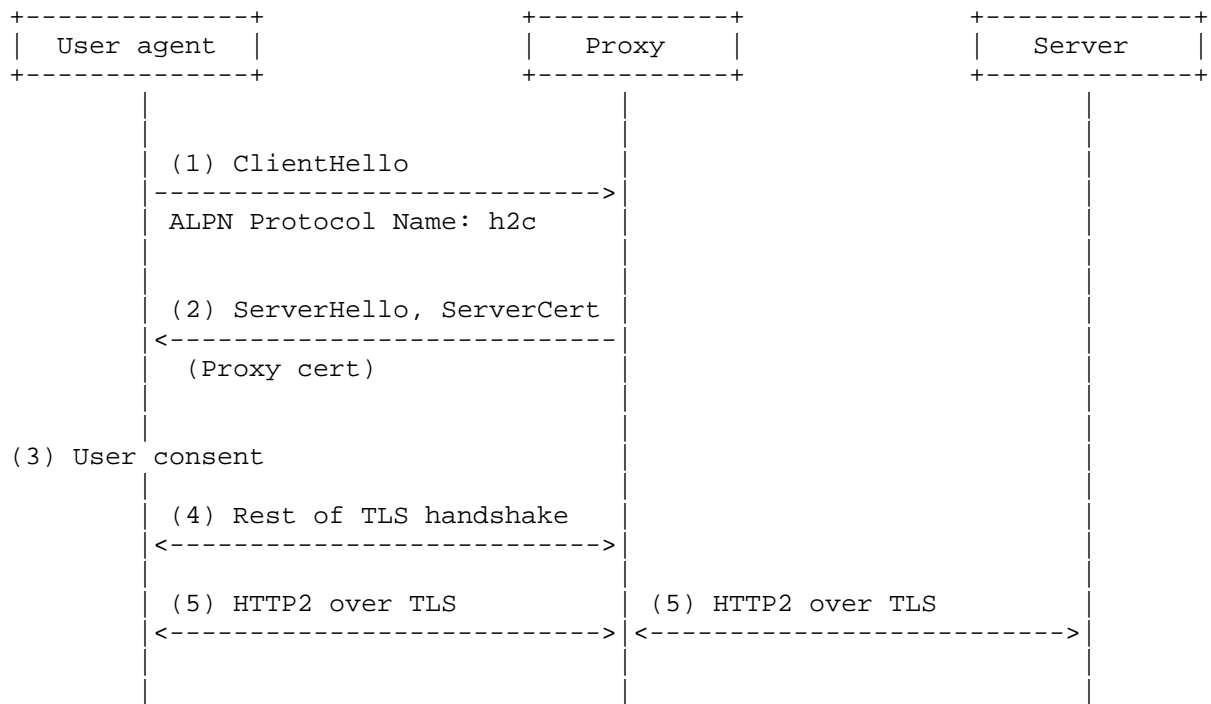


Figure 1: TLS Handshake with Proxy certificate

4. Connection to a mobile network

When a handset connects to a mobile network it is desirable to preserve the integrity of its exchange with the servers which host the services of this network entity. These use cases are described in [I-D.nottingham-http-proxy-problem] and in the [Proxy-User-Stories].

This section proposes a solution for such use cases. The proposal is inspired on the connection management specified in the section 9.1 of [I-D.ietf-httpbis-http2]. The connection with this proxy is used for all the servers' names listed in the "subjectAltName" field (<http://tools.ietf.org/html/rfc5280#page-35>) of the certificate of this proxy.

4.1. proxy discovery in a mobile network

At the network attachment, as usual, the network entity provides the handset with an IP address and with other pieces of information like DNS resolvers IP addresses. The network entity additionally provides the handset with the server name (e.g. pr.example.com) of the Explicitly Authenticated Proxy in charge of the domain names this network entity is authoritative on. These pieces of information are provided to the handset through a secure channel which preserves the integrity of the information.

5. Explicit Proxy behaviour

This section describes the role of the Explicitly Authenticated Proxy in helping the user to fetch http URI resources when the user has provided consent to the Explicitly Authenticated Proxy to be involved.

5.1. Explicitly Authenticated Forward Proxy towards HTTP2 origin server

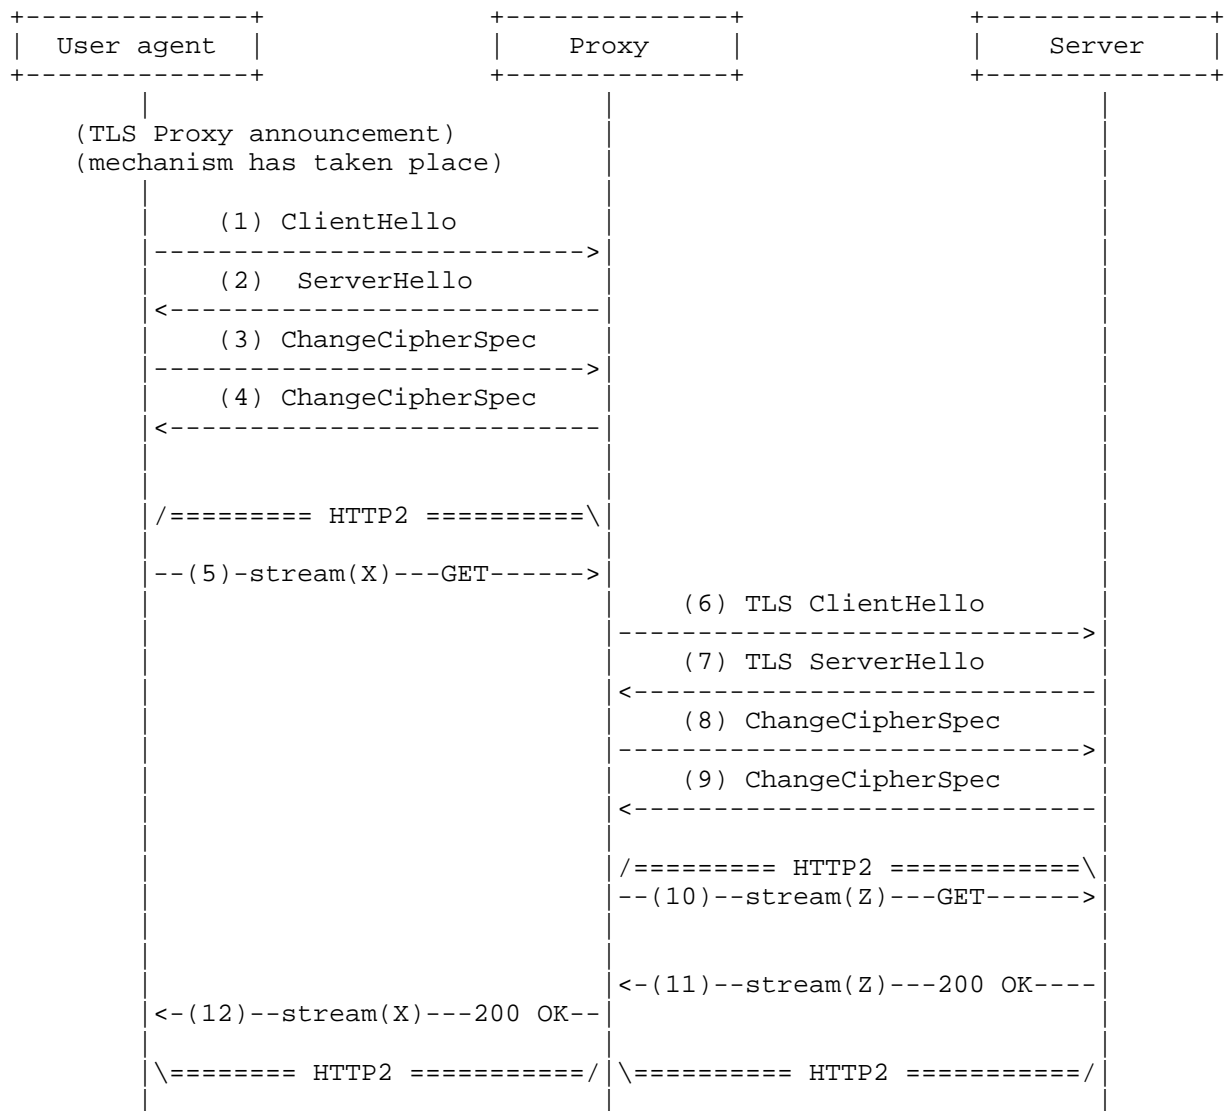


Figure 2: Requesting an HTTP resource

- (0) The TLS Proxy Announcement (Section 3) mechanism has already taken place, so the user agent is now configured in the proxy mode.
- (1)...(4) For each "http" URI resource towards a not yet contacted Server Origin, the user agent negotiates a new TLS session, using

the ALPN extension containing the "h2c" tag, to establish an HTTP2 connection.

(5) The user agent will then use the streams in the HTTP2 connection to request any resources hosted on that Origin Server.

(6)...(9) In the case the Proxy receives a request for a resource towards a not yet contacted Server Origin, the Explicitly Authenticated Proxy negotiates a new TLS session, using the ALPN extension containing the "h2c" ALPN identifier, to establish an HTTP2 connection.

(10) Once the Proxy has established the HTTP2 connection toward the origin, it picks one stream to forward the request

(11), (12) The Proxy forwards the answer it receives from the Origin Server to the user agent.

5.2. Explicitly Authenticated Forward Proxy towards HTTP/1.1 Origin Server

In the case the proxy has a privies knowledge about the fact that the "http" URI resources requested by the user agent will be only available over HTTP/1.1 or the proxy does not have a previous knowledge about it, the proxy will then attempt to contact the resource based on its knowledge.

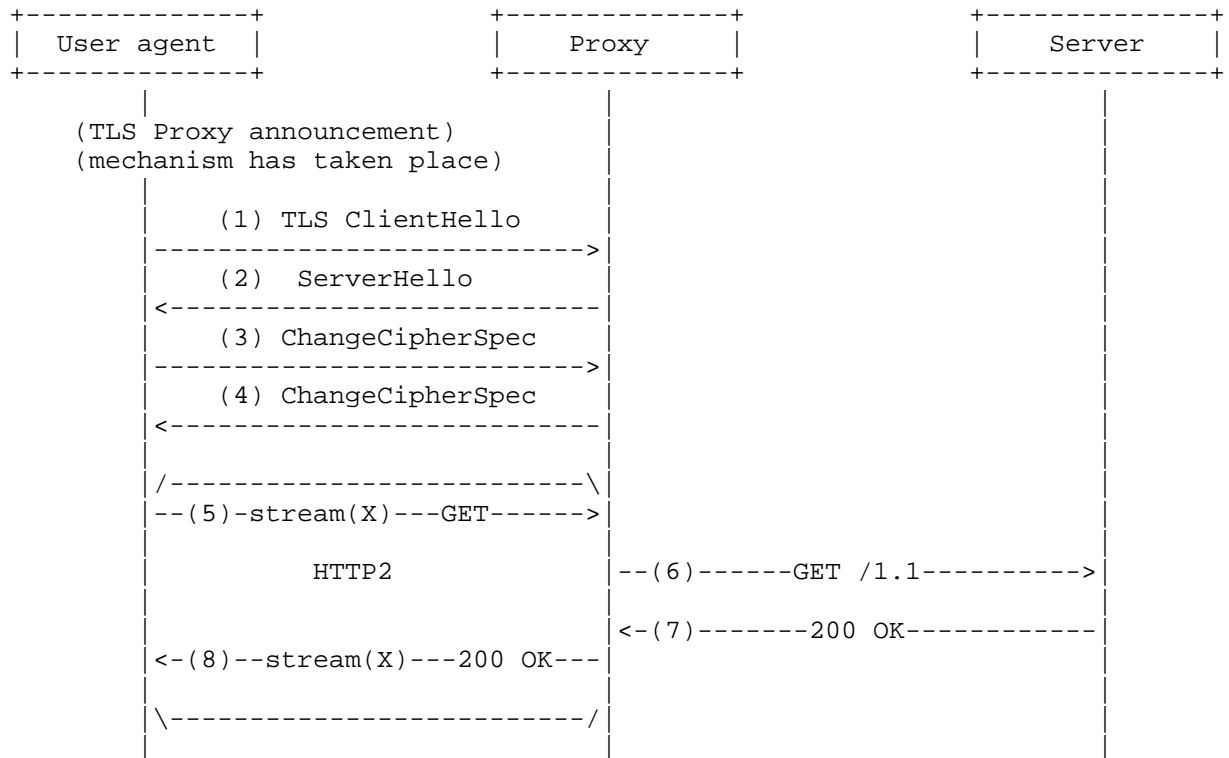


Figure 3: Origin server with only HTTP/1.1 support

5.3. Explicitly Authenticated Forward Proxy and https URIs

A user agent MUST NOT use "h2c" as ALPN extension field in request for https resources.

