

Workgroup: QUIC
Internet-Draft: draft-ietf-quic-http-28
Published: 20 May 2020
Intended Status: Standards Track
Expires: 21 November 2020
Authors: M. Bishop, Ed.
Akamai

Hypertext Transfer Protocol Version 3 (HTTP/3)

Abstract

The QUIC transport protocol has several features that are desirable in a transport for HTTP, such as stream multiplexing, per-stream flow control, and low-latency connection establishment. This document describes a mapping of HTTP semantics over QUIC. This document also identifies HTTP/2 features that are subsumed by QUIC, and describes how HTTP/2 extensions can be ported to HTTP/3.

Note to Readers

Discussion of this draft takes place on the QUIC working group mailing list (quic@ietf.org), which is archived at https://mailarchive.ietf.org/arch/search/?email_list=quic.

Working Group information can be found at <https://github.com/quicwg>; source code and issues list for this draft can be found at <https://github.com/quicwg/base-drafts/labels/-http>.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

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

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

This Internet-Draft will expire on 21 November 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Introduction](#)
 - [1.1. Prior versions of HTTP](#)
 - [1.2. Delegation to QUIC](#)
- [2. HTTP/3 Protocol Overview](#)
 - [2.1. Document Organization](#)
 - [2.2. Conventions and Terminology](#)
- [3. Connection Setup and Management](#)
 - [3.1. Draft Version Identification](#)
 - [3.2. Discovering an HTTP/3 Endpoint](#)
 - [3.2.1. HTTP Alternative Services](#)
 - [3.2.2. Other Schemes](#)
 - [3.3. Connection Establishment](#)
 - [3.4. Connection Reuse](#)
- [4. HTTP Request Lifecycle](#)
 - [4.1. HTTP Message Exchanges](#)
 - [4.1.1. Field Formatting and Compression](#)
 - [4.1.2. Request Cancellation and Rejection](#)
 - [4.1.3. Malformed Requests and Responses](#)
 - [4.2. The CONNECT Method](#)
 - [4.3. HTTP Upgrade](#)

- [4.4. Server Push](#)
- [5. Connection Closure](#)
 - [5.1. Idle Connections](#)
 - [5.2. Connection Shutdown](#)
 - [5.3. Immediate Application Closure](#)
 - [5.4. Transport Closure](#)
- [6. Stream Mapping and Usage](#)
 - [6.1. Bidirectional Streams](#)
 - [6.2. Unidirectional Streams](#)
 - [6.2.1. Control Streams](#)
 - [6.2.2. Push Streams](#)
 - [6.2.3. Reserved Stream Types](#)
- [7. HTTP Framing Layer](#)
 - [7.1. Frame Layout](#)
 - [7.2. Frame Definitions](#)
 - [7.2.1. DATA](#)
 - [7.2.2. HEADERS](#)
 - [7.2.3. CANCEL_PUSH](#)
 - [7.2.4. SETTINGS](#)
 - [7.2.5. PUSH_PROMISE](#)
 - [7.2.6. GOAWAY](#)
 - [7.2.7. MAX_PUSH_ID](#)
 - [7.2.8. Reserved Frame Types](#)
- [8. Error Handling](#)
 - [8.1. HTTP/3 Error Codes](#)
- [9. Extensions to HTTP/3](#)

[10. Security Considerations](#)

[10.1. Server Authority](#)

[10.2. Cross-Protocol Attacks](#)

[10.3. Intermediary Encapsulation Attacks](#)

[10.4. Cacheability of Pushed Responses](#)

[10.5. Denial-of-Service Considerations](#)

[10.5.1. Limits on Field Section Size](#)

[10.5.2. CONNECT Issues](#)

[10.6. Use of Compression](#)

[10.7. Padding and Traffic Analysis](#)

[10.8. Frame Parsing](#)

[10.9. Early Data](#)

[10.10. Migration](#)

[10.11. Privacy Considerations](#)

[11. IANA Considerations](#)

[11.1. Registration of HTTP/3 Identification String](#)

[11.2. New Registries](#)

[11.2.1. Frame Types](#)

[11.2.2. Settings Parameters](#)

[11.2.3. Error Codes](#)

[11.2.4. Stream Types](#)

[12. References](#)

[12.1. Normative References](#)

[12.2. Informative References](#)

[Appendix A. Considerations for Transitioning from HTTP/2](#)

[A.1. Streams](#)

[A.2. HTTP Frame Types](#)

[A.2.1. Prioritization Differences](#)

[A.2.2. Field Compression Differences](#)

[A.2.3. Guidance for New Frame Type Definitions](#)

[A.2.4. Mapping Between HTTP/2 and HTTP/3 Frame Types](#)

[A.3. HTTP/2 SETTINGS Parameters](#)

[A.4. HTTP/2 Error Codes](#)

[A.4.1. Mapping Between HTTP/2 and HTTP/3 Errors](#)

[Appendix B. Change Log](#)

[B.1. Since draft-ietf-quic-http-27](#)

[B.2. Since draft-ietf-quic-http-26](#)

[B.3. Since draft-ietf-quic-http-25](#)

[B.4. Since draft-ietf-quic-http-24](#)

[B.5. Since draft-ietf-quic-http-23](#)

[B.6. Since draft-ietf-quic-http-22](#)

[B.7. Since draft-ietf-quic-http-21](#)

[B.8. Since draft-ietf-quic-http-20](#)

[B.9. Since draft-ietf-quic-http-19](#)

[B.10. Since draft-ietf-quic-http-18](#)

[B.11. Since draft-ietf-quic-http-17](#)

[B.12. Since draft-ietf-quic-http-16](#)

[B.13. Since draft-ietf-quic-http-15](#)

[B.14. Since draft-ietf-quic-http-14](#)

[B.15. Since draft-ietf-quic-http-13](#)

[B.16. Since draft-ietf-quic-http-12](#)

[B.17. Since draft-ietf-quic-http-11](#)

[B.18. Since draft-ietf-quic-http-10](#)

[B.19. Since draft-ietf-quic-http-09](#)

[B.20. Since draft-ietf-quic-http-08](#)

[B.21. Since draft-ietf-quic-http-07](#)

[B.22. Since draft-ietf-quic-http-06](#)

[B.23. Since draft-ietf-quic-http-05](#)

[B.24. Since draft-ietf-quic-http-04](#)

[B.25. Since draft-ietf-quic-http-03](#)

[B.26. Since draft-ietf-quic-http-02](#)

[B.27. Since draft-ietf-quic-http-01](#)

[B.28. Since draft-ietf-quic-http-00](#)

[B.29. Since draft-shade-quic-http2-mapping-00](#)

[Acknowledgements](#)

[Author's Address](#)

1. Introduction

HTTP semantics [[SEMANTICS](#)] are used for a broad range of services on the Internet. These semantics have most commonly been used with two different TCP mappings, HTTP/1.1 and HTTP/2. HTTP/3 supports the same semantics over a new transport protocol, QUIC.

1.1. Prior versions of HTTP

HTTP/1.1 [[HTTP11](#)] is a TCP mapping which uses whitespace-delimited text fields to convey HTTP messages. While these exchanges are human-readable, using whitespace for message formatting leads to parsing complexity and motivates tolerance of variant behavior. Because each connection can transfer only a single HTTP request or response at a time in each direction, multiple parallel TCP connections are often used, reducing the ability of the congestion controller to effectively manage traffic between endpoints.

HTTP/2 [[HTTP2](#)] introduced a binary framing and multiplexing layer to improve latency without modifying the transport layer. However, because the parallel nature of HTTP/2's multiplexing is not visible to TCP's loss recovery mechanisms, a lost or reordered packet causes

all active transactions to experience a stall regardless of whether that transaction was directly impacted by the lost packet.

1.2. Delegation to QUIC

The QUIC transport protocol incorporates stream multiplexing and per-stream flow control, similar to that provided by the HTTP/2 framing layer. By providing reliability at the stream level and congestion control across the entire connection, it has the capability to improve the performance of HTTP compared to a TCP mapping. QUIC also incorporates TLS 1.3 [\[TLS13\]](#) at the transport layer, offering comparable security to running TLS over TCP, with the improved connection setup latency of TCP Fast Open [\[TFO\]](#).

This document defines a mapping of HTTP semantics over the QUIC transport protocol, drawing heavily on the design of HTTP/2. While delegating stream lifetime and flow control issues to QUIC, a similar binary framing is used on each stream. Some HTTP/2 features are subsumed by QUIC, while other features are implemented atop QUIC.

QUIC is described in [\[QUIC-TRANSPORT\]](#). For a full description of HTTP/2, see [\[HTTP2\]](#).

2. HTTP/3 Protocol Overview

HTTP/3 provides a transport for HTTP semantics using the QUIC transport protocol and an internal framing layer similar to HTTP/2.

Once a client knows that an HTTP/3 server exists at a certain endpoint, it opens a QUIC connection. QUIC provides protocol negotiation, stream-based multiplexing, and flow control. Discovery of an HTTP/3 endpoint is described in [Section 3.2](#).

Within each stream, the basic unit of HTTP/3 communication is a frame ([Section 7.2](#)). Each frame type serves a different purpose. For example, HEADERS and DATA frames form the basis of HTTP requests and responses ([Section 4.1](#)).

Multiplexing of requests is performed using the QUIC stream abstraction, described in Section 2 of [\[QUIC-TRANSPORT\]](#). Each request-response pair consumes a single QUIC stream. Streams are independent of each other, so one stream that is blocked or suffers packet loss does not prevent progress on other streams.

Server push is an interaction mode introduced in HTTP/2 [\[HTTP2\]](#) which permits a server to push a request-response exchange to a client in anticipation of the client making the indicated request. This trades off network usage against a potential latency gain.

Several HTTP/3 frames are used to manage server push, such as PUSH_PROMISE, MAX_PUSH_ID, and CANCEL_PUSH.

As in HTTP/2, request and response fields are compressed for transmission. Because HPACK [HPACK] relies on in-order transmission of compressed field sections (a guarantee not provided by QUIC), HTTP/3 replaces HPACK with QPACK [QPACK]. QPACK uses separate unidirectional streams to modify and track field table state, while encoded field sections refer to the state of the table without modifying it.

2.1. Document Organization

The following sections provide a detailed overview of the connection lifecycle and key concepts:

- *Connection Setup and Management ([Section 3](#)) covers how an HTTP/3 endpoint is discovered and a connection is established.
- *HTTP Request Lifecycle ([Section 4](#)) describes how HTTP semantics are expressed using frames.
- *Connection Closure ([Section 5](#)) describes how connections are terminated, either gracefully or abruptly.

The details of the wire protocol and interactions with the transport are described in subsequent sections:

- *Stream Mapping and Usage ([Section 6](#)) describes the way QUIC streams are used.
- *HTTP Framing Layer ([Section 7](#)) describes the frames used on most streams.
- *Error Handling ([Section 8](#)) describes how error conditions are handled and expressed, either on a particular stream or for the connection as a whole.

Additional resources are provided in the final sections:

- *Extensions to HTTP/3 ([Section 9](#)) describes how new capabilities can be added in future documents.
- *A more detailed comparison between HTTP/2 and HTTP/3 can be found in [Appendix A](#).

2.2. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and

"OPTIONAL" in this document are to be interpreted as described in BCP 14 [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

Field definitions are given in Augmented Backus-Naur Form (ABNF), as defined in [[RFC5234](#)].

This document uses the variable-length integer encoding from [[QUIC-TRANSPORT](#)].