The Proxy that intercepts the TLS ClientHello analyses the ALPN extension field and if it does not contain the "h2c" value it does not do anything and lets the TLS handshake continue and the TLS session be established between the user agent and the Server (see Figure 4).

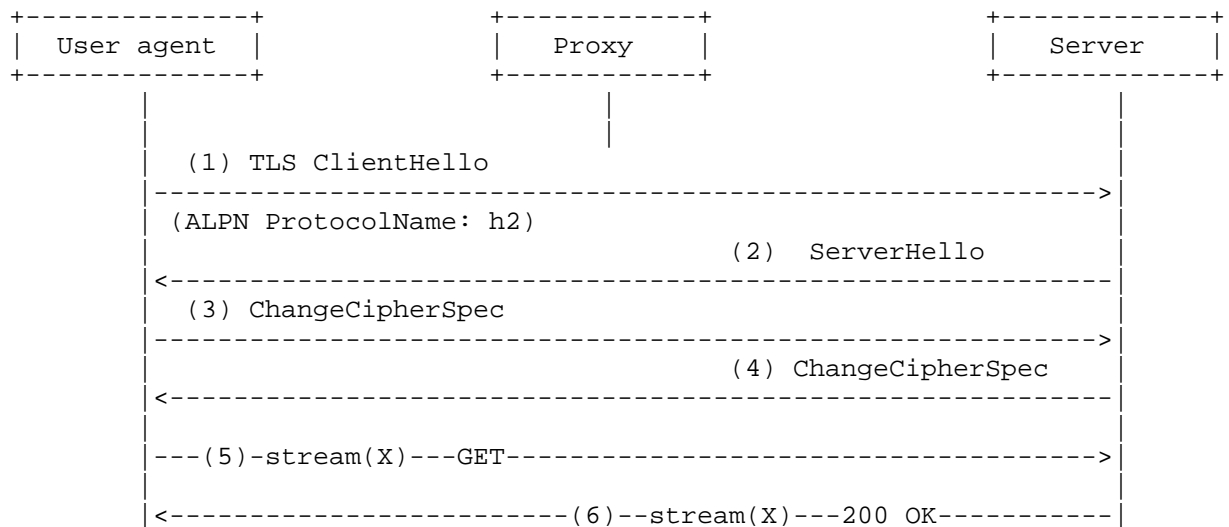


Figure 4: Explicitly Authenticated Proxy and https URI resources

6. User Consent

This document proposes an approach to making the presence of proxy explicit, explaining the functions it provides to users and letting them decide whether they accept that. A user can opt out and choose to bypass the proxy. This ensures that a proxy never acts as intermediary for HTTP2 traffic unless authorised by the user.

The user selection can be cached by the user agent. A consent SHOULD however be limited to the specific network access (such as APN or SSID) and may be limited to a single connection to that access or limited in time. How the consent information is stored is implementation specific, but as a network may have several proxies (for network resilience) it is RECOMMENDED that the consent is only tied to the Subject field of the proxy certificate so that the consent applies to all proxy certificates with the same name.

6.1. Expected behaviour if the user opts out/revokes consent

If the user does not give consent, or decides to opt out from the proxy for a specific connection, the user agent will negotiate HTTP2 connection using "h2" value in the ALPN extension field. The proxy will then treat the connection as an "https" connection and will forward the ClientHello message to the Server, establishing an end-to-end TLS connection between the user agent and the destination server.

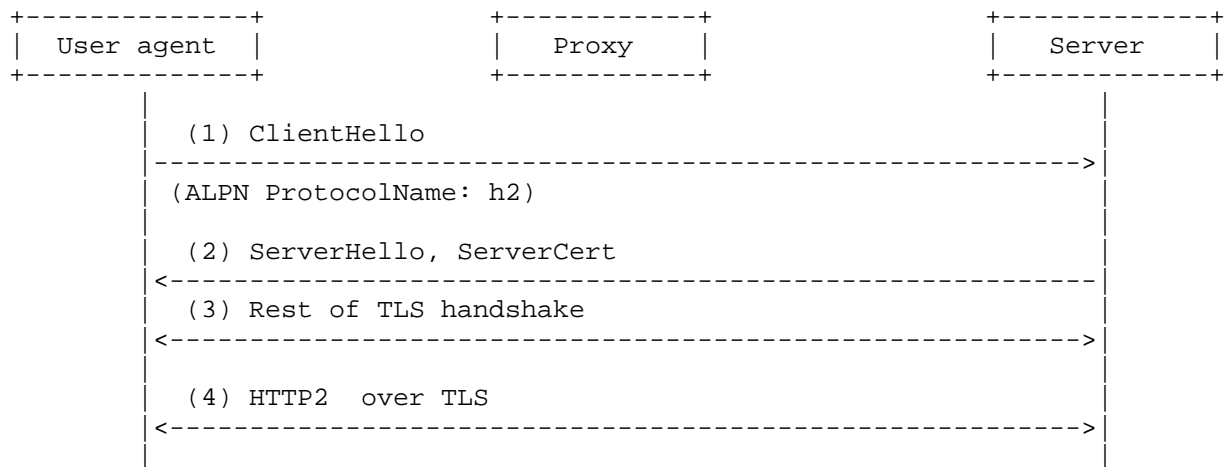


Figure 5: Opt Out

7. Signalling the presence of a Proxy in between

The presence of Explicitly Authenticated Proxy in between an user agent and the origin server must be signalled to the origin server using an already defined HTTP header.

The Explicitly Authenticated proxy MUST add, or update when already present, the Forwarded HTTP header field [I-D.ietf-appsawg-http-forwarded] "for" parameter.

8. Security Considerations

This document addresses Explicitly Authenticated proxies that act as intermediary for HTTP2 traffic and therefore the security and privacy implications of having those proxies in the path need to be considered. MITM [3], [I-D.nottingham-http-proxy-problem] and [I-D.vidya-httpbis-explicit-proxy-ps] discuss various security and privacy issues associated with the use of proxies.

It should however be noticed that the presence of the Explicitly Authenticated proxy as discussed in this document does not in any way affect "https" URI resources. Those resources are protected end-to-end between user agent and origin server as usual. Only for "http" URI resources the achievable security level of hop-by-hop protection may be different than end-to-end protection, because it is now also dependent on the security features/capabilities of the proxy as to what cipher suites it supports, which root CA certificates it trusts, how it checks certificate revocation status, etc. Users should also

be made aware that the proxy has visibility to the actual content they exchange with Web servers, including personal and sensitive information.

The TLS connection from the user agent to the Explicitly Authenticated proxy is always authenticated. In case the origin server only offers unauthenticated TLS (e.g. by using a self-signed certificate) the explicit Explicitly Authenticated proxy increases the security in the access network (e.g. an unencrypted hotspot) by ensuring that there is no unwanted MITMs in this part of the network.

To ensure the trustfulness of proxies, certification authorities validation procedure for issuing proxy certificates should be more rigorous than for issuing normal certificates and may also include technical details and processes relevant for the security assurance. The owner of the proxy could for example be obliged to apply security patches in a timely fashion.

When negotiating ciphersuite with the server, the Explicitly Authenticated proxy SHALL offer the ciphersuite negotiated between the user-agent and the proxy. Ciphersuites with a higher security level than the ciphersuite negotiated between the user-agent and proxy MAY be given a higher preference than the ciphersuite negotiated between the user-agent and proxy. Ciphersuites with a lower security level than the ciphersuite negotiated between the user-agent and proxy SHALL NOT be given a higher preference than the ciphersuite negotiated between the user-agent and proxy. While AES-256 is no weaker (and most probably much stronger) than AES-128, the relative security between different algorithms e.g. SHA-256 vs Keccak-256 is not that clear. With security level we mean the complexity of the best known attack on that ciphersuite. The Explicitly Authenticated proxy SHOULD therefore be up to date with the best current practices regarding TLS.

This document proposes an approach to making the presence of proxy explicit to users and letting them decide whether they accept that. A user can opt out and choose to bypass the proxy. This ensures that a proxy never acts as intermediary for HTTP2 traffic unless authorised by the user.

When the user has given consent to the presence of the proxy, the user agent switches to a Proxy mode in which it does not check the hostname of the origin server against the server's identity as presented in the Server Certificate message. However if any of the following checks fails the user agent should immediately exit this Proxy mode:

1. the server's certificate is issued by a trusted CA and the certificate is valid;
2. the Extended Key Usage extension is present in the certificate and indicates the owner of this certificate is a proxy;
3. the server possesses the private key corresponding to the certificate.

9. Acknowledgments

The authors wish to thank Yi Cheng, Goran Eriksson, Stefan Hakansson, Nicolas Mailhot, Martin Nilsson, Emile Stephan (Connection with prior knowledge) and Salman Taj for their ideas, technical suggestions and comments.

10. References

10.1. Normative References

- [I-D.ietf-appsawg-http-forwarded]
Petersson, A. and M. Nilsson, "Forwarded HTTP Extension", draft-ietf-appsawg-http-forwarded-10 (work in progress), October 2012.
- [I-D.ietf-httpbis-http2]
Belshe, M., Peon, R., and M. Thomson, "Hypertext Transfer Protocol version 2", draft-ietf-httpbis-http2-13 (work in progress), June 2014.
- [I-D.ietf-httpbis-pl-messaging]
Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", draft-ietf-httpbis-pl-messaging-26 (work in progress), February 2014.
- [I-D.ietf-tls-applayerprotoneg]
Friedl, S., Popov, A., Langley, A., and S. Emile, "Transport Layer Security (TLS) Application Layer Protocol Negotiation Extension", draft-ietf-tls-applayerprotoneg-05 (work in progress), March 2014.
- [I-D.nottingham-http-proxy-problem]
Nottingham, M., "Problems with Proxies in HTTP", draft-nottingham-http-proxy-problem-00 (work in progress), October 2013.

- [I-D.rpeon-httpbis-exproxy]
Peon, R., "Explicit Proxies for HTTP/2.0", draft-rpeon-httpbis-exproxy-00 (work in progress), June 2012.
- [I-D.vidya-httpbis-explicit-proxy-ps]
Narayanan, V., "Explicit Proxying in HTTP - Problem Statement And Goals", draft-vidya-httpbis-explicit-proxy-ps-00 (work in progress), October 2013.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3709] Santesson, S., Housley, R., and T. Freeman, "Internet X.509 Public Key Infrastructure: Logotypes in X.509 Certificates", RFC 3709, February 2004.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.

10.2. Informative References

- [RFC6455] Fette, I. and A. Melnikov, "The WebSocket Protocol", RFC 6455, December 2011.

10.3. URIs

- [1] <https://github.com/http2/http2-spec/wiki/Proxy-User-Stories>
- [2] <https://github.com/bizzbyster/TrustedProxy/wiki/Trusted-Proxy-Use-Case-Analysis-and-Alternatives>
- [3] Jarmoc, J., SSL/TLS Interception Proxies and Transitive Trust, 2012 https://www.grc.com/miscfiles/HTTPS_Interception_Proxies.pdf

Appendix A. Proxy certificate

To help HTTP user agents identify and distinguish Explicitly Authenticated proxies from other servers (e.g. web servers), Explicitly Authenticated proxies should have a certification authority issued public key certificate.

More specifically, the certification authority SHOULD use the Extended Key Usage extension as specified in [RFC5280] to indicate a key purpose "proxyAuthentication" (a new object identifier needs to be assigned by IANA for this key purpose). The certification authority also marks this Extended Key Usage extension as critical.

As the user needs to have high trust in the Proxy, it is desirable that the validation procedure for issuing proxy certificates be more rigorous than for issuing ordinary SSL certificates.

A proxy certificate MUST contain the SubjectAltName extension as defined in [RFC5280]. A name identifying the legal entity that is operating the proxy should be given in this extension.

To help end users understand the reason why the proxy is offered (in other words, the benefits of having the proxy in the path), a new X.509 certificate extension ProxyFunctions is introduced to list the functions the proxy is performing. More specifically, the ProxyFunction extension consists of a sequence of ProxyFunctionId which are object identifiers. The user agent should check the presence of this extension in the proxy certificate and present the proxy functions in a human readable format.