The following terms are used:

abort: An abrupt termination of a connection or stream, possibly due to an error condition.

client: The endpoint that initiates an HTTP/3 connection. Clients send HTTP requests and receive HTTP responses.

connection: A transport-layer connection between two endpoints, using QUIC as the transport protocol.

connection error: An error that affects the entire HTTP/3 connection.

endpoint: Either the client or server of the connection.

frame: The smallest unit of communication on a stream in HTTP/3, consisting of a header and a variable-length sequence of bytes structured according to the frame type. Protocol elements called "frames" exist in both this document and [[QUIC-TRANSPORT](#)]. Where frames from [[QUIC-TRANSPORT](#)] are referenced, the frame name will be prefaced with "QUIC." For example, "QUIC CONNECTION_CLOSE frames." References without this preface refer to frames defined in [Section 7.2](#).

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/3 connection. Servers receive HTTP requests and send HTTP responses.

stream: A bidirectional or unidirectional bytestream provided by the QUIC transport.

stream error: An error on the individual HTTP/3 stream.

The term "payload body" is defined in Section 6.3.3 of [\[SEMANTICS\]](#).

Finally, the terms "gateway", "intermediary", "proxy", and "tunnel" are defined in Section 2.2 of [\[SEMANTICS\]](#). Intermediaries act as both client and server at different times.

3. Connection Setup and Management

3.1. Draft Version Identification

RFC Editor's Note: Please remove this section prior to publication of a final version of this document.

HTTP/3 uses the token "h3" to identify itself in ALPN and Alt-Svc. Only implementations of the final, published RFC can identify themselves as "h3". Until such an RFC exists, implementations MUST NOT identify themselves using this string.

Implementations of draft versions of the protocol MUST add the string "-" and the corresponding draft number to the identifier. For example, draft-ietf-quit-http-01 is identified using the string "h3-01".

Draft versions MUST use the corresponding draft transport version as their transport. For example, the application protocol defined in draft-ietf-quit-http-25 uses the transport defined in draft-ietf-quit-transport-25.

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 based on draft-ietf-quit-http-09 which reserves an extra stream for unsolicited transmission of 1980s pop music might identify itself as "h3-09-rickroll". Note that any label MUST conform to the "token" syntax defined in Section 4.4.1.1 of [\[SEMANTICS\]](#). Experimenters are encouraged to coordinate their experiments on the quit@ietf.org mailing list.

3.2. Discovering an HTTP/3 Endpoint

HTTP relies on the notion of an authoritative response: a response that has been determined to be the most appropriate response for that request given the state of the target resource at the time of response message origination by (or at the direction of) the origin server identified within the target URI. Locating an authoritative server for an HTTP URL is discussed in Section 5.4 of [\[SEMANTICS\]](#).

The "https" scheme associates authority with possession of a certificate that the client considers to be trustworthy for the host identified by the authority component of the URL. If a server

presents a certificate and proof that it controls the corresponding private key, then a client will accept a secured connection to that server as being authoritative for all origins with the "https" scheme and a host identified in the certificate.

A client MAY attempt access to a resource with an "https" URI by resolving the host identifier to an IP address, establishing a QUIC connection to that address on the indicated port, and sending an HTTP/3 request message targeting the URI to the server over that secured connection.

Connectivity problems (e.g., blocking UDP) can result in QUIC connection establishment failure; clients SHOULD attempt to use TCP-based versions of HTTP in this case.

Servers MAY serve HTTP/3 on any UDP port; an alternative service advertisement always includes an explicit port, and URLs contain either an explicit port or a default port associated with the scheme.

3.2.1. HTTP Alternative Services

An HTTP origin advertises the availability of an equivalent HTTP/3 endpoint via the Alt-Svc HTTP response header field or the HTTP/2 ALTSVC frame ([[ALTSVC](#)]), using the ALPN token defined in [Section 3.3](#).

For example, an origin could indicate in an HTTP response that HTTP/3 was available on UDP port 50781 at the same hostname by including the following header field:

```
Alt-Svc: h3=":50781"
```

On receipt of an Alt-Svc record indicating HTTP/3 support, a client MAY attempt to establish a QUIC connection to the indicated host and port and, if successful, send HTTP requests using the mapping described in this document.

3.2.2. Other Schemes

Although HTTP is independent of the transport protocol, the "http" scheme associates authority with the ability to receive TCP connections on the indicated port of whatever host is identified within the authority component. Because HTTP/3 does not use TCP, HTTP/3 cannot be used for direct access to the authoritative server for a resource identified by an "http" URI. However, protocol extensions such as [[ALTSVC](#)] permit the authoritative server to identify other services which are also authoritative and which might be reachable over HTTP/3.

Prior to making requests for an origin whose scheme is not "https", the client MUST ensure the server is willing to serve that scheme. If the client intends to make requests for an origin whose scheme is "http", this means that it MUST obtain a valid http-opportunistic response for the origin as described in [\[RFC8164\]](#) prior to making any such requests. Other schemes might define other mechanisms.

3.3. Connection Establishment

HTTP/3 relies on QUIC version 1 as the underlying transport. The use of other QUIC transport versions with HTTP/3 MAY be defined by future specifications.

QUIC version 1 uses TLS version 1.3 or greater as its handshake protocol. HTTP/3 clients MUST support a mechanism to indicate the target host to the server during the TLS handshake. If the server is identified by a DNS name, clients MUST send the Server Name Indication (SNI) [\[RFC6066\]](#) TLS extension unless an alternative mechanism to indicate the target host is used.

QUIC connections are established as described in [\[QUIC-TRANSPORT\]](#). During connection establishment, HTTP/3 support is indicated by selecting the ALPN token "h3" in the TLS handshake. Support for other application-layer protocols MAY be offered in the same handshake.

While connection-level options pertaining to the core QUIC protocol are set in the initial crypto handshake, HTTP/3-specific settings are conveyed in the SETTINGS frame. After the QUIC connection is established, a SETTINGS frame ([Section 7.2.4](#)) MUST be sent by each endpoint as the initial frame of their respective HTTP control stream; see [Section 6.2.1](#).

3.4. Connection Reuse

HTTP/3 connections are persistent across multiple requests. For best performance, it is expected that 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.

Once a connection exists to a server endpoint, this connection MAY be reused for requests with multiple different URI authority components. In general, a server is considered authoritative for all URIs with the "https" scheme for which the hostname in the URI is present in the authenticated certificate provided by the server, either as the CN field of the certificate subject or as a `dnsName` in the `subjectAltName` field of the certificate; see [\[RFC6125\]](#). For a host that is an IP address, the client MUST verify that the address appears as an `iPAddress` in the `subjectAltName` field of the

certificate. If the hostname or address is not present in the certificate, the client **MUST NOT** consider the server authoritative for origins containing that hostname or address. See Section 5.4 of [\[SEMANTICS\]](#) for more detail on authoritative access.

Clients **SHOULD NOT** open more than one HTTP/3 connection to a given host and port pair, where the host is derived from a URI, a selected alternative service [\[ALTSVC\]](#), or a configured proxy. A client **MAY** open multiple connections to the same IP address and UDP port using different transport or TLS configurations 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 HTTP/3 session, the terminating endpoint **SHOULD** first send a GOAWAY frame ([Section 5.2](#)) so that both endpoints can reliably determine whether previously sent frames have been processed and gracefully complete or terminate any necessary remaining tasks.

A server that does not wish clients to reuse connections for a particular origin 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 of [\[HTTP2\]](#).

4. HTTP Request Lifecycle

4.1. HTTP Message Exchanges

A client sends an HTTP request on a client-initiated bidirectional QUIC stream. A client **MUST** send only a single request on a given stream. A server sends zero or more interim HTTP responses on the same stream as the request, followed by a single final HTTP response, as detailed below.

Pushed responses are sent on a server-initiated unidirectional QUIC stream; see [Section 6.2.2](#). A server sends zero or more interim HTTP responses, followed by a single final HTTP response, in the same manner as a standard response. Push is described in more detail in [Section 4.4](#).

On a given stream, receipt of multiple requests or receipt of an additional HTTP response following a final HTTP response **MUST** be treated as malformed ([Section 4.1.3](#)).

An HTTP message (request or response) consists of:

1. the header field section (see Section 4 of [\[SEMANTICS\]](#)), sent as a single HEADERS frame (see [Section 7.2.2](#)),

2. optionally, the payload body, if present (see Section 6.3.3 of [\[SEMANTICS\]](#)), sent as a series of DATA frames (see [Section 7.2.1](#)),
3. optionally, the trailer field section, if present (see Section 4.6 of [\[SEMANTICS\]](#)), sent as a single HEADERS frame.

Receipt of an invalid sequence of frames MUST be treated as a connection error of type H3_FRAME_UNEXPECTED ([Section 8](#)). In particular, a DATA frame before any HEADERS frame, or a HEADERS or DATA frame after the trailing HEADERS frame is considered invalid.

A server MAY send one or more PUSH_PROMISE frames (see [Section 7.2.5](#)) before, after, or interleaved with the frames of a response message. These PUSH_PROMISE frames are not part of the response; see [Section 4.4](#) for more details. These frames are not permitted in pushed responses; a pushed response which includes PUSH_PROMISE frames MUST be treated as a connection error of type H3_FRAME_UNEXPECTED.

Frames of unknown types ([Section 9](#)), including reserved frames ([Section 7.2.8](#)) MAY be sent on a request or push stream before, after, or interleaved with other frames described in this section.

The HEADERS and PUSH_PROMISE frames might reference updates to the QPACK dynamic table. While these updates are not directly part of the message exchange, they must be received and processed before the message can be consumed. See [Section 4.1.1](#) for more details.

The "chunked" transfer encoding defined in Section 7.1 of [\[HTTP11\]](#) MUST NOT be used.

A response MAY consist of multiple messages when and only when one or more informational responses (1xx; see Section 9.2 of [\[SEMANTICS\]](#)) precede a final response to the same request. Interim responses do not contain a payload body or trailers.

An HTTP request/response exchange fully consumes a client-initiated bidirectional QUIC stream. After sending a request, a client MUST close the stream for sending. Unless using the CONNECT method (see [Section 4.2](#)), clients MUST NOT make stream closure dependent on receiving a response to their request. After sending a final response, the server MUST close the stream for sending. At this point, the QUIC stream is fully closed.

When a stream is closed, this indicates the end of an HTTP message. Because some messages are large or unbounded, endpoints SHOULD begin processing partial HTTP messages once enough of the message has been received to make progress. If a client stream terminates without enough of the HTTP message to provide a complete response, the

server SHOULD abort its response with the error code H3_REQUEST_INCOMPLETE.

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 the server does not need to receive the remainder of the request, it MAY abort reading the request stream, send a complete response, and cleanly close the sending part of the stream. The error code H3_NO_ERROR SHOULD be used when requesting that the client stop sending on the request stream. Clients MUST NOT discard complete responses as a result of having their request terminated abruptly, though clients can always discard responses at their discretion for other reasons. If the server sends a partial or complete response but does not abort reading, clients SHOULD continue sending the body of the request and close the stream normally.

4.1.1. Field Formatting and Compression

HTTP messages carry metadata as a series of key-value pairs, called HTTP fields. For a listing of registered HTTP fields, see the "Hypertext Transfer Protocol (HTTP) Field Name Registry" maintained at <https://www.iana.org/assignments/http-fields/>.

As in previous versions of HTTP, field names are strings containing a subset of ASCII characters that are compared in a case-insensitive fashion. Properties of HTTP field names and values are discussed in more detail in Section 4.3 of [SEMANTICS]. As in HTTP/2, characters in field names MUST be converted to lowercase prior to their encoding. A request or response containing uppercase characters in field names MUST be treated as malformed ([Section 4.1.3](#)).

Like HTTP/2, HTTP/3 does not use the Connection header field to indicate connection-specific fields; in this protocol, connection-specific metadata is conveyed by other means. An endpoint MUST NOT generate an HTTP/3 field section containing connection-specific fields; any message containing connection-specific fields MUST be treated as malformed ([Section 4.1.3](#)).

The only exception to this is the TE header field, which MAY be present in an HTTP/3 request header; 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/3 will need to remove any fields nominated by the Connection field, along with the Connection field itself. Such intermediaries SHOULD also remove other connection-specific fields, such as Keep-Alive, Proxy-Connection, Transfer-Encoding, and Upgrade, even if they are not nominated by the Connection field.

4.1.1.1. Pseudo-Header Fields

Like HTTP/2, HTTP/3 employs a series of pseudo-header fields where the field name begins with the ':' character (ASCII 0x3a). These pseudo-header fields convey the target URI, the method of the request, and the status code for the response.

Pseudo-header fields are not HTTP fields. Endpoints **MUST NOT** generate pseudo-header fields other than those defined in this document, except as negotiated via an extension; see [Section 9](#).

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 4.1.3](#)).

All pseudo-header fields **MUST** appear in the header field section before regular header fields. Any request or response that contains a pseudo-header field that appears in a header field section after a regular header field **MUST** be treated as malformed ([Section 4.1.3](#)).

The following pseudo-header fields are defined for requests:

":method": Contains the HTTP method (Section 7 of [\[SEMANTICS\]](#))

":scheme": Contains the scheme portion of the target URI (Section 3.1 of [\[RFC3986\]](#))

":scheme" is not restricted to "http" and "https" schemed URIs. A proxy or gateway can translate requests for non-HTTP schemes, enabling the use of HTTP to interact with non-HTTP services.

":authority": Contains the authority portion of the target URI (Section 3.2 of [\[RFC3986\]](#)). The authority **MUST NOT** include the deprecated "userinfo" subcomponent for "http" or "https" schemed 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 Section 3.2 of [\[HTTP11\]](#). Clients that generate HTTP/3 requests directly **SHOULD** use the ":authority" pseudo-header field instead of the Host field. An intermediary that converts an HTTP/3 request to HTTP/1.1 **MUST** create a Host field if one is not present in a request by copying the value of the ":authority" pseudo-header field.

":path":

Contains the path and query parts of the target URI (the "path-absolute" production and optionally a '?' character followed by the "query" production; see Sections 3.3 and 3.4 of [\[URI\]](#)). 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 Section 3.2.4 of [\[HTTP11\]](#).

All HTTP/3 requests MUST include exactly one value for the ":method", ":scheme", and ":path" pseudo-header fields, unless it is a CONNECT request; see [Section 4.2](#).

If the ":scheme" pseudo-header field identifies a scheme which has a mandatory authority component (including "http" and "https"), the request MUST contain either an ":authority" pseudo-header field or a "Host" header field. If these fields are present, they MUST NOT be empty. If both fields are present, they MUST contain the same value. If the scheme does not have a mandatory authority component and none is provided in the request target, the request MUST NOT contain the ":authority" pseudo-header and "Host" header fields.

An HTTP request that omits mandatory pseudo-header fields or contains invalid values for those pseudo-header fields is malformed ([Section 4.1.3](#)).

HTTP/3 does not define a way to carry the version identifier that is included in the HTTP/1.1 request line.

For responses, a single ":status" pseudo-header field is defined that carries the HTTP status code; see Section 9 of [\[SEMANTICS\]](#). This pseudo-header field MUST be included in all responses; otherwise, the response is malformed ([Section 4.1.3](#)).

HTTP/3 does not define a way to carry the version or reason phrase that is included in an HTTP/1.1 status line.

4.1.1.2. Field Compression

HTTP/3 uses QPACK field compression as described in [\[QPACK\]](#), a variation of HPACK which allows the flexibility to avoid compression-induced head-of-line blocking. See that document for additional details.

To allow for better compression efficiency, the "Cookie" field [[RFC6265](#)] MAY be split into separate field lines, each with one or more cookie-pairs, before compression. If a decompressed field section contains multiple cookie field lines, 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 context other than HTTP/2 or HTTP/3, such as an HTTP/1.1 connection, or a generic HTTP server application.

4.1.1.3. Header Size Constraints

An HTTP/3 implementation MAY impose a limit on the maximum size of the message header it will accept on an individual HTTP message. A server that receives a larger header section 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 size of a field list is calculated based on the uncompressed size of fields, including the length of the name and value in bytes plus an overhead of 32 bytes for each field.

If an implementation wishes to advise its peer of this limit, it can be conveyed as a number of bytes in the SETTINGS_MAX_FIELD_SECTION_SIZE parameter. An implementation which has received this parameter SHOULD NOT send an HTTP message header which exceeds the indicated size, as the peer will likely refuse to process it. However, because this limit is applied at each hop, messages below this limit are not guaranteed to be accepted.

4.1.1.2. Request Cancellation and Rejection

Clients can cancel requests by resetting and aborting the request stream with an error code of H3_REQUEST_CANCELLED ([Section 8.1](#)). When the client aborts reading a response, it indicates that this response is no longer of interest. Implementations SHOULD cancel requests by abruptly terminating any directions of a stream that are still open.

When the server rejects a request without performing any application processing, it SHOULD abort its response stream with the error code H3_REQUEST_REJECTED. 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. The client can treat requests rejected by the server as though they had never been sent at all, thereby allowing them to be retried later on a new connection. Servers MUST NOT use the H3_REQUEST_REJECTED error code for requests which were partially or fully processed. When a server abandons a response after partial processing, it SHOULD abort its response stream with the error code H3_REQUEST_CANCELLED.

When a client resets a request with the error code `H3_REQUEST_CANCELLED`, a server MAY abruptly terminate the response using the error code `H3_REQUEST_REJECTED` if no processing was performed. Clients MUST NOT use the `H3_REQUEST_REJECTED` error code, except when a server has requested closure of the request stream with this error code.

If a stream is cancelled after receiving a complete response, the client MAY ignore the cancellation and use the response. However, if a stream is cancelled after receiving a partial response, the response SHOULD NOT be used. Automatically retrying such requests is not possible, unless this is otherwise permitted (e.g., idempotent actions like GET, PUT, or DELETE).

4.1.3. Malformed Requests and Responses

A malformed request or response is one that is an otherwise valid sequence of frames but is invalid due to:

- *the presence of prohibited fields or pseudo-header fields,
- *the absence of mandatory pseudo-header fields,
- *invalid values for pseudo-header fields,
- *pseudo-header fields after fields,
- *an invalid sequence of HTTP messages,
- *the inclusion of uppercase field names, or
- *the inclusion of invalid characters in field names or values

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 Section 6.3.3 of [\[SEMANTICS\]](#) can have a non-zero content-length 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 8](#)) of type `H3_GENERAL_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.

4.2. The CONNECT Method

The CONNECT method requests that the recipient establish a tunnel to the destination origin server identified by the request-target (Section 3.2 of [\[HTTP11\]](#)). It 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/1.x, CONNECT is used to convert an entire HTTP connection into a tunnel to a remote host. In HTTP/2 and HTTP/3, the CONNECT method is used to establish a tunnel over a single stream.

A CONNECT request MUST be constructed as follows:

- *The ":method" pseudo-header field is set to "CONNECT"
- *The ":scheme" and ":path" pseudo-header fields are omitted
- *The ":authority" pseudo-header field contains the host and port to connect to (equivalent to the authority-form of the request-target of CONNECT requests; see Section 5.3 of [\[HTTP11\]](#))

The request stream remains open at the end of the request to carry the data to be transferred. A CONNECT request that does not conform to these restrictions is malformed; see [Section 4.1.3](#).

A proxy that supports CONNECT establishes a TCP connection ([\[RFC0793\]](#)) to the server identified in the ":authority" pseudo-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 Section 9.3 of [\[SEMANTICS\]](#).

All DATA frames on the stream correspond to data sent or received on the TCP connection. Any DATA frame sent by the client is transmitted by the proxy to the TCP server; data received from the TCP server is packaged into DATA frames by the proxy. Note that the size and number of TCP segments is not guaranteed to map predictably to the size and number of HTTP DATA or QUIC STREAM frames.

Once the CONNECT method has completed, only DATA frames are permitted to be sent on the stream. Extension frames MAY be used if specifically permitted by the definition of the extension. Receipt of any other frame type MUST be treated as a connection error of type H3_FRAME_UNEXPECTED.

The TCP connection can be closed by either peer. When the client ends the request stream (that is, the receive stream at the proxy enters the "Data Recvd" state), the proxy will set the FIN bit on its connection to the TCP server. When the proxy receives a packet with the FIN bit set, it will terminate the send stream that it sends to the client. TCP connections which remain half-closed in a single direction are not invalid, but are often handled poorly by servers, so clients SHOULD NOT close a stream for sending while they still expect to receive data from the target of the CONNECT.

A TCP connection error is signaled by abruptly terminating the 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 of type H3_CONNECT_ERROR ([Section 8.1](#)). Correspondingly, if a proxy detects an error with the stream or the QUIC connection, it MUST close the TCP connection. If the underlying TCP implementation permits it, the proxy SHOULD send a TCP segment with the RST bit set.

4.3. HTTP Upgrade

HTTP/3 does not support the HTTP Upgrade mechanism (Section 9.9 of [[HTTP11](#)]) or 101 (Switching Protocols) informational status code (Section 9.2.2 of [[SEMANTICS](#)]).

4.4. Server Push

Server push is an interaction mode which permits a server to push a request-response exchange to a client in anticipation of the client making the indicated request. This trades off network usage against a potential latency gain. HTTP/3 server push is similar to what is described in HTTP/2 [[HTTP2](#)], but uses different mechanisms.

Each server push is identified by a unique Push ID. This Push ID is used in one or more PUSH_PROMISE frames (see [Section 7.2.5](#)) that carry the request fields, then included with the push stream which ultimately fulfills those promises. When the same Push ID is promised on multiple request streams, the decompressed request field sections MUST contain the same fields in the same order, and both the name and the value in each field MUST be exact matches.

Server push is only enabled on a connection when a client sends a MAX_PUSH_ID frame; see [Section 7.2.7](#). A server cannot use server push until it receives a MAX_PUSH_ID frame. A client sends additional MAX_PUSH_ID frames to control the number of pushes that a server can promise. A server SHOULD use Push IDs sequentially, starting at 0. A client MUST treat receipt of a push stream with a Push ID that is greater than the maximum Push ID as a connection error of type H3_ID_ERROR.

The header section of the request message is carried by a PUSH_PROMISE frame (see [Section 7.2.5](#)) on the request stream which generated the push. This allows the server push to be associated with a client request.

Not all requests can be pushed. A server MAY push requests which have the following properties:

- *cacheable; see Section 7.2.3 of [[SEMANTICS](#)]

- *safe; see Section 7.2.1 of [[SEMANTICS](#)]

- *does not include a request body or trailer section

The server MUST include a value in the ":authority" pseudo-header field for which the server is authoritative; see [Section 3.4](#).

Clients SHOULD send a CANCEL_PUSH frame upon receipt of a PUSH_PROMISE frame carrying a request which is not cacheable, is not known to be safe, that indicates the presence of a request body, or for which it does not consider the server authoritative.

Each pushed response is associated with one or more client requests. The push is associated with the request stream on which the PUSH_PROMISE frame was received. The same server push can be associated with additional client requests using a PUSH_PROMISE frame with the same Push ID on multiple request streams. These associations do not affect the operation of the protocol, but MAY be considered by user agents when deciding how to use pushed resources.

Ordering of a PUSH_PROMISE in relation to certain parts of the response is important. The server SHOULD send PUSH_PROMISE frames prior to sending HEADERS or DATA frames that reference the promised responses. This reduces the chance that a client requests a resource that will be pushed by the server.

When a server later fulfills a promise, the server push response is conveyed on a push stream; see [Section 6.2.2](#). The push stream identifies the Push ID of the promise that it fulfills, then contains a response to the promised request using the same format described for responses in [Section 4.1](#).