The user agent will provide the user with an opportunity to graphically view the results of a successful proxy certificate-based identification process leveraging on the usage of logotypes in public key certificates and attribute certificates as specified in [RFC3709].

Authors' Addresses

Salvatore Loreto (editor)
Ericsson
Hirsalantie 11
Jorvas 02420
Finland

Email: salvatore.loreto@ericsson.com

John Mattsson
Ericsson
Kista
Sweden

Email: john.mattsson@ericsson.com

Robert Skog
Ericsson
Kista
Sweden

Email: robert.skog@ericsson.com

Hans Spaak
Ericsson
Kista
Sweden

Email: hans.spaak@ericsson.com

Gus Bourg
AT&T

Email: gb3635@att.com

Dan Druta
AT&T

Email: dd5826@att.com

Mohammad Hafeez
AT&T

Email: mh2897@att.com

HTTPbis
Internet-Draft
Intended status: Informational
Expires: January 6, 2015

H. Nakajima
Keio University, W3C
July 5, 2014

HTTP/2 Interoperability Survey
draft-nakajima-httpbis-http2-interop-survey-00

Abstract

This document provides a survey of HTTP/2 [I-D.ietf-httpbis-http2] and HPACK [I-D.ietf-httpbis-header-compression] implementations and interoperability tests based on HTTP/2.0 Testing [I-D.trace-httpbis-http2-test]. Goals of this document are to help improving HTTP/2 specifications and HTTP/2 implementations and deployments.

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 January 6, 2015.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1. Requirements	2
2. Interoperability Test Target	2
3. Interoperability Test Area	3
3.1. Connect	3
3.2. Upgrade	3
3.3. Framing	3
3.4. Flow Control	3
3.5. Streams and Multiplexing	3
3.6. Header Compression	3
3.7. Connection Management	3
3.8. Stream Prioritization	3
3.9. Authentication	4
3.10. Server Push	4
3.11. TLS Negotiation	4
3.12. TLS Cipher Suite	4
3.13. Opportunistic Encryption	4
3.14. Alternative Services	4
4. Interoperability Test Results	4
5. Implementation Survey	4
6. Security Considerations	5
7. IANA Considerations	5
8. Contributors	5
9. References	5
9.1. Normative References	5
9.2. Informative References	5
9.3. URL References	5

1. Requirements

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

2. Interoperability Test Target

The following implementaions have been tested for interoperability test. Each implementation has been classified according to the roles "client", "server", "intermediary" which are defined in [RFC7230]. Table 1 shows name and roles of tested implementations.

Name	Roles
[nghttp2]	client,server,intermediary
[http2-katana]	client,server
[node-http2]	client,server
[MozillaFirefox]	client
[iiij-http2]	client
[AkamaiGhost]	intermediary
[Chromium]	client
[Twitter]	client,server
[http2-go]	client,server
[OkHttp2]	client
[mruby-http2]	client,server
[curl]	client
[cl-http2-protocol]	client,server
[Netty]	client,server
[Jetty]	client,server

Data from <https://github.com/http2/http2-spec/wiki/Implementations>

Table 1: Surveyed Implementations

3. Interoperability Test Area

3.1. Connect

Basic Connection Test is checking if the client and server(or intermediary) are able to establish HTTP/2 connection.

3.2. Upgrade

3.3. Framing

3.4. Flow Control

3.5. Streams and Multiplexing

3.6. Header Compression

3.7. Connection Management

3.8. Stream Prioritization

3.9. Authentication

3.10. Server Push

3.11. TLS Negotiation

3.12. TLS Cipher Suite

3.13. Opportunistic Encryption

3.14. Alternative Services

4. Interoperability Test Results

Table 2 shows interoperability test result.

	ngh ttp 2	http2 -kata na	node- http2	iiij- http2	Twit ter	mruby- http2	cl-http2 -protoco l
[nghhttp2]	\	x	x	o	o	o	o
[http2-katana]	x	\	x	x	x	x	x
[node-http2]	x	x	\	x	x	x	x
[Mozilla Firefox]	o	x	x	o	o	x	x
[Chromium]	o	x	x	o	o	x	o
[cl-http2-protocol]	o	x	x	x	x	o	\

Vertical axis: Client, Horizontal: Server

Table 2: Interoperability Test Result - Connection

5. Implementation Survey

TBD.

6. Security Considerations

TBD.

7. IANA Considerations

This document makes no request to IANA.

8. Contributors

9. References

9.1. Normative References

[I-D.ietf-httpbis-header-compression]

Peon, R. and H. Ruellan, "HPACK - Header Compression for HTTP/2", draft-ietf-httpbis-header-compression-08 (work in progress), June 2014.

[I-D.ietf-httpbis-http2]

Belshe, M., Peon, R., and M. Thomson, "Hypertext Transfer Protocol version 2", draft-ietf-httpbis-http2-13 (work in progress), June 2014.

[I-D.trace-httpbis-http2-test]

Lai, M., Jian, C., and R. Trace, "HTTP/2.0 Protocol Test", draft-trace-httpbis-http2-test-00 (work in progress), September 2013.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. Informative References

[RFC7230] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, June 2014.

9.3. URL References

[AkamaiGhost]

Akamai Technologies, Inc., "<https://github.com/http2/http2-spec/wiki/Akamaighost>", .

[Chromium]

The Chromium Projects,
"<https://sites.google.com/a/chromium.org/dev/spdy/http2>",
.

- [Jetty] The Eclipse Foundation, "<http://git.eclipse.org/c/jetty/org.eclipse.jetty.project.git/tree/?h=jetty-http2>", .
- [MozillaFirefox]
Mozilla Foundation, "<https://wiki.mozilla.org/Networking/http2>", .
- [Netty] The Netty project, "<http://netty.io/>", .
- [OkHttp2] Square, Inc., "<https://github.com/square/okhttp>", .
- [Twitter] Twitter Inc., "<https://twitter.com>", .
- [cl-http2-protocol]
Akamai Technologies, Inc., "<https://github.com/akamai/cl-http2-protocol>", .
- [curl] Stenberg, D., "<http://curl.haxx.se/>", .
- [http2-go]
"<https://github.com/Jxck/http2>", .
- [http2-katana]
Microsoft Open Technologies, Inc.,
"<https://github.com/Microsoft/http2-katana>", .
- [iij-http2]
Internet Initiative Japan Inc.,
"<https://github.com/shigeki/interop-iij-http2>", .
- [mruby-http2]
Matsumoto, R., "<https://github.com/matsumoto-r/mruby-http2>", .
- [nghttp2] Tsujikawa, T., "<https://nghttp2.org/>", .
- [node-http2]
Hurley, N., Belshe, M., and Y. Iwanaga,
"<https://github.com/molnarg/node-http2>", .

Author's Address

Hiroataka Nakajima
Keio University, W3C
5322 Endo
Fujisawa, Kanagawa
Japan

Phone: +81.466.49.3424
EMail: hiro@awa.sfc.keio.ac.jp

Network Working Group
Internet-Draft
Updates: 5789 (if approved)
Intended status: Informational
Expires: September 15, 2014

M. Nottingham
March 14, 2014

The 2NN Patch HTTP Status Code
draft-nottingham-http-patch-status-00

Abstract

This document specifies the 2NN Patch HTTP status code to allow servers to perform partial updates of stored responses in client caches.

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 September 15, 2014.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1. Introduction	3
1.1. Notational Conventions	3
2. The 2NN Patch Status Code	3
2.1. The Patched Header Field	5
3. IANA Considerations	5
3.1. 2NN Patch HTTP Status Code	5
3.2. Accept-Patch Header Field	5
3.3. Patched Header Field	5
4. Security Considerations	5
5. References	6
5.1. Normative References	6
5.2. Informative References	6
Appendix A. 2NN Patch and HTTP/2 Server Push	6
Author's Address	8

1. Introduction

[RFC5246] defines the HTTP PATCH method as a means of selectively updating the state of a resource on a server. This document complements that specification by specifying a means for a server to selectively update a stored response on a client - usually, in a cache [I-D.ietf-httpbis-p6-cache].

1.1. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

This document uses the Augmented BNF defined by [RFC5246], and additionally uses the entity-tag rule defined in [I-D.ietf-httpbis-p4-conditional].

2. The 2NN Patch Status Code

The 2NN (Patch) status code allows a response to patch a stored response in a cache, by reusing the patch formats of [RFC5789]. In some sense, it is the complement of the HTTP PATCH request method.

TODO: is this a 2NN or 3xx?

Clients can advertise support for 2NN (Patch), along with the patch formats supported in it, by using the Accept-Patch header field in requests. For example:

```
GET /foo HTTP/1.1
Host: api.example.com
Accept-Patch: application/patch+json
If-None-Match: "abcdef", "ghijkl"
User-Agent: Example/1.0
```

If the server can generate a patch for one of the entity tags provided in If-None-Match, in one of the accepted patch formats, it can generate a 2NN (Patch) response:

```
HTTP/1.1 2NN Patch
Content-Type: application/patch+json
Patched: "ghijkl"
ETag: "mnopqrs"
```

The entity tag carried by the ETag header field is associated with the selected representation - i.e., the stored response after the

patch is applied.

The Patched header field identifies the representation to apply the patch to, as indicated by the entity-tag provided in If-None-Match request header field; see Section 2.1.

Therefore, in the example above, the stored response "ghijkl" is being patched, with the resulting stored response having the entity tag "mnopqrs".

Application of a 2NN (Patch) response happens in a manner very similar to the process for freshening a stored response by applying a 304 (Not Modified), as described in [I-D.ietf-httpbis-p6-cache], Section 4.3.4.

In particular, the stored response to apply a 2NN (Patch) response to is the same; if none is selected, the patch fails, and the client MAY resubmit the request without an Accept-Patch header field, in order to get a full response.

If a stored response is selected, clients MUST update it in the following manner:

- o The value of the Content-Length header field MUST be adjusted to reflect the length of the patched response body.
- o The ETag header field MUST be replaced (or inserted, if not present) with the value of the Patched header field in the 2NN response (if present).
- o Other header fields in the 2NN response MUST update the stored response, in the same manner as described in [I-D.ietf-httpbis-p6-cache], Section 4.3.4. However, the following fields MUST not be updated: Content-Type, Patched.

The 2NN (Patch) status code SHOULD NOT be generated if the request did not include If-None-Match, unless conflicts are handled by the patch format itself (e.g., allowing a patch to append to an array), or externally.

Intermediaries MAY append the Accept-Patch header field to requests, or append new values to it, if they will process 2NN responses for the patch format(s) they add. Likewise, intermediaries MAY generate 2NN (Patch) responses under the conditions specified here.

The 2NN status code is not cacheable by default, and is not a representation of any identified resource.

2.1. The Patched Header Field

The Patched header field identifies the stored representation that a patch is to be applied to in a 2NN (Patch) response.

Patched = entity-tag

3. IANA Considerations

3.1. 2NN Patch HTTP Status Code

This document defines the 2NN (Patch) HTTP status code, as per [I-D.ietf-httpbis-p2-semantics].

- o Status Code (3 digits): TBD
- o Short Description: Patch
- o Pointer to specification text: Section 2

3.2. Accept-Patch Header Field

This document updates [RFC5789] to allow the Accept-Patch HTTP header field to be used in requests, which ought to be reflected in the registry.

3.3. Patched Header Field

This document defines a new HTTP header, field, "Patched", to be registered in the Permanent Message Header Registry, as per [RFC3864].

- o Header field name: Patched
- o Applicable protocol: http
- o Status: standard
- o Author/Change controller: IETF
- o Specification document(s): [this document]
- o Related information:

4. Security Considerations

2NN (Patch) can be brittle when the application of a patch fails, because the client has no way to report the failure of a patch to the server. This asymmetry might be exploited by an attacker, but can be mitigated by judicious use of strong ETags.

Some patch formats might have additional security considerations.

5. References

5.1. Normative References

- [I-D.ietf-httpbis-p4-conditional]
Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests",
draft-ietf-httpbis-p4-conditional-26 (work in progress),
February 2014.
- [I-D.ietf-httpbis-p6-cache]
Fielding, R., Nottingham, M., and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Caching",
draft-ietf-httpbis-p6-cache-26 (work in progress),
February 2014.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC5789] Dusseault, L. and J. Snell, "PATCH Method for HTTP", RFC 5789, March 2010.