Due to reordering, push stream data can arrive before the corresponding PUSH_PROMISE frame. When a client receives a new push stream with an as-yet-unknown Push ID, both the associated client request and the pushed request header fields are unknown. The client can buffer the stream data in expectation of the matching PUSH_PROMISE. The client can use stream flow control (see section 4.1 of [[QUIC-TRANSPORT](#)]) to limit the amount of data a server may commit to the pushed stream.

If a promised server push is not needed by the client, the client SHOULD send a CANCEL_PUSH frame. If the push stream is already open or opens after sending the CANCEL_PUSH frame, the client can abort reading the stream with an error code of H3_REQUEST_CANCELLED. This asks the server not to transfer additional data and indicates that it will be discarded upon receipt.

Pushed responses that are cacheable (see Section 3 of [[CACHING](#)]) 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 is present (Section 5.2.2.3 of [[CACHING](#)])) at the time the pushed response is received.

Pushed responses that are not cacheable MUST NOT be stored by any HTTP cache. They MAY be made available to the application separately.

5. Connection Closure

Once established, an HTTP/3 connection can be used for many requests and responses over time until the connection is closed. Connection closure can happen in any of several different ways.

5.1. Idle Connections

Each QUIC endpoint declares an idle timeout during the handshake. If the connection remains idle (no packets received) for longer than this duration, the peer will assume that the connection has been closed. HTTP/3 implementations will need to open a new connection for new requests if the existing connection has been idle for longer than the server's advertised idle timeout, and SHOULD do so if approaching the idle timeout.

HTTP clients are expected to request that the transport keep connections open while there are responses outstanding for requests or server pushes, as described in Section 19.2 of [[QUIC-TRANSPORT](#)]. If the client is not expecting a response from the server, allowing an idle connection to time out is preferred over expending effort maintaining a connection that might not be needed. A gateway MAY maintain connections in anticipation of need rather than incur the latency cost of connection establishment to servers. Servers SHOULD NOT actively keep connections open.

5.2. Connection Shutdown

Even when a connection is not idle, either endpoint can decide to stop using the connection and initiate a graceful connection close. Endpoints initiate the graceful shutdown of a connection by sending a GOAWAY frame ([Section 7.2.6](#)). The GOAWAY frame contains an

identifier that indicates to the receiver the range of requests or pushes that were or might be processed in this connection. The server sends a client-initiated bidirectional Stream ID; the client sends a Push ID. Requests or pushes with the indicated identifier or greater are rejected by the sender of the GOAWAY. This identifier MAY be zero if no requests or pushes were processed.

The information in the GOAWAY frame enables a client and server to agree on which requests or pushes were accepted prior to the connection shutdown. Upon sending a GOAWAY frame, the endpoint SHOULD explicitly cancel (see [Section 4.1.2](#) and [Section 7.2.3](#)) any requests or pushes that have identifiers greater than or equal to that indicated, in order to clean up transport state for the affected streams. The endpoint SHOULD continue to do so as more requests or pushes arrive.

Endpoints MUST NOT initiate new requests or promise new pushes on the connection after receipt of a GOAWAY frame from the peer. Clients MAY establish a new connection to send additional requests.

Some requests or pushes might already be in transit:

- *Upon receipt of a GOAWAY frame, if the client has already sent requests with a Stream ID greater than or equal to the identifier received in a GOAWAY frame, those requests will not be processed. Clients can safely retry unprocessed requests on a different connection.

Requests on Stream IDs less than the Stream ID in a GOAWAY frame from the server might have been processed; their status cannot be known until a response is received, the stream is reset individually, another GOAWAY is received, or the connection terminates.

Servers MAY reject individual requests on streams below the indicated ID if these requests were not processed.

- *If a server receives a GOAWAY frame after having promised pushes with a Push ID greater than or equal to the identifier received in a GOAWAY frame, those pushes will not be accepted.

Servers SHOULD send a GOAWAY frame when the closing of a connection is known in advance, even if the advance notice is small, so that the remote peer can know whether a request has been partially processed or not. For example, if an HTTP client sends a POST at the same time that a server closes a QUIC 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.

A client that is unable to retry requests loses all requests that are in flight when the server closes the connection. An endpoint MAY send multiple GOAWAY frames indicating different identifiers, but MUST NOT increase the identifier value they send, since clients might already have retried unprocessed requests on another connection.

An endpoint that is attempting to gracefully shut down a connection can send a GOAWAY frame with a value set to the maximum possible value ($2^{62}-4$ for servers, $2^{62}-1$ for clients). This ensures that the peer stops creating new requests or pushes. After allowing time for any in-flight requests or pushes to arrive, the endpoint can send another GOAWAY frame indicating which requests or pushes it might accept before the end of the connection. This ensures that a connection can be cleanly shut down without losing requests.

A client has more flexibility in the value it chooses for the Push ID in a GOAWAY that it sends. A value of $2^{62} - 1$ indicates that the server can continue fulfilling pushes which have already been promised, and the client can continue granting push credit as needed; see [Section 7.2.7](#). A smaller value indicates the client will reject pushes with Push IDs greater than or equal to this value. Like the server, the client MAY send subsequent GOAWAY frames so long as the specified Push ID is strictly smaller than all previously sent values.

Even when a GOAWAY indicates that a given request or push will not be processed or accepted upon receipt, the underlying transport resources still exist. The endpoint that initiated these requests can cancel them to clean up transport state.

Once all accepted requests and pushes have been processed, the endpoint can permit the connection to become idle, or MAY initiate an immediate closure of the connection. An endpoint that completes a graceful shutdown SHOULD use the H3_NO_ERROR code when closing the connection.

If a client has consumed all available bidirectional stream IDs with requests, the server need not send a GOAWAY frame, since the client is unable to make further requests.

5.3. Immediate Application Closure

An HTTP/3 implementation can immediately close the QUIC connection at any time. This results in sending a QUIC CONNECTION_CLOSE frame to the peer indicating that the application layer has terminated the connection. The application error code in this frame indicates to the peer why the connection is being closed. See [Section 8](#) for error codes which can be used when closing a connection in HTTP/3.

Before closing the connection, a GOAWAY frame MAY be sent to allow the client to retry some requests. Including the GOAWAY frame in the same packet as the QUIC CONNECTION_CLOSE frame improves the chances of the frame being received by clients.

5.4. Transport Closure

For various reasons, the QUIC transport could indicate to the application layer that the connection has terminated. This might be due to an explicit closure by the peer, a transport-level error, or a change in network topology which interrupts connectivity.

If a connection terminates without a GOAWAY frame, clients MUST assume that any request which was sent, whether in whole or in part, might have been processed.

6. Stream Mapping and Usage

A QUIC stream provides reliable in-order delivery of bytes, but makes no guarantees about order of delivery with regard to bytes on other streams. On the wire, data is framed into QUIC STREAM frames, but this framing is invisible to the HTTP framing layer. The transport layer buffers and orders received QUIC STREAM frames, exposing the data contained within as a reliable byte stream to the application. Although QUIC permits out-of-order delivery within a stream, HTTP/3 does not make use of this feature.

QUIC streams can be either unidirectional, carrying data only from initiator to receiver, or bidirectional. Streams can be initiated by either the client or the server. For more detail on QUIC streams, see Section 2 of [[QUIC-TRANSPORT](#)].

When HTTP fields and data are sent over QUIC, the QUIC layer handles most of the stream management. HTTP does not need to do any separate multiplexing when using QUIC - data sent over a QUIC stream always maps to a particular HTTP transaction or connection context.

6.1. Bidirectional Streams

All client-initiated bidirectional streams are used for HTTP requests and responses. A bidirectional stream ensures that the response can be readily correlated with the request. This means that the client's first request occurs on QUIC stream 0, with subsequent requests on stream 4, 8, and so on. In order to permit these streams to open, an HTTP/3 server SHOULD configure non-zero minimum values for the number of permitted streams and the initial stream flow control window. So as to not unnecessarily limit parallelism, at least 100 requests SHOULD be permitted at a time.

HTTP/3 does not use server-initiated bidirectional streams, though an extension could define a use for these streams. Clients **MUST** treat receipt of a server-initiated bidirectional stream as a connection error of type `H3_STREAM_CREATION_ERROR` unless such an extension has been negotiated.

6.2. Unidirectional Streams

Unidirectional streams, in either direction, are used for a range of purposes. The purpose is indicated by a stream type, which is sent as a variable-length integer at the start of the stream. The format and structure of data that follows this integer is determined by the stream type.

```
Unidirectional Stream Header {  
    Stream Type (i),  
}
```

Figure 1: Unidirectional Stream Header

Some stream types are reserved ([Section 6.2.3](#)). Two stream types are defined in this document: control streams ([Section 6.2.1](#)) and push streams ([Section 6.2.2](#)). [QPACK] defines two additional stream types. Other stream types can be defined by extensions to HTTP/3; see [Section 9](#) for more details.

The performance of HTTP/3 connections in the early phase of their lifetime is sensitive to the creation and exchange of data on unidirectional streams. Endpoints that excessively restrict the number of streams or the flow control window of these streams will increase the chance that the remote peer reaches the limit early and becomes blocked. In particular, implementations should consider that remote peers may wish to exercise reserved stream behavior ([Section 6.2.3](#)) with some of the unidirectional streams they are permitted to use. To avoid blocking, the transport parameters sent by both clients and servers **MUST** allow the peer to create at least one unidirectional stream for the HTTP control stream plus the number of unidirectional streams required by mandatory extensions (three being the minimum number required for the base HTTP/3 protocol and QPACK), and **SHOULD** provide at least 1,024 bytes of flow control credit to each stream.

Note that an endpoint is not required to grant additional credits to create more unidirectional streams if its peer consumes all the initial credits before creating the critical unidirectional streams. Endpoints **SHOULD** create the HTTP control stream as well as the unidirectional streams required by mandatory extensions (such as the QPACK encoder and decoder streams) first, and then create additional streams as allowed by their peer.

If the stream header indicates a stream type which is not supported by the recipient, the remainder of the stream cannot be consumed as the semantics are unknown. Recipients of unknown stream types MAY abort reading of the stream with an error code of `H3_STREAM_CREATION_ERROR`, but MUST NOT consider such streams to be a connection error of any kind.

Implementations MAY send stream types before knowing whether the peer supports them. However, stream types which could modify the state or semantics of existing protocol components, including QPACK or other extensions, MUST NOT be sent until the peer is known to support them.

A sender can close or reset a unidirectional stream unless otherwise specified. A receiver MUST tolerate unidirectional streams being closed or reset prior to the reception of the unidirectional stream header.

6.2.1. Control Streams

A control stream is indicated by a stream type of `0x00`. Data on this stream consists of HTTP/3 frames, as defined in [Section 7.2](#).

Each side MUST initiate a single control stream at the beginning of the connection and send its SETTINGS frame as the first frame on this stream. If the first frame of the control stream is any other frame type, this MUST be treated as a connection error of type `H3_MISSING_SETTINGS`. Only one control stream per peer is permitted; receipt of a second stream which claims to be a control stream MUST be treated as a connection error of type `H3_STREAM_CREATION_ERROR`. The sender MUST NOT close the control stream, and the receiver MUST NOT request that the sender close the control stream. If either control stream is closed at any point, this MUST be treated as a connection error of type `H3_CLOSED_CRITICAL_STREAM`.

A pair of unidirectional streams is used rather than a single bidirectional stream. This allows either peer to send data as soon as it is able. Depending on whether 0-RTT is enabled on the connection, either client or server might be able to send stream data first after the cryptographic handshake completes.

6.2.2. Push Streams

Server push is an optional feature introduced in HTTP/2 that allows a server to initiate a response before a request has been made. See [Section 4.4](#) for more details.