5.2. Informative References

- [I-D.ietf-httpbis-http2]
Belshe, M., Peon, R., and M. Thomson, "Hypertext Transfer Protocol version 2", draft-ietf-httpbis-http2-10 (work in progress), February 2014.
- [I-D.ietf-httpbis-p2-semantics]
Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content",
draft-ietf-httpbis-p2-semantics-26 (work in progress),
February 2014.
- [RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", BCP 90, RFC 3864, September 2004.

Appendix A. 2NN Patch and HTTP/2 Server Push

In HTTP/2 [I-D.ietf-httpbis-http2], it is possible to "push" a request/response pair into a client's cache. 2NN (Patch) can be used with this mechanism to perform partial updates on stored responses.

For example, if a cache has this response stored for "http://example.com/list":

```
200 OK
Content-Type: application/json
Cache-Control: max-age=3600
ETag: "aaa"

{ "items": ["a"]}
```

A HTTP/2 server could partially update it by sending the request/response pair (using pseudo-HTTP/1 syntax for purposes of illustration):

```
GET /list
Host: example.com
If-None-Match: "aaa"
Accept-Patch: application/patch+json

2NN Patch
Content-Type: application/patch+json
ETag: "aab"
Patched: "aaa"

[
  { "op": "add", "path": "/items/1", "value": "b" }
]
```

Once the patch is applied, the stored response is now:

```
200 OK
Content-Type: application/json
Cache-Control: max-age=3600
ETag: "aab"

{ "items": ["a", "b"]}
```

Note that this approach requires a server pushing partial responses to know the stored response's ETag, since the client cache will silently ignore the push if it does not match that provided in "Patched". Likewise, clients that are not conformant to this specification will silently drop such pushes, since the status code is not recognised (as per [I-D.ietf-httpbis-p6-cache]).

However, it is possible to do some partial updates without strong consistency. For example, if the stored response is as above, and the server simply wishes to append a value to an array, without regard for the current content of the array (because, presumably,

ordering of its content is not important), it can push:

```
GET /list
Host: example.com
Accept-Patch: application/patch+json

2NN Patch
Content-Type: application/patch+json

[
  { "op": "add", "path": "/items/-", "value": "b" }
]
```

Here, the resulting document would be as above, but since entity tags are not provided, the operation will succeed as long as the patch application succeeds.

Author's Address

Mark Nottingham

Email: mnot@mnot.net

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

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 05, 2015

M. Nottingham
July 04, 2014

Problems with Proxies in HTTP
draft-nottingham-http-proxy-problem-01

Abstract

This document discusses the use and configuration of proxies in HTTP, pointing out problems in the currently deployed Web infrastructure along the way. It then offers a few principles to base further discussion upon, and lists some potential avenues for further exploration.

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 January 05, 2015.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1. Introduction	2
1.1. Notational Conventions	3
2. Why Proxy?	3
2.1. Application Layer Gatewaying	4
2.2. Caching	4
2.3. Network Policy Enforcement	4
2.4. Content Filtering (a.k.a. Content Policy Enforcement) . .	4
2.5. Content Modification	5
3. How Proxies are Interposed	6
3.1. Manual Configuration	6
3.2. proxy.pac and WPAD	6
3.3. Interception	7
3.4. Configuration As Side Effect	8
4. Second-Order Effects of Proxy Deployment	8
4.1. Proxies and HTTP	8
4.2. Proxies and TLS	9
5. Principles for Consideration	9
5.1. Proxies Have a Legitimate Place	10
5.2. Security Should be Encouraged	10
5.3. Users Need to be Informed of Proxies	10
5.4. Users Need to be able to Tunnel through Proxies	11
5.5. Proxies Can say "No"	11
5.6. Changes Need to be Detectable	11
5.7. Proxies Need to be Easy	11
5.8. Proxies Need to Communicate to Users	11
5.9. Users Require Simple Interfaces	12
5.10. User Agents Are Diverse	12
5.11. RFC2119 Doesn't Define Reality	12
5.12. It Needs to be Deployable	13
6. Potential Areas to Investigate	13
6.1. Improving Proxy.Pac	13
6.2. TLS Errors for Proxies	13
6.3. HTTP Errors for Proxies	13
6.4. TLS for Proxy Connections	14
6.5. Improved Network Information	14
6.6. Improving Trust	14
6.7. HTTP Signatures	14
7. Security Considerations	15
8. Acknowledgements	15
9. References	15
9.1. Normative References	15
9.2. Informative References	15
Author's Address	16

1. Introduction

HTTP/1.1 [RFC7230] was designed to accommodate proxies. It allows them (and other components) to cache content expansively, and allows for proxies to break "semantic transparency" by changing message content, within broad constraints.

As the Web has matured, more networks have taken advantage of this by deploying proxies for a variety of reasons, in a number of different ways. Section 2 is a survey of the different ways that proxies are used, and Section 3 shows how they are interposed into communication.

Some uses of proxies cause problems (or the perception of them) for origin servers and end users. While some uses are obviously undesirable from the perspective of an end users and/or origin server, other effects of their deployment are more subtle; these are examined in Section 4.

These tensions between the interests of the stakeholders in every HTTP connection - the end users, the origin servers and the networks they use - has led to decreased trust for proxies, then increasing deployment of encryption, then workarounds for encryption, and so forth.

Left unchecked, this escalation can erode the value of the Web itself. Therefore, Section 5 proposes straw-man principals to base further discussion upon.

Finally, Section 6 proposes some areas of technical investigation that might yield solutions (or at least mitigations) for some of these problems.

Note that this document is explicitly about "proxies" in the sense that HTTP defines them. Intermediaries that are interposed by the server (e.g., gateways and so-called "Reverse Proxies", as used in Content Delivery Networks) are out of scope.

1.1. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Why Proxy?

HTTP proxies are interposed between user agents and origin servers for a variety of purposes; some of them are with the full knowledge and consent of end users, to their benefit, and some are solely for the purposes of the network operator - sometimes even against the interests of the end users.

This section attempts to identify the different motivations networks have for deploying proxies.

2.1. Application Layer Gatewaying

Some networks do not have direct Internet connectivity for Web browsing. These networks can deploy proxies that do have Internet connectivity and then configure clients to use them.

Such gatewaying between networks were some of the first uses for proxies.

2.2. Caching

An extremely common use of proxies is to interpose a HTTP cache, in order to save bandwidth, improve end-user perceived latency, increase reliability, or some combination of these purposes.

HTTP defines a detailed model for caching (see [RFC7234]); however, some lesser-known aspects of the caching model can cause operational issues. For example, it allows caches to go into an "offline" mode where most content can be served stale.

Also, proxy caches sometimes fail to honor the HTTP caching model, reusing content when it should not have been. This can cause interoperability issues, with the end user seeing overly "stale" content, or applications not operating correctly.

2.3. Network Policy Enforcement

Some proxies are deployed to aid in network policy enforcement; for example, to control access to the network, requiring a login (as allowed explicitly by HTTP's proxy authentication mechanism), bandwidth shaping of HTTP access, quotas, etc. This includes so-called "Captive Portals" used for network login.

Some uses of proxies for policy enforcement cause problems; e.g., when a proxy uses URL rewriting to send a user a message (e.g., a "blocked" page), they can make it appear as if the origin server is sending that message - especially when the user agent isn't a browser (e.g., a software update process).

2.4. Content Filtering (a.k.a. Content Policy Enforcement)

Some networks attempt to filter HTTP messages (both request and response) based upon network-specific criteria. For example, they might wish to stop users from downloading content that contains malware, or that violates site policies on appropriate content, or that violates local law.

Intermediary proxies as a mechanism for enforcing content restrictions are often easy to circumvent. For example, a device might become infected by using a different network, or a VPN. Nevertheless, they are commonly used for this purpose.

Some content policy enforcement is also done locally to the user agent; for example, several Operating Systems have machine-local proxies built in that scan content.

Content filtering is often seen as controversial, often depending on the context it is used within and how it is performed.

2.5. Content Modification

Some networks modify HTTP messages (both request and response) as they pass through proxies. This might include the message body, headers, request-target, method or status code.

Motivation for content modification varies. For example, some mobile networks interpose proxies that modify content in an attempt to save bandwidth, improve perceived performance, or transcode content to formats that limited-resource devices can more easily consume.

Modifications also include adding metadata in headers for accounting purposes, or removing metadata such as Accept-Encoding to make virus scanning easier.

In other cases, content modification is performed to make more substantial modifications. This could include inserting advertisements, or changing the layout of content in an attempt to make it easier to use.

Content modification is very controversial, often depending on the context it is used within and how it is performed. Many feel that, without the explicit consent of either the end user or the origin server, a proxy that modifies content violates their relationship, thereby degrading trust in the Web overall.

However, it should be noted that Section 5.7.2 of [RFC7230] explicitly allows "non-transparent" proxies that modify content in certain ways. Such proxies are required to honor the "no-transform" directive, giving both user agents and origin servers a mechanism to

"opt out" of modifications ([RFC7234], Section 5.2.1.6); however, it is not technically enforced.

[W3C.NOTE-ct-guidelines-20101026] is a product of the W3C Mobile Web Best Practices Working Group that attempts to set guidelines for content modification proxies. Again, it is a policy document, without technical enforcement measures.

3. How Proxies are Interposed

How a proxy is interposed into a network flow often has great affect on perceptions of its operation by end users and origin servers. This section catalogues the ways that this happens, and potential problems with each.

3.1. Manual Configuration

The original way to interpose a proxy was to manually configure it into the user agent. For example, most browsers still have the ability to have a proxy hostname and port configured for HTTP; many Operating Systems have system-wide proxy settings.

Unfortunately, manual configuration suffers from several problems:

- o Users often lack the expertise to manually configure proxies.
- o When the user changes networks, they must manually change proxy settings, a laborious task. This makes manual configuration impractical in a modern, mobile-driven world.
- o Not all HTTP stacks support manual proxy configuration. Therefore, a proxy administrator cannot rely upon this method.

3.2. proxy.pac and WPAD

The limitations of manual configuration were recognized long ago. The solution that evolved was a format called "proxy.pac" [proxypac] that allowed the proxy configuration to be automated, once the user agent had loaded it.

Proxy.pac is a JavaScript format; before each request is made, it is dispatched to a function in the file that returns a string that denotes whether a proxy is to be used, and if so, which one to use.

Discovery of the appropriate proxy.pac file for a given network can be made using a DHCP extension, [wpad]. WPAD started as a simple protocol; it conveys a URL that locates the proxy.pac file for the network.

Unfortunately, the proxy.pac/WPAD combination has several operational issues that limit its deployment:

- o The proxy.pac format does not define timeouts or failover behavior precisely, leading to wide divergence between implementations. This makes supporting multiple user agents reliably difficult for the network.
- o WPAD is not widely implemented by user agents; some only implement proxy.pac.
- o In those user agents where it is implemented, WPAD is often not the default, meaning that users need to configure its use.
- o Neither proxy.pac nor WPAD have been standardized, leading to implementation divergence and resulting interoperability problems.
- o There are DNS-based variants of WPAD, adding to to confusion.
- o DHCP options generally require tight integration with the operating system to pass the results to HTTP-based applications. While this level of integration is found between O/Ses and their provided applications, the interface may or may not be available to third parties.
- o WPAD can be spoofed, allowing attackers to interpose a proxy and intercept traffic. This is a blocking issue for implementation.

3.3. Interception

The problems with manual configuration and proxy.pac/WPAD have led to the wide deployment of a third style of interposition; interception proxies.

Interception occurs when lower-layer protocols are configured to route HTTP traffic to a host other than the origin server for the URI in question. It requires no client configuration (hence its popularity over other methods). See [RFC3040] for an example of an interception-related protocol.

Interception is also strongly motivated when it is necessary to assure that the proxy is always used, e.g., to enforce policy.

Interception is problematic, however, because it is often done without the consent of either the end user or the origin server. This means that a response that appears to be coming from the origin server is actually coming from the intercepting proxy. This makes it difficult to support features like proxy authentication, as the

unexpected status code breaks many clients (e.g., non-interactive applications like software installers).

Furthermore, interception is a "layer violation"; i.e., misusing lower-layer protocols to enforce a higher-layer (often expressed as "layer 8") requirement.

In addition, as adoption of multi-path TCP (MPTCP) [RFC6824] increases, the ability of intercepting proxies to offer a consistent service degrades.

3.4. Configuration As Side Effect

More recently, it's become more common for a proxy to be interposed as a side effect of another choice by the user.

For example, the user might decide to add virus scanning - either as installed software, or a service that they configure from their provider - that is interposed as a proxy. Indeed, almost all desktop virus scanners and content filters operate in this fashion.

This approach has the merits of both being easy and obtaining explicit user consent. However, in some cases, the end user might not understand the consequences of use of the proxy, especially upon security and interoperability.

4. Second-Order Effects of Proxy Deployment

4.1. Proxies and HTTP

Deployment of proxies has an effect on the HTTP protocol itself. Because a proxy implements both a server and a client, any limitations or bugs in their implementation impact the protocol's use.

For example, HTTP has a defined mechanism for upgrading the protocol of a connection, to aid in the deployment of new versions of HTTP (such as HTTP/2) or completely different protocol (e.g., [RFC6455]).

However, operational experience has shown that a significant number of proxy implementations do not correctly implement it, leading to dangerous situations where two ends of a HTTP connection think different protocols are being spoken.

Another example is the Expect/100-continue mechanism in HTTP/1.1, which is often incorrectly implemented. Likewise, differences in support for trailers limits protocol extensions.

4.2. Proxies and TLS

It has become more common for Web sites to use TLS [RFC5246] in an attempt to avoid many of the problems above. Many have advocated use of TLS more broadly; for example, see the EFF's HTTPS Everywhere [https-everywhere] program, and SPDY's default use of TLS [I-D.mbelshe-httpbis-spdy].

However, doing so engenders a few problems.

Firstly, TLS as used on the Web is not a perfectly secure protocol, and using it to protect all traffic gives proxies a strong incentive to work around it, e.g., by deploying a certificate authority directly into browsers, or buying a sub-root certificate.

Secondly, it removes the opportunity for the proxy to inform the user agent of relevant information; for example, conditions of access, access denials, login interfaces, and so on. User Agents currently do not display any feedback from proxy, even in the CONNECT response (e.g., a 4xx or 5xx error), limiting their ability to have informed users of what's going on.

Finally, it removes the opportunity for services provided by a proxy that the end user might wish to opt into. For example, consider when a remote village shares a proxy server to cache content, thereby helping to overcome the limitations of their Internet connection. TLS-protected HTTP traffic cannot be cached by intermediaries, removing much of the benefit of the Web to what is arguably one of its most important target audiences.

It is now becoming more common for a proxy to man-in-the-middle TLS connections (see [tls-mitm] for an overview), to gain access to the application message flows. This represents a serious degradation in the trust infrastructure of the Web.

Worse is the situation where proxies provide a certificate where they inure the user to a certificate warning that they then need to ignore in order to receive service.

5. Principles for Consideration

Every HTTP connection has at least three major stakeholders; the user (through their agent), the origin server (possibly using gateways such as a CDN) and the networks between them.

Currently, the capabilities of these stakeholders are defined by how the Web is deployed. Most notably, networks sometimes change content. If they change it too much, origin servers will start using

encryption. Changing the way that HTTP operates therefore has the potential to re-balance the capabilities of the various stakeholders.

This section proposes several straw-man principles for consideration as the basis of those changes. Their sole purpose here is to provoke discussion.

5.1. Proxies Have a Legitimate Place

As illustrated above, there are many legitimate uses for proxies, and they are a necessary part of the architecture of the Web. While all uses of proxies are not legitimate - especially when they're interposed without the knowledge or consent of the end user and the origin - undesirable intermediaries (i.e., those that break the reasonable expectations of other stakeholders) are a small portion of those deployed used.

Note that while proxies have a legitimate place, it does not imply that they are an equal stakeholder to other parties in all ways; e.g., they do not have a natural right to access encrypted content, for example.

5.2. Security Should be Encouraged

Any solution needs to give all stakeholders - end users, networks and origin servers - a strong incentive towards security.

This has subtle implications. If networks are disempowered disproportionately, they might react by blocking secure connections, discouraging origin servers (who often have even stronger profit incentives) from deploying encryption, which would result in a net loss of security.

On the other hand, if networks are given carte blanche, it can destroy trust in the Web altogether. In particular, making it too easy to interpose a proxy (even if the user is "informed" by clicking through a dialogue) degrades the infrastructure in an unacceptable way.

5.3. Users Need to be Informed of Proxies

When a proxy is interposed, the user needs to be informed about it, so they have the opportunity to change their configuration (e.g., attempt to introduce encryption), or not use the network at all.

Proxies also need to be strongly authenticated; i.e., users need to be able to verify who the proxy is.

5.4. Users Need to be able to Tunnel through Proxies

When a proxy is interposed, the user needs to be able to tunnel any request through it without its content (or that of the response) being exposed to the proxy.

This includes both "https://" and "http://" URIs.

5.5. Proxies Can say "No"

A proxy can refuse to forward any request. Currently, the granularity of that "no" is per-URI for unencrypted requests, and per-IP (perhaps per-SNI) for encrypted requests.

5.6. Changes Need to be Detectable

Any changes to the message body, request URI, method, status code, or representation header fields of an HTTP message need to be detectable by the origin server or user agent, as appropriate, if they desire it.

This allows a proxy to be trusted, but its integrity to be verified.

5.7. Proxies Need to be Easy

It must be possible to configure a proxy extremely easily; the adoption of interception over proxy.pac/WPAD illustrates this very clearly.

5.8. Proxies Need to Communicate to Users

There are many situations where a proxy needs to communicate with the end user; for example, to gather network authentication credentials, communicate network policy, report that access to content has been denied, and so on.

Currently, HTTP has poor facilities for doing so. The proxy authentication mechanism is extremely limited, and while there are a few status codes that are defined as being from a proxy rather than the origin, they do not cover all necessary situations.

The Warning header field ([RFC7234], Section 5.5) was designed as a very limited form of communication between proxies and end users, but it has not been widely adopted, nor exposed by User Agents.

Importantly, proxies also need a limited communication channel when TLS is in use, for similar purposes.

Equally as important, the communication needs to clearly come from the proxy, rather than the origin, and be strongly authenticated.

5.9. Users Require Simple Interfaces

While some users are sophisticated in their understanding of Web security, they are in a vanishingly small minority. The concepts and implications of many decisions regarding security are subtle, and require an understanding of how the Web works; describing these trade-offs in a modal dialogue box that gets in the way of the content the user wants has been proven not to work.

Similarly, while some Web publishers are sophisticated regarding security, the vast majority are not (as can be proven by the prevalence of cross-site scripting attacks).

Therefore, any changes cannot rely upon perfect understanding by these parties, or even any great effort upon their part. This implies that user interface will be one of the biggest challenges faced, both in the browser and for any changes server-side.

Notably, the most widely understood indicator of security today is the "lock icon" that shows when a connection is protected by TLS. Any erosion of the commonly-understood semantics of that indicator, as well as "https://" URIs, is likely to be extremely controversial, because it changes the already-understood security properties of the Web.

Another useful emerging convention is that of "Incognito" or "private" mode, where the end user has requested enhanced privacy and security. This might be used to introduce higher requirements for the interposition of intermediaries, or even to prohibit their use without full encryption.

5.10. User Agents Are Diverse

HTTP is used in a wide variety of environments. As such there can be no assumption that a user is sitting on the other end to interpret information or answer questions from proxies.

5.11. RFC2119 Doesn't Define Reality

It's very tempting for a committee to proclaim that proxies "MUST" do this and "SHOULD NOT" do that, but the reality is that the proxies, like any other actor in a networked system, will do what they can, not what they're told to do, if they have an incentive to do it.

Therefore, it's not enough to say that (for example), "proxies have to honor no-transform" as HTTP/1.1 does. Instead, the protocol needs to be designed in a way so that either transformations aren't possible, or if they are, they can be detected (with appropriate handling by User Agents defined).

5.12. It Needs to be Deployable

Any improvements to the proxy ecosystem **MUST** be incrementally deployable, so that existing clients can continue to function.

6. Potential Areas to Investigate

Finally, this section lists some areas of potential future investigation, bearing the principles suggested above in mind.

6.1. Improving Proxy.Pac

Many of the flaws in proxy.pac can be fixed by careful specification and standardization, with active participation by both implementers and those that deploy it.

6.2. TLS Errors for Proxies

HTTP's use of TLS [RFC2818] currently offers no way for an interception proxy to communicate with the user agent on its own behalf. This might be necessary for network authentication, notification of filtering by hostname, etc.

The challenge in defining such a mechanism is avoiding the opening of new attack vectors; if unauthenticated content can be served as if it were from the origin server, or the user can be encouraged to "click through" a dialog, it has severe security implications. As such, the user experience would need to be carefully considered.

6.3. HTTP Errors for Proxies

HTTP currently defines two status codes that are explicitly generated by a proxy:

- o 504 Gateway Timeout ([RFC7231], Section 6.6.5) - when a proxy (or gateway) times out going forward
- o 511 Network Authentication Required ([RFC6585], Section 6) - when authentication information is necessary to access the network

It might be interesting to discuss whether a separate user experience can be formed around proxy-specific status codes, along with the definition of new ones as necessary.

6.4. TLS for Proxy Connections

While TLS can be used end-to-end for "https://" URIs, support for connecting to a proxy itself using TLS (e.g., for "http://" URIs) is spotty. Using a proxy without strong proof of its identity introduces security issues, and if a proxy can legitimately insert itself into communication, its identity needs to be verifiable.

6.5. Improved Network Information

Many of the use cases for proxies that modify content is transcoding or otherwise adapting that which is too "heavy" for the network it is transiting through.

If network operators made better, more fine-grained and timely information about their operational characteristics freely available, endpoints (server and client) could adapt requests and responses to reflect it, thereby removing the need for intermediation.

6.6. Improving Trust

Currently, it is possible to exploit the mismatched incentives and other flaws in the CA system to cause a browser to trust a proxy as authoritative for a "https://" URI without full user knowledge. This needs to be remedied; otherwise, proxies will continue to man-in-the-middle TLS.

6.7. HTTP Signatures

Signatures for HTTP content - both requests and responses - have been discussed on and off for some time.

Of particular interest here, signed responses would allow a user-agent to verify that the origin's content has not been modified in transit, whilst still allowing it to be cached by intermediaries.

Likewise, if header values can be signed, the caching policy (as expressed by Cache-Control, Date, Last-Modified, Age, etc.) can be signed, meaning it can be verified as being adhered to.

Note that properly designed, a signature mechanism could work over TLS, separating the trust relationship between the UA and the origin server and that of the UA and its proxy (with appropriate consent).

There are significant challenges in designing a robust, widely-deployable HTTP signature mechanism. One of the largest is an issue of user interface - what ought the UA do when encountering a bad signature?

7. Security Considerations

Plenty of them, I suspect.

8. Acknowledgements

This document benefits from conversations and feedback from many people, including Amos Jeffries, Willy Tarreau, Patrick McManus, Roberto Peon, Guy Podjarny, Eliot Lear, Brad Hill, Martin Nilsson and Julian Reschke.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. Informative References

- [I-D.mbelshe-httpbis-spy]
Belshe, M. and R. Peon, "SPDY Protocol", draft-mbelshe-httpbis-spy-00 (work in progress), February 2012.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.
- [RFC3040] Cooper, I., Melve, I., and G. Tomlinson, "Internet Web Replication and Caching Taxonomy", RFC 3040, January 2001.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC6455] Fette, I. and A. Melnikov, "The WebSocket Protocol", RFC 6455, December 2011.
- [RFC6585] Nottingham, M. and R. Fielding, "Additional HTTP Status Codes", RFC 6585, April 2012.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, January 2013.

- [RFC7230] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, June 2014.
- [RFC7231] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, June 2014.
- [RFC7234] Fielding, R., Nottingham, M., and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Caching", RFC 7234, June 2014.
- [W3C.NOTE-ct-guidelines-20101026]
Rabin, J., "Guidelines for Web Content Transformation Proxies 1.0", World Wide Web Consortium NOTE NOTE-ct-guidelines-20101026, October 2010,
<<http://www.w3.org/TR/2010/NOTE-ct-guidelines-20101026>>.
- [https-everywhere]
EFF, ., "HTTPS Everywhere", 2013, <<https://www.eff.org/https-everywhere>>.
- [proxypac]
various, ., "Proxy Auto-Config", 2013,
<http://en.wikipedia.org/wiki/Proxy_auto-config>.
- [tls-mitm]
Jarmoc, J., "SSL/TLS Interception Proxies and Transitive Trust", 2012, <https://www.grc.com/miscfiles/HTTPS_Interception_Proxies.pdf>.
- [wpad] Cohen, J., "Web Proxy Auto-Discovery Protocol", 1999,
<<http://tools.ietf.org/html/draft-ietf-wrec-wpad-01>>.