A push stream is indicated by a stream type of `0x01`, followed by the Push ID of the promise that it fulfills, encoded as a variable-length integer. The remaining data on this stream consists of HTTP/3

frames, as defined in [Section 7.2](#), and fulfills a promised server push by zero or more interim HTTP responses followed by a single final HTTP response, as defined in [Section 4.1](#). Server push and Push IDs are described in [Section 4.4](#).

Only servers can push; if a server receives a client-initiated push stream, this MUST be treated as a connection error of type H3_STREAM_CREATION_ERROR.

```
Push Stream Header {  
  Stream Type (i) = 0x01,  
  Push ID (i),  
}
```

Figure 2: Push Stream Header

Each Push ID MUST only be used once in a push stream header. If a push stream header includes a Push ID that was used in another push stream header, the client MUST treat this as a connection error of type H3_ID_ERROR.

6.2.3. Reserved Stream Types

Stream types of the format $0x1f * N + 0x21$ for non-negative integer values of N are reserved to exercise the requirement that unknown types be ignored. These streams have no semantics, and can be sent when application-layer padding is desired. They MAY also be sent on connections where no data is currently being transferred. Endpoints MUST NOT consider these streams to have any meaning upon receipt.

The payload and length of the stream are selected in any manner the implementation chooses. Implementations MAY terminate these streams cleanly, or MAY abruptly terminate them. When terminating abruptly, the error code H3_NO_ERROR or a reserved error code ([Section 8.1](#)) SHOULD be used.

7. HTTP Framing Layer

HTTP frames are carried on QUIC streams, as described in [Section 6](#). HTTP/3 defines three stream types: control stream, request stream, and push stream. This section describes HTTP/3 frame formats and the streams types on which they are permitted; see [Table 1](#) for an overview. A comparison between HTTP/2 and HTTP/3 frames is provided in [Appendix A.2](#).

Frame	Control Stream	Request Stream	Push Stream	Section
DATA	No	Yes	Yes	

Frame	Control Stream	Request Stream	Push Stream	Section
				Section 7.2.1
HEADERS	No	Yes	Yes	Section 7.2.2
CANCEL_PUSH	Yes	No	No	Section 7.2.3
SETTINGS	Yes (1)	No	No	Section 7.2.4
PUSH_PROMISE	No	Yes	No	Section 7.2.5
GOAWAY	Yes	No	No	Section 7.2.6
MAX_PUSH_ID	Yes	No	No	Section 7.2.7
Reserved	Yes	Yes	Yes	Section 7.2.8

Table 1: HTTP/3 Frames and Stream Type Overview

Certain frames can only occur as the first frame of a particular stream type; these are indicated in [Table 1](#) with a (1). Specific guidance is provided in the relevant section.

Note that, unlike QUIC frames, HTTP/3 frames can span multiple packets.

7.1. Frame Layout

All frames have the following format:

```
HTTP/3 Frame Format {
  Type (i),
  Length (i),
  Frame Payload (..),
}
```

Figure 3: HTTP/3 Frame Format

A frame includes the following fields:

Type: A variable-length integer that identifies the frame type.

Length: A variable-length integer that describes the length in bytes of the Frame Payload.

Frame Payload: A payload, the semantics of which are determined by the Type field.

Each frame's payload MUST contain exactly the fields identified in its description. A frame payload that contains additional bytes after the identified fields or a frame payload that terminates before the end of the identified fields MUST be treated as a connection error of type H3_FRAME_ERROR.

When a stream terminates cleanly, if the last frame on the stream was truncated, this MUST be treated as a connection error ([Section 8](#)) of type H3_FRAME_ERROR. Streams which terminate abruptly may be reset at any point in a frame.

7.2. Frame Definitions

7.2.1. DATA

DATA frames (type=0x0) convey arbitrary, variable-length sequences of bytes associated with an HTTP request or response payload.

DATA frames MUST be associated with an HTTP request or response. If a DATA frame is received on a control stream, the recipient MUST respond with a connection error ([Section 8](#)) of type H3_FRAME_UNEXPECTED.

```
DATA Frame {  
    Type (i) = 0x0,  
    Length (i),  
    Data (..),  
}
```

Figure 4: DATA Frame

7.2.2. HEADERS

The HEADERS frame (type=0x1) is used to carry an HTTP field section, encoded using QPACK. See [[QPACK](#)] for more details.

```
HEADERS Frame {  
    Type (i) = 0x1,  
    Length (i),  
    Encoded Field Section (..),  
}
```

Figure 5: HEADERS Frame

HEADERS frames can only be sent on request / push streams. If a HEADERS frame is received on a control stream, the recipient MUST respond with a connection error ([Section 8](#)) of type H3_FRAME_UNEXPECTED.

7.2.3. CANCEL_PUSH

The CANCEL_PUSH frame (type=0x3) is used to request cancellation of a server push prior to the push stream being received. The CANCEL_PUSH frame identifies a server push by Push ID (see [Section 7.2.5](#)), encoded as a variable-length integer.

When a client sends CANCEL_PUSH, it is indicating that it does not wish to receive the promised resource. The server SHOULD abort sending the resource, but the mechanism to do so depends on the state of the corresponding push stream. If the server has not yet created a push stream, it does not create one. If the push stream is open, the server SHOULD abruptly terminate that stream. If the push stream has already ended, the server MAY still abruptly terminate the stream or MAY take no action.

When a server sends CANCEL_PUSH, it is indicating that it will not be fulfilling a promise and has not created a push stream. The client should not expect the corresponding promise to be fulfilled.

Sending CANCEL_PUSH has no direct effect on the state of existing push streams. A server SHOULD NOT send a CANCEL_PUSH when it has already created a corresponding push stream, and a client SHOULD NOT send a CANCEL_PUSH when it has already received a corresponding push stream. If a push stream arrives after a client has sent CANCEL_PUSH, this MAY be treated as a stream error of type H3_STREAM_CREATION_ERROR.

A CANCEL_PUSH frame is sent on the control stream. Receiving a CANCEL_PUSH frame on a stream other than the control stream MUST be treated as a connection error of type H3_FRAME_UNEXPECTED.

```
CANCEL_PUSH Frame {  
  Type (i) = 0x3,  
  Length (i),  
  Push ID (..),  
}
```

Figure 6: CANCEL_PUSH Frame

The CANCEL_PUSH frame carries a Push ID encoded as a variable-length integer. The Push ID identifies the server push that is being cancelled; see [Section 7.2.5](#). If a CANCEL_PUSH frame is received which references a Push ID greater than currently allowed on the connection, this MUST be treated as a connection error of type H3_ID_ERROR.

If the client receives a CANCEL_PUSH frame, that frame might identify a Push ID that has not yet been mentioned by a PUSH_PROMISE

frame due to reordering. If a server receives a CANCEL_PUSH frame for a Push ID that has not yet been mentioned by a PUSH_PROMISE frame, this MUST be treated as a connection error of type H3_ID_ERROR.

7.2.4. SETTINGS

The SETTINGS frame (type=0x4) conveys configuration parameters that affect how endpoints communicate, such as preferences and constraints on peer behavior. Individually, a SETTINGS parameter can also be referred to as a "setting"; the identifier and value of each setting parameter can be referred to as a "setting identifier" and a "setting value".

SETTINGS frames always apply to a connection, never a single stream. A SETTINGS frame MUST be sent as the first frame of each control stream (see [Section 6.2.1](#)) by each peer, and MUST NOT be sent subsequently. If an endpoint receives a second SETTINGS frame on the control stream, the endpoint MUST respond with a connection error of type H3_FRAME_UNEXPECTED.

SETTINGS frames MUST NOT be sent on any stream other than the control stream. If an endpoint receives a SETTINGS frame on a different stream, the endpoint MUST respond with a connection error of type H3_FRAME_UNEXPECTED.

SETTINGS parameters are not negotiated; they describe characteristics of the sending peer, which can be used by the receiving peer. However, a negotiation can be implied by the use of SETTINGS - each peer uses SETTINGS to advertise a set of supported values. The definition of the setting would describe how each peer combines the two sets to conclude which choice will be used. SETTINGS does not provide a mechanism to identify when the choice takes effect.

Different values for the same parameter can be advertised by each peer. For example, a client might be willing to consume a very large response field section, while servers are more cautious about request size.

The same setting identifier MUST NOT occur more than once in the SETTINGS frame. A receiver MAY treat the presence of duplicate setting identifiers as a connection error of type H3_SETTINGS_ERROR.

The payload of a SETTINGS frame consists of zero or more parameters. Each parameter consists of a setting identifier and a value, both encoded as QUIC variable-length integers.

```

Setting {
    Identifier (i),
    Value (i),
}

SETTINGS Frame {
    Type (i) = 0x4,
    Length (i),
    Setting (...) ...,
}

```

Figure 7: SETTINGS Frame

An implementation MUST ignore the contents for any SETTINGS identifier it does not understand.

7.2.4.1. Defined SETTINGS Parameters

The following settings are defined in HTTP/3:

SETTINGS_MAX_FIELD_SECTION_SIZE (0x6): The default value is unlimited. See [Section 4.1.1](#) for usage.

Setting identifiers of the format $0x1f * N + 0x21$ for non-negative integer values of N are reserved to exercise the requirement that unknown identifiers be ignored. Such settings have no defined meaning. Endpoints SHOULD include at least one such setting in their SETTINGS frame. Endpoints MUST NOT consider such settings to have any meaning upon receipt.

Because the setting has no defined meaning, the value of the setting can be any value the implementation selects.

Additional settings can be defined by extensions to HTTP/3; see [Section 9](#) for more details.

7.2.4.2. Initialization

An HTTP implementation MUST NOT send frames or requests which would be invalid based on its current understanding of the peer's settings.

All settings begin at an initial value. Each endpoint SHOULD use these initial values to send messages before the peer's SETTINGS frame has arrived, as packets carrying the settings can be lost or delayed. When the SETTINGS frame arrives, any settings are changed to their new values.

This removes the need to wait for the SETTINGS frame before sending messages. Endpoints MUST NOT require any data to be received from

the peer prior to sending the SETTINGS frame; settings MUST be sent as soon as the transport is ready to send data.

For servers, the initial value of each client setting is the default value.

For clients using a 1-RTT QUIC connection, the initial value of each server setting is the default value. 1-RTT keys will always become available prior to SETTINGS arriving, even if the server sends SETTINGS immediately. Clients SHOULD NOT wait indefinitely for SETTINGS to arrive before sending requests, but SHOULD process received datagrams in order to increase the likelihood of processing SETTINGS before sending the first request.

When a 0-RTT QUIC connection is being used, the initial value of each server setting is the value used in the previous session. Clients SHOULD store the settings the server provided in the connection where resumption information was provided, but MAY opt not to store settings in certain cases (e.g., if the session ticket is received before the SETTINGS frame). A client MUST comply with stored settings - or default values, if no values are stored - when attempting 0-RTT. Once a server has provided new settings, clients MUST comply with those values.

A server can remember the settings that it advertised, or store an integrity-protected copy of the values in the ticket and recover the information when accepting 0-RTT data. A server uses the HTTP/3 settings values in determining whether to accept 0-RTT data. If the server cannot determine that the settings remembered by a client are compatible with its current settings, it MUST NOT accept 0-RTT data. Remembered settings are compatible if a client complying with those settings would not violate the server's current settings.

A server MAY accept 0-RTT and subsequently provide different settings in its SETTINGS frame. If 0-RTT data is accepted by the server, its SETTINGS frame MUST NOT reduce any limits or alter any values that might be violated by the client with its 0-RTT data. The server MUST include all settings which differ from their default values. If a server accepts 0-RTT but then sends settings that are not compatible with the previously specified settings, this MUST be treated as a connection error of type H3_SETTINGS_ERROR. If a server accepts 0-RTT but then sends a SETTINGS frame that omits a setting value that the client understands (apart from reserved setting identifiers) that was previously specified to have a non-default value, this MUST be treated as a connection error of type H3_SETTINGS_ERROR.