Author's Address

Mark Nottingham

Email: mnot@mnot.netURI: <http://www.mnot.net/>

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2015

J. Reschke
greenbytes
March 9, 2015

Hypertext Transfer Protocol (HTTP) Client-Initiated Content-Encoding
draft-reschke-http-cice-02

Abstract

In HTTP, "Content Codings" allow for payload encodings such as for compression or integrity checks. In particular, the "gzip" content coding is widely used for payload data sent in response messages.

Content Codings can be used in request messages as well, however discoverability is not on par with response messages. This document extends the HTTP "Accept-Encoding" header field for use in responses.

Editorial Note (To be removed by RFC Editor before publication)

Distribution of this document is unlimited. Although this is not a work item of the HTTPbis Working Group, comments should be sent to the Hypertext Transfer Protocol (HTTP) mailing list at ietf-http-wg@w3.org [1], which may be joined by sending a message with subject "subscribe" to ietf-http-wg-request@w3.org [2].

Discussions of the HTTPbis Working Group are archived at <http://lists.w3.org/Archives/Public/ietf-http-wg/>.

XML versions and latest edits for this document are available from <http://greenbytes.de/tech/webdav/#draft-reschke-http-cice>.

The changes in this draft are summarized in Appendix A.2.

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 September 10, 2015.

Copyright Notice

Copyright (c) 2015 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.

Table of Contents

1. Introduction	3
2. Notational Conventions	3
3. Extensions to 'Accept-Encoding' Header Field	3
4. Example	4
5. Deployment Considerations	5
6. Security Considerations	5
7. IANA Considerations	5
8. Acknowledgements	5
9. References	6
9.1. Normative References	6
9.2. Informative References	6
Appendix A. Change Log (to be removed by RFC Editor before publication)	6
A.1. draft-reschke-http-cice-00	6
A.2. draft-reschke-http-cice-01	6

1. Introduction

In HTTP, "Content Codings" allow for payload encodings such as for compression or integrity checks ([RFC7231], Section 3.1.2). In particular, the "gzip" content coding is widely used for payload data sent in response messages.

Content Codings can be used in request messages as well, however discoverability is not on par with response messages. This document extends the HTTP "Accept-Encoding" header field ([RFC7231], Section 5.3.4) for use in responses.

2. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

This document reuses terminology used in the base HTTP specifications, namely Section 2 of [RFC7230] and Section 3.1.2 of [RFC7231].

3. Extensions to 'Accept-Encoding' Header Field

Section 5.3.4 of [RFC7231] defines "Accept-Encoding" as a request header field only.

This specification extends that definition to allow "Accept-Encoding" as a response header field as well. When present, it indicates what content codings a resource was willing to accept at the time of the response. A field value that only contains "identity" implies that no content codings are supported.

Note that this information is specific to the specific request. The set of supported encodings might be different for other resources on the same server, could also change depending on other aspects of the request (such as the request method), or might change in the future.

Section 6.5.13 of [RFC7231] defines status code 415 (Unsupported Media Type) to apply to both media type and content coding related problems.

Servers that fail a request due to an unsupported content coding SHOULD respond with a 415 status and SHOULD include an "Accept-Encoding" header field in that response, allowing clients to distinguish between content coding related issues and media type related issues. In order to avoid confusion with media type related problems, servers that fail a request with a 415 status for reasons

unrelated to content codings SHOULD NOT include the "Accept-Encoding" header field.

While sending "Accept-Encoding" in a 415 (Unsupported Media Type) response will be the most common use case, it is not restricted to this particular status code. For instance, a server might include it in a 2xx response when a request payload was big enough to justify use of a compression coding, but the client failed to do so.

4. Example

Client submits a POST request using Content-Encoding "compress" ([RFC7231], Section 3.1.2.1):

```
POST /edit/ HTTP/1.1
Host: example.org
Content-Type: application/atom+xml;type=entry
Content-Encoding: compress
```

...compressed payload...

Server rejects request because it only allows the "gzip" content coding:

```
HTTP/1.1 415 Unsupported Media Type
Date: Fri, 09 May 2014 11:43:53 GMT
Accept-Encoding: gzip
Content-Length: 68
Content-Type: text/plain
```

This resource only supports the "gzip" content coding in requests.

...at which point the client can retry the request with the supported "gzip" content coding.

Alternatively, a server that does not support any content codings in requests could answer with:

```
HTTP/1.1 415 Unsupported Media Type
Date: Fri, 09 May 2014 11:43:53 GMT
Accept-Encoding: identity
Content-Length: 61
Content-Type: text/plain
```

This resource does not support content codings in requests.

5. Deployment Considerations

Servers that do not support content codings in requests already are required to fail a request that does use a content coding. Section 6.5.13 of [RFC7231] recommends to use the status code 415 (Unsupported Media Type), so the only change needed is to include the "Accept-Encoding" header field with value "identity" in that response.

Servers that do support some content codings are required to fail requests with unsupported content codings as well. To be compliant with this specification, servers will need to use the status code 415 (Unsupported Media Type) to signal the problem, and will have to include an "Accept-Encoding" header field that enumerates the content codings that are supported. As the set of supported content codings usually is static and small, adding the header field ought to be trivial.

6. Security Considerations

This specification does not introduce any new security considerations beyond those discussed in Section 9 of [RFC7231].

7. IANA Considerations

HTTP header fields are registered within the "Message Headers" registry located at <http://www.iana.org/assignments/message-headers>, as defined by [BCP90].

This document updates the definition of the "Accept-Encoding" header field, so the "Permanent Message Header Field Names" registry shall be updated accordingly:

Header Field Name	Protocol	Status	Reference
Accept-Encoding	http	standard	[RFC7231], Section 5.3.4, extended by Section 3 of this document

8. Acknowledgements

Thanks go to the members of the and HTTPbis Working Group, namely Amos Jeffries, Mark Nottingham, and Ted Hardie.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, June 2014.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, June 2014.

9.2. Informative References

- [BCP90] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", BCP 90, RFC 3864, September 2004.

URIs

- [1] <mailto:ietf-http-wg@w3.org>
- [2] <mailto:ietf-http-wg-request@w3.org?subject=subscribe>

Appendix A. Change Log (to be removed by RFC Editor before publication)

A.1. draft-reschke-http-cice-00

Clarified that the information returned in Accept-Encoding is per resource, not per server.

Added some deployment considerations.

Updated HTTP/1.1 references.

A.2. draft-reschke-http-cice-01

Restrict the scope of A-E from "future requests" to "at the time of this request".

Mention use of A-E in responses other than 415.

Recommend not to include A-E in a 415 response unless there was actually a problem related to content coding.

Author's Address

Julian F. Reschke
greenbytes GmbH
Hafenweg 16
Muenster, NW 48155
Germany

EMail: julian.reschke@greenbytes.de
URI: <http://greenbytes.de/tech/webdav/>

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 3, 2015

J. Reschke
greenbytes
July 2, 2014

A JSON Encoding for HTTP Header Field Values
draft-reschke-http-jfv-00

Abstract

This document establishes a convention for use of JSON-encoded field values in HTTP header fields.

Editorial Note (To be removed by RFC Editor before publication)

Distribution of this document is unlimited. Although this is not a work item of the HTTPbis Working Group, comments should be sent to the Hypertext Transfer Protocol (HTTP) mailing list at ietf-http-wg@w3.org [1], which may be joined by sending a message with subject "subscribe" to ietf-http-wg-request@w3.org [2].

Discussions of the HTTPbis Working Group are archived at <http://lists.w3.org/Archives/Public/ietf-http-wg/>.

XML versions and latest edits for this document are available from <http://greenbytes.de/tech/webdav/#draft-reschke-http-jfv>.

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 January 3, 2015.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1. Introduction	3
2. Data Model and Format	3
3. Sender Requirements	4
4. Recipient Requirements	5
5. Using this Format in Header Field Definitions	5
6. Examples	5
6.1. Content-Length	5
6.2. Content-Disposition	6
6.3. WWW-Authenticate	7
7. Discussion	8
8. Deployment Considerations	8
9. Internationalization Considerations	8
10. Security Considerations	8
11. References	8
11.1. Normative References	8
11.2. Informative References	9
Appendix A. Sample Code and Test Cases	9

1. Introduction

Defining syntax for new HTTP header fields ([RFC7230], Section 3.2) is non-trivial. Among the commonly encountered problems are:

- o There is no common syntax for complex field values. Several well-known header fields do use a similarly looking syntax, but it is hard to write generic parsing code that will both correctly handle valid field values but also fail on invalid ones.
- o The HTTP message format allows header fields to repeat, so field syntax needs to be designed in a way that these cases are either meaningful, or can be unambiguously detected and rejected.
- o HTTP/1.1 does not define a character encoding scheme ([RFC6365], Section 2), so header fields are either stuck with US-ASCII ([USASCII]), or need out-of-band information to decide what encoding scheme is used. Furthermore, APIs usually assume a default encoding scheme in order to map from octet sequences to strings (for instance, [XMLHttpRequest] uses the IDL type "ByteString", effectively resulting in the ISO-8859-1 character encoding scheme [ISO-8859-1] being used).

(See Section 8.3.1 of [RFC7231] for a summary of considerations for new header fields.)

This specification addresses the issues listed above by defining both a generic JSON-based ([RFC7159]) data model and a concrete wire format that can be used in definitions of new header fields.

2. Data Model and Format

In HTTP, header fields with the same field name can occur multiple times within a single message (Section 3.2.2 of [RFC7230]). When this happens, recipients are allowed to combine the field values using commas as delimiter. This rule matches nicely JSON's array format (Section 5 of [RFC7159]). Thus, the basic data model used here is the JSON array.

Header field definitions that need only a single value can restrict themselves to arrays of length 1, and are encouraged to define error handling in case more values are received (such as "first wins", "last wins", or "abort with fatal error message").

JSON arrays are mapped to field values by creating a sequence of serialized member elements, separated by commas and optionally whitespace. This is equivalent to using the full JSON array format, while leaving out the "begin-array" ('[') and "end-array" (']')

delimiters.

The ABNF character names and classes below are used (copied from [RFC5234], Appendix B.1):

CR	= %x0D	; carriage return
HTAB	= %x09	; horizontal tab
LF	= %x0A	; line feed
SP	= %x20	; space
VCHAR	= %x21-7E	; visible (printing) characters

Characters in JSON strings that are not allowed or discouraged in HTTP header field values -- that is, not in the "VCHAR" definition -- need to be represented using JSON's "backslash" escaping mechanism ([RFC7159], Section 7).

The control characters CR, LF, and HTAB do not appear inside JSON strings, but can be used outside (line breaks, indentation etc). These characters can be either stripped or replaced by space characters (ABNF "SP").

Formally, using the HTTP specification's ABNF extensions defined in Section 7 of [RFC7230]:

```
json-field-value = #json-field-item
json-field-item = JSON-Text
                  ; see [RFC7159], Section 2,
                  ; post-processed so that only VCHAR characters
                  ; are used
```

3. Sender Requirements

[[anchor3: The text below assumes we're starting with a JSON-formatted sequence of characters, not octets; need to clarify.]] To map a JSON array to an HTTP header field value, process each array element separately by:

1. generating the JSON representation,
2. stripping all JSON control characters (CR, HTAB, LF), or replacing them by space ("SP") characters,
3. replacing all remaining non-VSPACE characters by the equivalent backslash-escape sequence ([RFC7159], Section 7).

The resulting list of strings is transformed into an HTTP field value by combining them using comma (%x2C) plus optional SP as delimiter, and encoding the resulting string into an octet sequence using the

US-ASCII character encoding scheme.

4. Recipient Requirements

To map a set of HTTP header field instances to a JSON array:

1. remove all header field instances that only contain whitespace (SP / HTAB) and "comma" characters [[anchor5: either drop this or make it more precise]],
2. combine all header field instances into a single field as per Section 3.2.2 of [RFC7230],
3. add a leading begin-array ("[") octet and a trailing end-array ("]") octet, then
4. run the resulting octet sequence through a JSON parser.

The result of the parsing operation is either an error (in which case the header field values needs to be considered invalid), or a JSON array.

5. Using this Format in Header Field Definitions

[[anchor7: Explain what a definition of a new header field needs to do precisely to use this format]]

6. Examples

This section shows how some of the existing HTTP header fields would look like if they would use the format defined by this specification.

6.1. Content-Length

"Content-Length" is defined in Section 3.3.2 of [RFC7230], with the field value's ABNF being:

Content-Length = 1*DIGIT

So the field value is similar to a JSON number ([RFC7230], Section 6).

Content-Length is restricted to a single field instance, as it doesn't use the list production (as per Section 3.2.2 of [RFC7230]). However, in practice multiple instances do occur, and the definition of the header field does indeed discuss how to handle these cases.

If Content-Length was defined using the JSON format discussed here,

the ABNF would be something like:

```
Content-Length = #number
                ; number: [RFC7159], Section 6
```

...and the prose definition would:

- o restrict all numbers to be non-negative integers without fractions, and
- o require that the array of values is of length 1 (but allow the case where the array is longer, but all members represent the same value)

6.2. Content-Disposition

Content-Disposition field values, defined in [RFC6266], consist of a "disposition type" (a string), plus multiple parameters, of which at least one ("filename") sometime needs to carry non-ASCII characters.

For instance, the first example in Section 5 of [RFC6266]:

```
Attachment; filename=example.html
```

has a disposition type of "Attachment", with filename parameter value "example.html". A JSON representation of this information might be:

```
{
  "Attachment": {
    "filename" : "example.html"
  }
}
```

which would translate to a header field value of:

```
{ "Attachment": { "filename" : "example.html" } }
```

The third example in Section 5 of [RFC6266] uses a filename parameter containing non-US-ASCII characters:

```
attachment; filename*=UTF-8''%e2%82%ac%20rates
```

Note that in this case, the "filename*" parameter uses the encoding defined in [RFC5987], representing a filename starting with the Unicode character U+20AC (EURO SIGN), followed by " rates". If the definition of Content-Disposition would have used the format proposed here, the workaround involving the "parameter*" syntax would not have been needed at all.

The JSON representation of this value could then be:

```
{ "attachment": { "filename" : "\u20AC rates" } }
```

6.3. WWW-Authenticate

The WWW-Authenticate is defined in Section 4.1 of [RFC7235] as a list of "challenges":

```
WWW-Authenticate = 1#challenge
```

...where a challenge consists of a scheme with optional parameters:

```
challenge      = auth-scheme [ 1*SP ( token68 / #auth-param ) ]
```

An example for a complex header field value given in the definition of the header field is:

```
Newauth realm="apps", type=1, title="Login to \"apps\"",  
Basic realm="simple"
```

(line break added for readability)

A possible JSON representation of this field value would be the array below:

```
[  
  {  
    "Newauth" : {  
      "realm": "apps",  
      "type" : 1,  
      "title" : "Login to \"apps\""  
    }  
  },  
  {  
    "Basic" : {  
      "realm": "simple"  
    }  
  }  
]
```

...which would translate to a header field value of:

```
{ "Newauth" : { "realm": "apps", "type" : 1,  
                "title": "Login to \"apps\" " } },  
{ "Basic" : { "realm": "simple" } }
```


7. Discussion

This approach uses a default of "JSON array", using implicit array markers. An alternative would be a default of "JSON object". This would simplify the syntax for non-list-typed headers, but all the benefits of having the same data model for both types of header fields would be gone. A hybrid approach might make sense, as long as it doesn't require any heuristics on the recipient's side.

[[anchor9: Use of generic libs vs compactness of field values...]]

8. Deployment Considerations

[[anchor11: Mention that some code might be refused by double quotes not being used for quoted-string.]]

9. Internationalization Considerations

[[anchor13: TBD, mention migration path to message format that is robust wrt UTF-8, or other binary encodings of JSON]]

10. Security Considerations

[[anchor15: TBD]]

11. References

11.1. Normative References

- | | |
|-----------|--|
| [RFC5234] | Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008. |
| [RFC7159] | Bray, T., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, March 2014. |
| [RFC7230] | Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, June 2014. |
| [RFC7231] | Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, June 2014. |
| [USASCII] | American National Standards Institute, "Coded Character Set -- 7-bit American Standard Code for Information Interchange", ANSI X3.4, 1986. |

11.2. Informative References

- [ISO-8859-1] International Organization for Standardization, "Information technology -- 8-bit single-byte coded graphic character sets -- Part 1: Latin alphabet No. 1", ISO/IEC 8859-1:1998, 1998.
- [RFC5987] Reschke, J., "Character Set and Language Encoding for Hypertext Transfer Protocol (HTTP) Header Field Parameters", RFC 5987, August 2010.
- [RFC6266] Reschke, J., "Use of the Content-Disposition Header Field in the Hypertext Transfer Protocol (HTTP)", RFC 6266, June 2011.
- [RFC6365] Hoffman, P. and J. Klensin, "Terminology Used in Internationalization in the IETF", BCP 166, RFC 6365, September 2011.
- [RFC7235] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Authentication", RFC 7235, June 2014.
- [XMLHttpRequest] van Kesteren, A., Aubourg, J., Song, J., and H. Steen, "XMLHttpRequest Level 1", W3C Working Draft WD-XMLHttpRequest-20140130, January 2014, <<http://www.w3.org/TR/2014/WD-XMLHttpRequest-20140130/>>.
- Latest version available at
<<http://www.w3.org/TR/XMLHttpRequest/>>.

URIs

- [1] <<mailto:ietf-http-wg@w3.org>>
- [2] <<mailto:ietf-http-wg-request@w3.org?subject=subscribe>>

Appendix A. Sample Code and Test Cases

[[anchor19: TBD]]

Author's Address

Julian F. Reschke
greenbytes GmbH
Hafenweg 16
Muenster, NW 48155
Germany

EMail: julian.reschke@greenbytes.de
URI: <http://greenbytes.de/tech/webdav/>

Network Working Group
Internet-Draft
Obsoletes: 5987 (if approved)
Intended status: Standards Track
Expires: January 3, 2015

J. Reschke
greenbytes
July 2, 2014

Indicating Character Encoding and Language for HTTP Header Field
Parameters
draft-reschke-rfc5987bis-07

Abstract

By default, message header field parameters in Hypertext Transfer Protocol (HTTP) messages cannot carry characters outside the ISO-8859-1 character set. RFC 2231 defines an encoding mechanism for use in Multipurpose Internet Mail Extensions (MIME) headers. This document specifies an encoding suitable for use in HTTP header fields that is compatible with a profile of the encoding defined in RFC 2231.

Editorial Note (To be removed by RFC Editor before publication)

Distribution of this document is unlimited. Although this is not a work item of the HTTPbis Working Group, comments should be sent to the Hypertext Transfer Protocol (HTTP) mailing list at ietf-http-wg@w3.org [1], which may be joined by sending a message with subject "subscribe" to ietf-http-wg-request@w3.org [2].

Discussions of the HTTPbis Working Group are archived at <http://lists.w3.org/Archives/Public/ietf-http-wg/>.

XML versions, latest edits, diffs, and the issues list for this document are available from <http://greenbytes.de/tech/webdav/#draft-reschke-rfc5987bis>. A collection of test cases is available at <http://greenbytes.de/tech/tc2231/>.

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 January 3, 2015.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1. Introduction	4
2. Notational Conventions	4
3. Comparison to RFC 2231 and Definition of the Encoding	4
3.1. Parameter Continuations	5
3.2. Parameter Value Character Encoding and Language Information	5
3.2.1. Definition	5
3.2.2. Historical Notes	7
3.2.3. Examples	8
3.3. Language Specification in Encoded Words	8
4. Guidelines for Usage in HTTP Header Field Definitions	9
4.1. When to Use the Extension	9
4.2. Error Handling	9
5. Security Considerations	10
6. IANA Considerations	10
7. Acknowledgements	10
8. References	11
8.1. Normative References	11
8.2. Informative References	11
Appendix A. Changes from RFC 5987	12
Appendix B. Implementation Report	13
Appendix C. Change Log (to be removed by RFC Editor before publication)	13
C.1. Since RFC5987	13
C.2. Since draft-reschke-rfc5987bis-00	13
C.3. Since draft-reschke-rfc5987bis-01	14
C.4. Since draft-reschke-rfc5987bis-02	14
C.5. Since draft-reschke-rfc5987bis-03	14
C.6. Since draft-reschke-rfc5987bis-04	14
C.7. Since draft-reschke-rfc5987bis-05	14
C.8. Since draft-reschke-rfc5987bis-06	14
Appendix D. Open issues (to be removed by RFC Editor prior to publication)	14
D.1. edit	14
D.2. httpbis	14

1. Introduction

By default, message header field parameters in HTTP ([RFC2616]) messages cannot carry characters outside the ISO-8859-1 coded character set ([ISO-8859-1]). RFC 2231 ([RFC2231]) defines an encoding mechanism for use in MIME headers. This document specifies an encoding suitable for use in HTTP header fields that is compatible with a profile of the encoding defined in RFC 2231.

This document obsoletes [RFC5987] and moves it to "historic" status; the changes are summarized in Appendix A.

Note: in the remainder of this document, RFC 2231 is only referenced for the purpose of explaining the choice of features that were adopted; they are therefore purely informative.

Note: this encoding does not apply to message payloads transmitted over HTTP, such as when using the media type "multipart/form-data" ([RFC2388]).

2. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

This specification uses the ABNF (Augmented Backus-Naur Form) notation defined in [RFC5234]. The following core rules are included by reference, as defined in [RFC5234], Appendix B.1: ALPHA (letters), DIGIT (decimal 0-9), HEXDIG (hexadecimal 0-9/A-F/a-f), and LWSP (linear whitespace).

This specification uses terminology defined in [RFC6365], namely: "character encoding scheme" (below abbreviated to "character encoding"), "charset" and "coded character set".

Note that this differs from RFC 2231, which uses the term "character set" for "character encoding scheme".

3. Comparison to RFC 2231 and Definition of the Encoding

RFC 2231 defines several extensions to MIME. The sections below discuss if and how they apply to HTTP header fields.

In short:

- o Parameter Continuations aren't needed (Section 3.1),

- o Character Encoding and Language Information are useful, therefore a simple subset is specified (Section 3.2), and
- o Language Specifications in Encoded Words aren't needed (Section 3.3).

3.1. Parameter Continuations

Section 3 of [RFC2231] defines a mechanism that deals with the length limitations that apply to MIME headers. These limitations do not apply to HTTP ([RFC7231], Appendix A.6).

Thus, parameter continuations are not part of the encoding defined by this specification.

3.2. Parameter Value Character Encoding and Language Information

Section 4 of [RFC2231] specifies how to embed language information into parameter values, and also how to encode non-ASCII characters, dealing with restrictions both in MIME and HTTP header field parameters.

However, RFC 2231 does not specify a mandatory-to-implement character encoding, making it hard for senders to decide which encoding to use. Thus, recipients implementing this specification MUST support the "UTF-8" character encoding [RFC3629].

Furthermore, RFC 2231 allows the character encoding information to be left out. The encoding defined by this specification does not allow that.

3.2.1. Definition

The syntax for parameters is defined in Section 3.6 of [RFC2616] (with RFC 2616 implied LWS translated to RFC 5234 LWSP):

parameter = attribute LWSP "=" LWSP value

attribute = token

value = token / quoted-string

quoted-string = <quoted-string, defined in [RFC7230], Section 3.2.6>

token = <token, defined in [RFC7230], Section 3.2.6>

In order to include character encoding and language information, this specification modifies the RFC 2616 grammar to be:

```

parameter      = reg-parameter / ext-parameter

reg-parameter  = parmname LWSP "=" LWSP value

ext-parameter  = parmname "*" LWSP "=" LWSP ext-value

parmname       = 1*attr-char

ext-value      = charset "'" [ language ] "'" value-chars
                ; like RFC 2231's <extended-initial-value>
                ; (see [RFC2231], Section 7)

charset        = "UTF-8" / mime-charset

mime-charset   = 1*mime-charsetc
mime-charsetc  = ALPHA / DIGIT
                / "!" / "#" / "$" / "%" / "&"
                / "+" / "-" / "^" / "_" / "`"
                / "{" / "}" / "~"
                ; as <mime-charset> in Section 2.3 of [RFC2978]
                ; except that the single quote is not included
                ; SHOULD be registered in the IANA charset registry

language       = <Language-Tag, defined in [RFC5646], Section 2.1>

value-chars    = *( pct-encoded / attr-char )

pct-encoded    = "%" HEXDIG HEXDIG
                ; see [RFC3986], Section 2.1

attr-char      = ALPHA / DIGIT
                / "!" / "#" / "$" / "&" / "+" / "-" / "."
                / "^" / "_" / "`" / "|" / "~"
                ; token except ( "*" / "'" / "%" )

```

Thus, a parameter is either a regular parameter (reg-parameter), as previously defined in Section 3.6 of [RFC2616], or an extended parameter (ext-parameter).