7.2.5. PUSH_PROMISE

The PUSH_PROMISE frame (type=0x5) is used to carry a promised request header field section from server to client on a request stream, as in HTTP/2.

```
PUSH_PROMISE Frame {  
    Type (i) = 0x5,  
    Length (i),  
    Push ID (i),  
    Encoded Field Section (..),  
}
```

Figure 8: PUSH_PROMISE Frame

The payload consists of:

Push ID: A variable-length integer that identifies the server push operation. A Push ID is used in push stream headers ([Section 4.4](#)), CANCEL_PUSH frames ([Section 7.2.3](#)).

Encoded Field Section: QPACK-encoded request header fields for the promised response. See [[QPACK](#)] for more details.

A server MUST NOT use a Push ID that is larger than the client has provided in a MAX_PUSH_ID frame ([Section 7.2.7](#)). A client MUST treat receipt of a PUSH_PROMISE frame that contains a larger Push ID than the client has advertised as a connection error of H3_ID_ERROR.

A server MAY use the same Push ID in multiple PUSH_PROMISE frames. If so, the decompressed request header sets MUST contain the same fields in the same order, and both the name and the value in each field MUST be exact matches. Clients SHOULD compare the request header sections for resources promised multiple times. If a client receives a Push ID that has already been promised and detects a mismatch, it MUST respond with a connection error of type H3_GENERAL_PROTOCOL_ERROR. If the decompressed field sections match exactly, the client SHOULD associate the pushed content with each stream on which a PUSH_PROMISE was received.

Allowing duplicate references to the same Push ID is primarily to reduce duplication caused by concurrent requests. A server SHOULD avoid reusing a Push ID over a long period. Clients are likely to consume server push responses and not retain them for reuse over time. Clients that see a PUSH_PROMISE that uses a Push ID that they have already consumed and discarded are forced to ignore the PUSH_PROMISE.

If a PUSH_PROMISE frame is received on the control stream, the client MUST respond with a connection error ([Section 8](#)) of type H3_FRAME_UNEXPECTED.

A client MUST NOT send a PUSH_PROMISE frame. A server MUST treat the receipt of a PUSH_PROMISE frame as a connection error of type H3_FRAME_UNEXPECTED.

See [Section 4.4](#) for a description of the overall server push mechanism.

7.2.6. GOAWAY

The GOAWAY frame (type=0x7) is used to initiate graceful shutdown of a connection by either endpoint. GOAWAY allows an endpoint to stop accepting new requests or pushes while still finishing processing of previously received requests and pushes. This enables administrative actions, like server maintenance. GOAWAY by itself does not close a connection.

```
GOAWAY Frame {  
  Type (i) = 0x7,  
  Length (i),  
  Stream ID/Push ID (..),  
}
```

Figure 9: GOAWAY Frame

The GOAWAY frame is always sent on the control stream. In the server to client direction, it carries a QUIC Stream ID for a client-initiated bidirectional stream encoded as a variable-length integer. A client MUST treat receipt of a GOAWAY frame containing a Stream ID of any other type as a connection error of type H3_ID_ERROR.

In the client to server direction, the GOAWAY frame carries a Push ID encoded as a variable-length integer.

The GOAWAY frame applies to the connection, not a specific stream. A client MUST treat a GOAWAY frame on a stream other than the control stream as a connection error ([Section 8](#)) of type H3_FRAME_UNEXPECTED.

See [Section 5.2](#) for more information on the use of the GOAWAY frame.

7.2.7. MAX_PUSH_ID

The MAX_PUSH_ID frame (type=0xD) is used by clients to control the number of server pushes that the server can initiate. This sets the maximum value for a Push ID that the server can use in PUSH_PROMISE

and CANCEL_PUSH frames. Consequently, this also limits the number of push streams that the server can initiate in addition to the limit maintained by the QUIC transport.

The MAX_PUSH_ID frame is always sent on the control stream. Receipt of a MAX_PUSH_ID frame on any other stream MUST be treated as a connection error of type H3_FRAME_UNEXPECTED.

A server MUST NOT send a MAX_PUSH_ID frame. A client MUST treat the receipt of a MAX_PUSH_ID frame as a connection error of type H3_FRAME_UNEXPECTED.

The maximum Push ID is unset when a connection is created, meaning that a server cannot push until it receives a MAX_PUSH_ID frame. A client that wishes to manage the number of promised server pushes can increase the maximum Push ID by sending MAX_PUSH_ID frames as the server fulfills or cancels server pushes.

```
MAX_PUSH_ID Frame {  
  Type (i) = 0x1,  
  Length (i),  
  Push ID (..),  
}
```

Figure 10: MAX_PUSH_ID Frame Payload

The MAX_PUSH_ID frame carries a single variable-length integer that identifies the maximum value for a Push ID that the server can use; see [Section 7.2.5](#). A MAX_PUSH_ID frame cannot reduce the maximum Push ID; receipt of a MAX_PUSH_ID that contains a smaller value than previously received MUST be treated as a connection error of type H3_ID_ERROR.

7.2.8. Reserved Frame Types

Frame types of the format $0x1f * N + 0x21$ for non-negative integer values of N are reserved to exercise the requirement that unknown types be ignored ([Section 9](#)). These frames have no semantics, and can be sent on any open stream when application-layer padding is desired. They MAY also be sent on connections where no data is currently being transferred. Endpoints MUST NOT consider these frames to have any meaning upon receipt.

The payload and length of the frames are selected in any manner the implementation chooses.

Frame types which were used in HTTP/2 where there is no corresponding HTTP/3 frame have also been reserved ([Section 11.2.1](#)).

These frame types MUST NOT be sent, and receipt MAY be treated as an error of type H3_FRAME_UNEXPECTED.

8. Error Handling

QUIC allows the application to abruptly terminate (reset) individual streams or the entire connection when an error is encountered. These are referred to as "stream errors" or "connection errors" and are described in more detail in [\[QUIC-TRANSPORT\]](#).

An endpoint MAY choose to treat a stream error as a connection error under certain circumstances. Implementations need to consider the impact on outstanding requests before making this choice.

Because new error codes can be defined without negotiation (see [Section 9](#)), use of an error code in an unexpected context or receipt of an unknown error code MUST be treated as equivalent to H3_NO_ERROR. However, closing a stream can have other effects regardless of the error code; see [Section 4.1](#).

This section describes HTTP/3-specific error codes which can be used to express the cause of a connection or stream error.

8.1. HTTP/3 Error Codes

The following error codes are defined for use when abruptly terminating streams, aborting reading of streams, or immediately closing connections.

H3_NO_ERROR (0x100): No error. This is used when the connection or stream needs to be closed, but there is no error to signal.

H3_GENERAL_PROTOCOL_ERROR (0x101): Peer violated protocol requirements in a way which doesn't match a more specific error code, or endpoint declines to use the more specific error code.

H3_INTERNAL_ERROR (0x102): An internal error has occurred in the HTTP stack.

H3_STREAM_CREATION_ERROR (0x103): The endpoint detected that its peer created a stream that it will not accept.

H3_CLOSED_CRITICAL_STREAM (0x104): A stream required by the connection was closed or reset.

H3_FRAME_UNEXPECTED (0x105): A frame was received which was not permitted in the current state or on the current stream.

H3_FRAME_ERROR (0x106): A frame that fails to satisfy layout requirements or with an invalid size was received.

H3_EXCESSIVE_LOAD (0x107):

The endpoint detected that its peer is exhibiting a behavior that might be generating excessive load.

H3_ID_ERROR (0x108): A Stream ID or Push ID was used incorrectly, such as exceeding a limit, reducing a limit, or being reused.

H3_SETTINGS_ERROR (0x109): An endpoint detected an error in the payload of a SETTINGS frame.

H3_MISSING_SETTINGS (0x10A): No SETTINGS frame was received at the beginning of the control stream.

H3_REQUEST_REJECTED (0x10B): A server rejected a request without performing any application processing.

H3_REQUEST_CANCELLED (0x10C): The request or its response (including pushed response) is cancelled.

H3_REQUEST_INCOMPLETE (0x10D): The client's stream terminated without containing a fully-formed request.

H3_CONNECT_ERROR (0x10F): The connection established in response to a CONNECT request was reset or abnormally closed.

H3_VERSION_FALLBACK (0x110): The requested operation cannot be served over HTTP/3. The peer should retry over HTTP/1.1.

Error codes of the format $0x1f * N + 0x21$ for non-negative integer values of N are reserved to exercise the requirement that unknown error codes be treated as equivalent to H3_NO_ERROR ([Section 9](#)). Implementations SHOULD select an error code from this space with some probability when they would have sent H3_NO_ERROR.

9. Extensions to HTTP/3

HTTP/3 permits extension of the protocol. Within the limitations described in this section, protocol extensions can be used to provide additional services or alter any aspect of the protocol. Extensions are effective only within the scope of a single HTTP/3 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 fields.

Extensions are permitted to use new frame types ([Section 7.2](#)), new settings ([Section 7.2.4.1](#)), new error codes ([Section 8](#)), or new unidirectional stream types ([Section 6.2](#)). Registries are established for managing these extension points: frame types

([Section 11.2.1](#)), settings ([Section 11.2.2](#)), error codes ([Section 11.2.3](#)), and stream types ([Section 11.2.4](#)).

Implementations MUST ignore unknown or unsupported values in all extensible protocol elements. Implementations MUST discard frames and unidirectional streams 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, where a known frame type is required to be in a specific location, such as the SETTINGS frame as the first frame of the control stream (see [Section 6.2.1](#)), an unknown frame type does not satisfy that requirement and SHOULD be treated as an 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. Coordinating when such a revised layout comes into effect could prove complex. As such, allocating new identifiers for new definitions of existing protocol elements is likely to be more effective.

This document doesn't mandate a specific method for negotiating the use of an extension but notes that a setting ([Section 7.2.4.1](#)) 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 default value MUST be defined in such a fashion that the extension is disabled if the setting is omitted.

10. Security Considerations

The security considerations of HTTP/3 should be comparable to those of HTTP/2 with TLS. However, many of the considerations from Section 10 of [[HTTP2](#)] apply to [[QUIC-TRANSPORT](#)] and are discussed in that document.

10.1. Server Authority

HTTP/3 relies on the HTTP definition of authority. The security considerations of establishing authority are discussed in Section 11.1 of [[SEMANTICS](#)].

10.2. Cross-Protocol Attacks

The use of ALPN in the TLS and QUIC handshakes establishes the target application protocol before application-layer bytes are processed. Because all QUIC packets are encrypted, it is difficult for an attacker to control the plaintext bytes of an HTTP/3

connection which could be used in a cross-protocol attack on a plaintext protocol.

10.3. Intermediary Encapsulation Attacks

The HTTP/3 field encoding allows the expression of names that are not valid field names in the syntax used by HTTP (Section 4.3 of [\[SEMANTICS\]](#)). Requests or responses containing invalid field names MUST be treated as malformed ([Section 4.1.3](#)). An intermediary therefore cannot translate an HTTP/3 request or response containing an invalid field name into an HTTP/1.1 message.

Similarly, HTTP/3 allows field values that are not valid. While most of the values that can be encoded will not alter 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 field value MUST be treated as malformed ([Section 4.1.3](#)). Valid characters are defined by the "field-content" ABNF rule in Section 4.4 of [\[SEMANTICS\]](#).

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 3.4](#)) MUST NOT be used or cached.

10.5. Denial-of-Service Considerations

An HTTP/3 connection can demand a greater commitment of resources to operate than an HTTP/1.1 or HTTP/2 connection. The use of field 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 constrained in a similar fashion. A client that accepts server push SHOULD limit the number of Push IDs it issues at a time.

Processing capacity cannot be guarded as effectively as state capacity.

The ability to send undefined protocol elements which the peer is required to ignore can be abused to cause a peer to expend additional processing time. This might be done by setting multiple undefined SETTINGS parameters, unknown frame types, or unknown stream types. Note, however, that some uses are entirely legitimate, such as optional-to-understand extensions and padding to increase resistance to traffic analysis.

Compression of field sections also offers some opportunities to waste processing resources; see Section 7 of [\[QPACK\]](#) for more details on potential abuses.

All these features - i.e., server push, unknown protocol elements, field 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 8](#)) of type H3_EXCESSIVE_LOAD, but false positives will result in disrupting valid connections and requests.

10.5.1. Limits on Field Section Size

A large field section ([Section 4.1](#)) 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 field section, which prevents streaming of the header field section to its ultimate destination. This ordering and other reasons, such as ensuring cache correctness, mean that an endpoint likely needs to buffer the entire header field section. Since there is no hard limit to the size of a field section, 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 ([Section 7.2.4.1](#)) setting to advise peers of limits that might apply on the size of field sections. This setting is only advisory, so endpoints MAY choose to send field sections 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 field section 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.

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. Therefore, a proxy cannot rely on QUIC stream limits alone to control 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/3 enables compression of fields ([Section 4.1.1](#)); the following concerns also apply to the use of HTTP compressed content-codings; see Section 6.1.2 of [[SEMANTICS](#)].

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.

Further considerations regarding the compression of fields sections are described in [[QPACK](#)].

10.7. Padding and Traffic Analysis

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 (e.g., [[BREACH](#)]).

Where HTTP/2 employs PADDING frames and Padding fields in other frames to make a connection more resistant to traffic analysis, HTTP/3 can either rely on transport-layer padding or employ the reserved frame and stream types discussed in [Section 7.2.8](#) and [Section 6.2.3](#). These methods of padding produce different results in terms of the granularity of padding, how padding is arranged in relation to the information that is being protected, whether padding is applied in the case of packet loss, and how an implementation might control padding. Redundant padding could even be counterproductive.

To mitigate attacks that rely on compression, disabling or limiting compression might be preferable to padding as a countermeasure.

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-sized boundary, which could be possible if an attacker can control plaintext.

10.8. Frame Parsing

Several protocol elements contain nested length elements, typically in the form of frames with an explicit length containing variable-length integers. This could pose a security risk to an incautious implementer. An implementation **MUST** ensure that the length of a frame exactly matches the length of the fields it contains.

10.9. Early Data

The use of 0-RTT with HTTP/3 creates an exposure to replay attack. The anti-replay mitigations in [[HTTP-REPLAY](#)] **MUST** be applied when using HTTP/3 with 0-RTT.

10.10. Migration

Certain HTTP implementations use the client address for logging or access-control purposes. Since a QUIC client's address might change during a connection (and future versions might support simultaneous use of multiple addresses), such implementations will need to either actively retrieve the client's current address or addresses when they are relevant or explicitly accept that the original address might change.

10.11. Privacy Considerations

Several characteristics of HTTP/3 provide an observer an opportunity to correlate actions of a single client or server over time. These include the value of settings, the timing of reactions to stimulus, and the handling of any features that are controlled by settings.

As far as these create observable differences in behavior, they could be used as a basis for fingerprinting a specific client.

HTTP/3's preference for using a single QUIC connection allows correlation of a user's activity on a site. Reusing connections for different origins allows for correlation of activity across those origins.

Several features of QUIC solicit immediate responses and can be used by an endpoint to measure latency to their peer; this might have privacy implications in certain scenarios.

11. IANA Considerations

This document registers a new ALPN protocol ID ([Section 11.1](#)) and creates new registries that manage the assignment of codepoints in HTTP/3.

11.1. Registration of HTTP/3 Identification String

This document creates a new registration for the identification of HTTP/3 in the "Application Layer Protocol Negotiation (ALPN) Protocol IDs" registry established in [[RFC7301](#)].

The "h3" string identifies HTTP/3:

Protocol: HTTP/3

Identification Sequence: 0x68 0x33 ("h3")

Specification: This document

11.2. New Registries

New registries created in this document operate under the QUIC registration policy documented in Section 22.1 of [[QUIC-TRANSPORT](#)]. These registries all include the common set of fields listed in Section 22.1.1 of [[QUIC-TRANSPORT](#)].

The initial allocations in these registries created in this document are all assigned permanent status and list as contact both the IESG (iesg@ietf.org) and the HTTP working group (ietf-http-wg@w3.org).

11.2.1. Frame Types

This document establishes a registry for HTTP/3 frame type codes. The "HTTP/3 Frame Type" registry governs a 62-bit space. This registry follows the QUIC registry policy; see [Section 11.2](#). Permanent registrations in this registry are assigned using the Specification Required policy [[RFC8126](#)], except for values between 0x00 and 0x3f (in hexadecimal; inclusive), which are assigned using Standards Action or IESG Approval as defined in Section 4.9 and 4.10 of [[RFC8126](#)].

While this registry is separate from the "HTTP/2 Frame Type" registry defined in [[HTTP2](#)], it is preferable that the assignments parallel each other where the code spaces overlap. If an entry is present in only one registry, every effort SHOULD be made to avoid assigning the corresponding value to an unrelated operation.

In addition to common fields as described in [Section 11.2](#), permanent registrations in this registry MUST include the following field:

Frame Type: A name or label for the frame type.

Specifications of frame types MUST include a description of the frame layout and its semantics, including any parts of the frame that are conditionally present.

The entries in [Table 2](#) are registered by this document.

Frame Type	Value	Specification
DATA	0x0	Section 7.2.1
HEADERS	0x1	Section 7.2.2
Reserved	0x2	N/A
CANCEL_PUSH	0x3	Section 7.2.3
SETTINGS	0x4	Section 7.2.4
PUSH_PROMISE	0x5	Section 7.2.5
Reserved	0x6	N/A
GOAWAY	0x7	Section 7.2.6
Reserved	0x8	N/A
Reserved	0x9	N/A
MAX_PUSH_ID	0xD	Section 7.2.7

Table 2: Initial HTTP/3 Frame Types

Additionally, each code of the format $0x1f * N + 0x21$ for non-negative integer values of N (that is, 0x21, 0x40, ..., through 0xFFFFFFFFFFFFFFFE) MUST NOT be assigned by IANA.

11.2.2. Settings Parameters

This document establishes a registry for HTTP/3 settings. The "HTTP/3 Settings" registry governs a 62-bit space. This registry follows the QUIC registry policy; see [Section 11.2](#). Permanent registrations in this registry are assigned using the Specification Required policy [[RFC8126](#)], except for values between 0x00 and 0x3f (in hexadecimal; inclusive), which are assigned using Standards Action or IESG Approval as defined in Section 4.9 and 4.10 of [[RFC8126](#)].

While this registry is separate from the "HTTP/2 Settings" registry defined in [[HTTP2](#)], it is preferable that the assignments parallel each other. If an entry is present in only one registry, every effort SHOULD be made to avoid assigning the corresponding value to an unrelated operation.

In addition to common fields as described in [Section 11.2](#), permanent registrations in this registry MUST include the following fields:

Setting Name: A symbolic name for the setting. Specifying a setting name is optional.

Default: The value of the setting unless otherwise indicated. A default SHOULD be the most restrictive possible value.

The entries in [Table 3](#) are registered by this document.

Setting Name	Value	Specification	Default
Reserved	0x2	N/A	N/A
Reserved	0x3	N/A	N/A
Reserved	0x4	N/A	N/A
Reserved	0x5	N/A	N/A
MAX_FIELD_SECTION_SIZE	0x6	Section 7.2.4.1	Unlimited

Table 3: Initial HTTP/3 Settings

Additionally, each code of the format $0x1f * N + 0x21$ for non-negative integer values of N (that is, 0x21, 0x40, ..., through 0x3FFFFFFFFFFFFFFE) MUST NOT be assigned by IANA.

11.2.3. Error Codes

This document establishes a registry for HTTP/3 error codes. The "HTTP/3 Error Code" registry manages a 62-bit space. This registry follows the QUIC registry policy; see [Section 11.2](#). Permanent registrations in this registry are assigned using the Specification Required policy [[RFC8126](#)], except for values between 0x00 and 0x3f (in hexadecimal; inclusive), which are assigned using Standards

Action or IESG Approval as defined in Section 4.9 and 4.10 of [\[RFC8126\]](#).

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.

In addition to common fields as described in [Section 11.2](#), permanent registrations in this registry MUST include the following fields:

Name: A name for the error code. Specifying an error code name is optional.

Description: A brief description of the error code semantics.

The entries in the [Table 4](#) are registered by this document.

Name	Value	Description	Specification
H3_NO_ERROR	0x0100	No error	Section 8.1
H3_GENERAL_PROTOCOL_ERROR	0x0101	General protocol error	Section 8.1
H3_INTERNAL_ERROR	0x0102	Internal error	Section 8.1
H3_STREAM_CREATION_ERROR	0x0103	Stream creation error	Section 8.1
H3_CLOSED_CRITICAL_STREAM	0x0104	Critical stream was closed	Section 8.1
H3_FRAME_UNEXPECTED	0x0105	Frame not permitted in the current state	Section 8.1
H3_FRAME_ERROR	0x0106	Frame violated layout or size rules	Section 8.1
H3_EXCESSIVE_LOAD	0x0107	Peer generating excessive load	Section 8.1
H3_ID_ERROR	0x0108	An identifier was used incorrectly	Section 8.1
H3_SETTINGS_ERROR	0x0109	SETTINGS frame contained invalid values	Section 8.1
H3_MISSING_SETTINGS	0x010A	No SETTINGS frame received	Section 8.1
H3_REQUEST_REJECTED	0x010B	Request not processed	Section 8.1
H3_REQUEST_CANCELLED	0x010C	Data no longer needed	Section 8.1
H3_REQUEST_INCOMPLETE	0x010D	Stream terminated early	Section 8.1

Name	Value	Description	Specification
H3_CONNECT_ERROR	0x010F	TCP reset or error on CONNECT request	Section 8.1
H3_VERSION_FALLBACK	0x0110	Retry over HTTP/1.1	Section 8.1

Table 4: Initial HTTP/3 Error Codes

Additionally, each code of the format $0x1f * N + 0x21$ for non-negative integer values of N (that is, 0x21, 0x40, ..., through 0x3FFFFFFFFFFFFFFE) MUST NOT be assigned by IANA.

11.2.4. Stream Types

This document establishes a registry for HTTP/3 unidirectional stream types. The "HTTP/3 Stream Type" registry governs a 62-bit space. This registry follows the QUIC registry policy; see [Section 11.2](#). Permanent registrations in this registry are assigned using the Specification Required policy [[RFC8126](#)], except for values between 0x00 and 0x3f (in hexadecimal; inclusive), which are assigned using Standards Action or IESG Approval as defined in Section 4.9 and 4.10 of [[RFC8126](#)].

In addition to common fields as described in [Section 11.2](#), permanent registrations in this registry MUST include the following fields:

Stream Type: A name or label for the stream type.

Sender: Which endpoint on a connection may initiate a stream of this type. Values are "Client", "Server", or "Both".

Specifications for permanent registrations MUST include a description of the stream type, including the layout semantics of the stream contents.

The entries in the following table are registered by this document.

Stream Type	Value	Specification	Sender
Control Stream	0x00	Section 6.2.1	Both
Push Stream	0x01	Section 4.4	Server

Table 5

Additionally, each code of the format $0x1f * N + 0x21$ for non-negative integer values of N (that is, 0x21, 0x40, ..., through 0x3FFFFFFFFFFFFFFE) MUST NOT be assigned by IANA.

12. References

12.1. Normative References

[ALTSVC]

Nottingham, M., McManus, P., and J. Reschke, "HTTP Alternative Services", RFC 7838, DOI 10.17487/RFC7838, April 2016, <<https://www.rfc-editor.org/info/rfc7838>>.