Extended parameters are those where the left-hand side of the assignment ends with an asterisk character.

The value part of an extended parameter (ext-value) is a token that consists of three parts: the REQUIRED character encoding name (charset), the OPTIONAL language information (language), and a character sequence representing the actual value (value-chars), separated by single quote characters. Note that both character encoding names and language tags are restricted to the US-ASCII coded

character set, and are matched case-insensitively (see [RFC2978], Section 2.3 and [RFC5646], Section 2.1.1).

Inside the value part, characters not contained in attr-char are encoded into an octet sequence using the specified character encoding. That octet sequence is then percent-encoded as specified in Section 2.1 of [RFC3986].

Producers MUST use the "UTF-8" ([RFC3629]) character encoding. Extension character encodings (mime-charset) are reserved for future use.

Note: recipients should be prepared to handle encoding errors, such as malformed or incomplete percent escape sequences, or non-decodable octet sequences, in a robust manner. This specification does not mandate any specific behavior, for instance, the following strategies are all acceptable:

- * ignoring the parameter,
- * stripping a non-decodable octet sequence,
- * substituting a non-decodable octet sequence by a replacement character, such as the Unicode character U+FFFD (Replacement Character).

3.2.2. Historical Notes

The RFC 7230 token production ([RFC7230], Section 3.2.6) differs from the production used in RFC 2231 (imported from Section 5.1 of [RFC2045]) in that curly braces ("{" and "}") are excluded. Thus, these two characters are excluded from the attr-char production as well.

The <mime-charset> ABNF defined here differs from the one in Section 2.3 of [RFC2978] in that it does not allow the single quote character (see also RFC Errata ID 1912 [Err1912]). In practice, no character encoding names using that character have been registered at the time of this writing.

For backwards compatibility with RFC 2231, the encoding defined by this specification deviates from common parameter syntax in that the quoted-string notation is not allowed. Implementations using generic parser components might not be able to detect the use of quoted-string notation and thus might accept that format, although invalid, as well.

[RFC5987] did require support for ISO-8859-1, too; for compatibility

with legacy code, recipients are encouraged to support this encoding as well.

3.2.3. Examples

Non-extended notation, using "token":

```
foo: bar; title=Economy
```

Non-extended notation, using "quoted-string":

```
foo: bar; title="US-$ rates"
```

Extended notation, using the Unicode character U+00A3 (POUND SIGN):

```
foo: bar; title*=utf-8'en'%C2%A3%20rates
```

Note: the Unicode pound sign character U+00A3 was encoded into the octet sequence C2 A3 using the UTF-8 character encoding, then percent-encoded. Also, note that the space character was encoded as %20, as it is not contained in attr-char.

Extended notation, using the Unicode characters U+00A3 (POUND SIGN) and U+20AC (EURO SIGN):

```
foo: bar; title*=UTF-8''%c2%a3%20and%20%e2%82%ac%20rates
```

Note: the Unicode pound sign character U+00A3 was encoded into the octet sequence C2 A3 using the UTF-8 character encoding, then percent-encoded. Likewise, the Unicode euro sign character U+20AC was encoded into the octet sequence E2 82 AC, then percent-encoded. Also note that HEXDIG allows both lowercase and uppercase characters, so recipients must understand both, and that the language information is optional, while the character encoding is not.

3.3. Language Specification in Encoded Words

Section 5 of [RFC2231] extends the encoding defined in [RFC2047] to also support language specification in encoded words. RFC 2616, the now-obsolete HTTP/1.1 specification, did refer to RFC 2047 ([RFC2616], Section 2.2). However, it wasn't clear to which header field it applied. Consequently, the current revision of the HTTP/1.1 specification has deprecated use of the encoding forms defined in RFC 2047 (see Section 3.2.4 of [RFC7230]).

Thus, this specification does not include this feature.

4. Guidelines for Usage in HTTP Header Field Definitions

Specifications of HTTP header fields that use the extensions defined in Section 3.2 ought to clearly state that. A simple way to achieve this is to normatively reference this specification, and to include the ext-value production into the ABNF for that header field.

For instance:

```
foo-header  = "foo" LWSP ":" LWSP token ";" LWSP title-param
title-param = "title" LWSP "=" LWSP value
            / "title*" LWSP "=" LWSP ext-value
ext-value   = <see RFC 5987, Section 3.2>
```

Note: The Parameter Value Continuation feature defined in Section 3 of [RFC2231] makes it impossible to have multiple instances of extended parameters with identical parmname components, as the processing of continuations would become ambiguous. Thus, specifications using this extension are advised to disallow this case for compatibility with RFC 2231.

Note: This specification does not automatically assign a new interpretation to parameter names ending in an asterisk. As pointed out above, it's up to the specification for the non-extended parameter to "opt in" to the syntax defined here. That being said, some existing implementations are known to automatically switch to the use of this notation when a parameter name ends with an asterisk, thus using parameter names ending in an asterisk for something else is likely to cause interoperability problems.

4.1. When to Use the Extension

Section 4.2 of [RFC2277] requires that protocol elements containing human-readable text are able to carry language information. Thus, the ext-value production ought to be always used when the parameter value is of textual nature and its language is known.

Furthermore, the extension ought to also be used whenever the parameter value needs to carry characters not present in the US-ASCII ([USASCII]) coded character set (note that it would be unacceptable to define a new parameter that would be restricted to a subset of the Unicode character set).

4.2. Error Handling

Header field specifications need to define whether multiple instances of parameters with identical parmname components are allowed, and how

they should be processed. This specification suggests that a parameter using the extended syntax takes precedence. This would allow producers to use both formats without breaking recipients that do not understand the extended syntax yet.

Example:

```
foo: bar; title="EURO exchange rates";  
      title*=utf-8''%e2%82%ac%20exchange%20rates
```

In this case, the sender provides an ASCII version of the title for legacy recipients, but also includes an internationalized version for recipients understanding this specification -- the latter obviously ought to prefer the new syntax over the old one.

Note: at the time of this writing, many implementations failed to ignore the form they do not understand, or prioritize the ASCII form although the extended syntax was present.

5. Security Considerations

The format described in this document makes it possible to transport non-ASCII characters, and thus enables character "spoofing" scenarios, in which a displayed value appears to be something other than it is.

Furthermore, there are known attack scenarios relating to decoding UTF-8.

See Section 10 of [RFC3629] for more information on both topics.

In addition, the extension specified in this document makes it possible to transport multiple language variants for a single parameter, and such use might allow spoofing attacks, where different language versions of the same parameter are not equivalent. Whether this attack is useful as an attack depends on the parameter specified.

6. IANA Considerations

There are no IANA Considerations related to this specification.

7. Acknowledgements

Thanks to Martin Duerst and Frank Ellermann for help figuring out ABNF details, to Graham Klyne and Alexey Melnikov for general review, to Chris Newman for pointing out an RFC 2231 incompatibility, and to Benjamin Carlyle, Roar Lauritzsen, Eric Lawrence, and James Manger

for implementer's feedback.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [RFC2978] Freed, N. and J. Postel, "IANA Charset Registration Procedures", BCP 19, RFC 2978, October 2000.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [RFC5646] Phillips, A., Ed. and M. Davis, Ed., "Tags for Identifying Languages", BCP 47, RFC 5646, September 2009.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, June 2014.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, June 2014.
- [USASCII] American National Standards Institute, "Coded Character Set -- 7-bit American Standard Code for Information Interchange", ANSI X3.4, 1986.

8.2. Informative References

- [Err1912] RFC Errata, "Errata ID 1912, RFC 2978", <<http://www.rfc-editor.org>>.

- [ISO-8859-1] International Organization for Standardization, "Information technology -- 8-bit single-byte coded graphic character sets -- Part 1: Latin alphabet No. 1", ISO/IEC 8859-1:1998, 1998.
- [RFC2045] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, November 1996.
- [RFC2047] Moore, K., "MIME (Multipurpose Internet Mail Extensions) Part Three: Message Header Extensions for Non-ASCII Text", RFC 2047, November 1996.
- [RFC2231] Freed, N. and K. Moore, "MIME Parameter Value and Encoded Word Extensions: Character Sets, Languages, and Continuations", RFC 2231, November 1997.
- [RFC2277] Alvestrand, H., "IETF Policy on Character Sets and Languages", BCP 18, RFC 2277, January 1998.
- [RFC2388] Masinter, L., "Returning Values from Forms: multipart/form-data", RFC 2388, August 1998.
- [RFC5987] Reschke, J., "Character Set and Language Encoding for Hypertext Transfer Protocol (HTTP) Header Field Parameters", RFC 5987, August 2010.
- [RFC5988] Nottingham, M., "Web Linking", RFC 5988, October 2010.
- [RFC6266] Reschke, J., "Use of the Content-Disposition Header Field in the Hypertext Transfer Protocol (HTTP)", RFC 6266, June 2011.
- [RFC6365] Hoffman, P. and J. Klensin, "Terminology Used in Internationalization in the IETF", BCP 166, RFC 6365, September 2011.

URIs

- [1] <<mailto:ietf-http-wg@w3.org>>
- [2] <<mailto:ietf-http-wg-request@w3.org?subject=subscribe>>

Appendix A. Changes from RFC 5987

This section summarizes the changes compared to [RFC5987]:

- o The document title was changed to "Indicating Character Encoding and Language for HTTP Header Field Parameters".
- o The requirement to support the "ISO-8859-1" encoding was removed.

Appendix B. Implementation Report

The encoding defined in this document currently is used for two different HTTP header fields:

- o "Content-Disposition", defined in [RFC6266], and
- o "Link", defined in [RFC5988].

As the encoding is a profile/clarification of the one defined in [RFC2231] in 1997, many user agents already supported it for use in "Content-Disposition" when [RFC5987] got published.

Since the publication of [RFC5987], three more popular desktop user agents have added support for this encoding; see <<http://purl.org/NET/http/content-disposition-tests#encoding-2231-char>> for details. At this time, the current versions of all major desktop user agents support it.

Note that the implementation in Internet Explorer 9 does not support the ISO-8859-1 character encoding; this document revision acknowledges that UTF-8 is sufficient for expressing all code points, and removes the requirement to support ISO-8859-1.

The "Link" header field, on the other hand, was only recently specified in [RFC5988]. At the time of this writing, no shipping User Agent except Firefox supported the "title*" parameter (starting with release 15).

Appendix C. Change Log (to be removed by RFC Editor before publication)

C.1. Since RFC5987

Only editorial changes for the purpose of starting the revision process (obs5987).

C.2. Since draft-reschke-rfc5987bis-00

Resolved issues "iso-8859-1" and "title" (title simplified). Added and resolved issue "historic5987".

C.3. Since draft-reschke-rfc5987bis-01

Added issues "httpbis", "parmsyntax", "terminology" and "valuesyntax". Closed issue "impls".

C.4. Since draft-reschke-rfc5987bis-02

Resolved issue "terminology".

C.5. Since draft-reschke-rfc5987bis-03

In Section 3.2, pull historical notes into a separate subsection. Resolved issues "valuesyntax" and "parmsyntax".

C.6. Since draft-reschke-rfc5987bis-04

Update status of Firefox support in HTTP Link Header field.

C.7. Since draft-reschke-rfc5987bis-05

Update status of Firefox support in HTTP Link Header field.

C.8. Since draft-reschke-rfc5987bis-06

Update status with respect to Safari 6.

Started work on update with respect to RFC 723x.

Appendix D. Open issues (to be removed by RFC Editor prior to publication)

D.1. edit

Type: edit

julian.reschke@greenbytes.de (2011-04-15): Umbrella issue for editorial fixes/enhancements.

D.2. httpbis

Type: edit

julian.reschke@greenbytes.de (2011-09-17): The document refers normatively to RFC 2616. Should it continue to do so, or should we wait for HTTPbis? This may affect edge case in the ABNF, such as the definition of linear white space or the characters allowed in "token".

Author's Address

Julian F. Reschke
greenbytes GmbH
Hafenweg 16
Muenster, NW 48155
Germany

EMail: julian.reschke@greenbytes.de
URI: <http://greenbytes.de/tech/webdav/>