[CACHING] Fielding, R., Nottingham, M., and J. Reschke, "HTTP Caching", Work in Progress, Internet-Draft, draft-ietf-httpbis-cache-07, 7 March 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-httpbis-cache-07.txt>>.

[HTTP-REPLAY] Thomson, M., Nottingham, M., and W. Tarreau, "Using Early Data in HTTP", RFC 8470, DOI 10.17487/RFC8470, September 2018, <<https://www.rfc-editor.org/info/rfc8470>>.

[QPACK] Krasic, C., Bishop, M., and A. Frindell, Ed., "QPACK: Header Compression for HTTP over QUIC", Work in Progress, Internet-Draft, draft-ietf-quic-qpack-15, 20 May 2020, <<https://tools.ietf.org/html/draft-ietf-quic-qpack-15>>.

[QUIC-TRANSPORT] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", Work in Progress, Internet-Draft, draft-ietf-quic-transport-28, 20 May 2020, <<https://tools.ietf.org/html/draft-ietf-quic-transport-28>>.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/info/rfc793>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

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

[RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <<https://www.rfc-editor.org/info/rfc5234>>.

[RFC6066] Eastlake 3rd, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", RFC 6066, DOI

10.17487/RFC6066, January 2011, <<https://www.rfc-editor.org/info/rfc6066>>.

[RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <<https://www.rfc-editor.org/info/rfc6125>>.

[RFC6265] Barth, A., "HTTP State Management Mechanism", RFC 6265, DOI 10.17487/RFC6265, April 2011, <<https://www.rfc-editor.org/info/rfc6265>>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

[RFC8164] Nottingham, M. and M. Thomson, "Opportunistic Security for HTTP/2", RFC 8164, DOI 10.17487/RFC8164, May 2017, <<https://www.rfc-editor.org/info/rfc8164>>.

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

[SEMANTICS] Fielding, R., Nottingham, M., and J. Reschke, "HTTP Semantics", Work in Progress, Internet-Draft, draft-ietf-httpbis-semantics-07, 7 March 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-httpbis-semantics-07.txt>>.

12.2. Informative References

[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>>.

[HPACK] Peon, R. and H. Ruellan, "HPACK: Header Compression for HTTP/2", RFC 7541, DOI 10.17487/RFC7541, May 2015, <<https://www.rfc-editor.org/info/rfc7541>>.

[HTTP11] Fielding, R., Nottingham, M., and J. Reschke, "HTTP/1.1 Messaging", Work in Progress, Internet-Draft, draft-ietf-httpbis-messaging-07, 7 March 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-httpbis-messaging-07.txt>>.

[HTTP2] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", RFC 7540, DOI

10.17487/RFC7540, May 2015, <<https://www.rfc-editor.org/info/rfc7540>>.

- [RFC6585] Nottingham, M. and R. Fielding, "Additional HTTP Status Codes", RFC 6585, DOI 10.17487/RFC6585, April 2012, <<https://www.rfc-editor.org/info/rfc6585>>.
- [RFC7301] Friedl, S., Popov, A., Langley, A., and E. Stephan, "Transport Layer Security (TLS) Application-Layer Protocol Negotiation Extension", RFC 7301, DOI 10.17487/RFC7301, July 2014, <<https://www.rfc-editor.org/info/rfc7301>>.
- [TF0] Cheng, Y., Chu, J., Radhakrishnan, S., and A. Jain, "TCP Fast Open", RFC 7413, DOI 10.17487/RFC7413, December 2014, <<https://www.rfc-editor.org/info/rfc7413>>.
- [TLS13] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [URI] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.

Appendix A. Considerations for Transitioning from HTTP/2

HTTP/3 is strongly informed by HTTP/2, and bears many similarities. This section describes the approach taken to design HTTP/3, points out important differences from HTTP/2, and describes how to map HTTP/2 extensions into HTTP/3.

HTTP/3 begins from the premise that similarity to HTTP/2 is preferable, but not a hard requirement. HTTP/3 departs from HTTP/2 where QUIC differs from TCP, either to take advantage of QUIC features (like streams) or to accommodate important shortcomings (such as a lack of total ordering). These differences make HTTP/3 similar to HTTP/2 in key aspects, such as the relationship of requests and responses to streams. However, the details of the HTTP/3 design are substantially different than HTTP/2.

These departures are noted in this section.

A.1. Streams

HTTP/3 permits use of a larger number of streams ($2^{62}-1$) than HTTP/2. The considerations about exhaustion of stream identifier space apply, though the space is significantly larger such that it is

likely that other limits in QUIC are reached first, such as the limit on the connection flow control window.

In contrast to HTTP/2, stream concurrency in HTTP/3 is managed by QUIC. QUIC considers a stream closed when all data has been received and sent data has been acknowledged by the peer. HTTP/2 considers a stream closed when the frame containing the `END_STREAM` bit has been committed to the transport. As a result, the stream for an equivalent exchange could remain "active" for a longer period of time. HTTP/3 servers might choose to permit a larger number of concurrent client-initiated bidirectional streams to achieve equivalent concurrency to HTTP/2, depending on the expected usage patterns.

Due to the presence of other unidirectional stream types, HTTP/3 does not rely exclusively on the number of concurrent unidirectional streams to control the number of concurrent in-flight pushes. Instead, HTTP/3 clients use the `MAX_PUSH_ID` frame to control the number of pushes received from an HTTP/3 server.

A.2. HTTP Frame Types

Many framing concepts from HTTP/2 can be elided on QUIC, because the transport deals with them. Because frames are already on a stream, they can omit the stream number. Because frames do not block multiplexing (QUIC's multiplexing occurs below this layer), the support for variable-maximum-length packets can be removed. Because stream termination is handled by QUIC, an `END_STREAM` flag is not required. This permits the removal of the `Flags` field from the generic frame layout.

Frame payloads are largely drawn from [\[HTTP2\]](#). However, QUIC includes many features (e.g., flow control) which are also present in HTTP/2. In these cases, the HTTP mapping does not re-implement them. As a result, several HTTP/2 frame types are not required in HTTP/3. Where an HTTP/2-defined frame is no longer used, the frame ID has been reserved in order to maximize portability between HTTP/2 and HTTP/3 implementations. However, even equivalent frames between the two mappings are not identical.

Many of the differences arise from the fact that HTTP/2 provides an absolute ordering between frames across all streams, while QUIC provides this guarantee on each stream only. As a result, if a frame type makes assumptions that frames from different streams will still be received in the order sent, HTTP/3 will break them.

Some examples of feature adaptations are described below, as well as general guidance to extension frame implementors converting an HTTP/2 extension to HTTP/3.

A.2.1. Prioritization Differences

HTTP/2 specifies priority assignments in PRIORITY frames and (optionally) in HEADERS frames. HTTP/3 does not provide a means of signaling priority.

Note that while there is no explicit signaling for priority, this does not mean that prioritization is not important for achieving good performance.

A.2.2. Field Compression Differences

HPACK was designed with the assumption of in-order delivery. A sequence of encoded field sections must arrive (and be decoded) at an endpoint in the same order in which they were encoded. This ensures that the dynamic state at the two endpoints remains in sync.

Because this total ordering is not provided by QUIC, HTTP/3 uses a modified version of HPACK, called QPACK. QPACK uses a single unidirectional stream to make all modifications to the dynamic table, ensuring a total order of updates. All frames which contain encoded fields merely reference the table state at a given time without modifying it.

[[QPACK](#)] provides additional details.

A.2.3. Guidance for New Frame Type Definitions

Frame type definitions in HTTP/3 often use the QUIC variable-length integer encoding. In particular, Stream IDs use this encoding, which allows for a larger range of possible values than the encoding used in HTTP/2. Some frames in HTTP/3 use an identifier rather than a Stream ID (e.g., Push IDs). Redefinition of the encoding of extension frame types might be necessary if the encoding includes a Stream ID.

Because the Flags field is not present in generic HTTP/3 frames, those frames which depend on the presence of flags need to allocate space for flags as part of their frame payload.

Other than this issue, frame type HTTP/2 extensions are typically portable to QUIC simply by replacing Stream 0 in HTTP/2 with a control stream in HTTP/3. HTTP/3 extensions will not assume ordering, but would not be harmed by ordering, and would be portable to HTTP/2 in the same manner.

A.2.4. Mapping Between HTTP/2 and HTTP/3 Frame Types

DATA (0x0): Padding is not defined in HTTP/3 frames. See [Section 7.2.1](#).

HEADERS (0x1):

The PRIORITY region of HEADERS is not defined in HTTP/3 frames. Padding is not defined in HTTP/3 frames. See [Section 7.2.2](#).

PRIORITY (0x2): As described in [Appendix A.2.1](#), HTTP/3 does not provide a means of signaling priority.

RST_STREAM (0x3): RST_STREAM frames do not exist, since QUIC provides stream lifecycle management. The same code point is used for the CANCEL_PUSH frame ([Section 7.2.3](#)).

SETTINGS (0x4): SETTINGS frames are sent only at the beginning of the connection. See [Section 7.2.4](#) and [Appendix A.3](#).

PUSH_PROMISE (0x5): The PUSH_PROMISE does not reference a stream; instead the push stream references the PUSH_PROMISE frame using a Push ID. See [Section 7.2.5](#).

PING (0x6): PING frames do not exist, since QUIC provides equivalent functionality.

GOAWAY (0x7): GOAWAY does not contain an error code. In the client to server direction, it carries a Push ID instead of a server initiated stream ID. See [Section 7.2.6](#).

WINDOW_UPDATE (0x8): WINDOW_UPDATE frames do not exist, since QUIC provides flow control.

CONTINUATION (0x9): CONTINUATION frames do not exist; instead, larger HEADERS/PUSH_PROMISE frames than HTTP/2 are permitted.

Frame types defined by extensions to HTTP/2 need to be separately registered for HTTP/3 if still applicable. The IDs of frames defined in [\[HTTP2\]](#) have been reserved for simplicity. Note that the frame type space in HTTP/3 is substantially larger (62 bits versus 8 bits), so many HTTP/3 frame types have no equivalent HTTP/2 code points. See [Section 11.2.1](#).

A.3. HTTP/2 SETTINGS Parameters

An important difference from HTTP/2 is that settings are sent once, as the first frame of the control stream, and thereafter cannot change. This eliminates many corner cases around synchronization of changes.

Some transport-level options that HTTP/2 specifies via the SETTINGS frame are superseded by QUIC transport parameters in HTTP/3. The HTTP-level options that are retained in HTTP/3 have the same value as in HTTP/2.

Below is a listing of how each HTTP/2 SETTINGS parameter is mapped:

SETTINGS_HEADER_TABLE_SIZE: See [\[QPACK\]](#).

SETTINGS_ENABLE_PUSH: This is removed in favor of the MAX_PUSH_ID which provides a more granular control over server push.

SETTINGS_MAX_CONCURRENT_STREAMS: QUIC controls the largest open Stream ID as part of its flow control logic. Specifying SETTINGS_MAX_CONCURRENT_STREAMS in the SETTINGS frame is an error.

SETTINGS_INITIAL_WINDOW_SIZE: QUIC requires both stream and connection flow control window sizes to be specified in the initial transport handshake. Specifying SETTINGS_INITIAL_WINDOW_SIZE in the SETTINGS frame is an error.

SETTINGS_MAX_FRAME_SIZE: This setting has no equivalent in HTTP/3. Specifying it in the SETTINGS frame is an error.

SETTINGS_MAX_FIELD_SECTION_SIZE: See [Section 7.2.4.1](#).

In HTTP/3, setting values are variable-length integers (6, 14, 30, or 62 bits long) rather than fixed-length 32-bit fields as in HTTP/2. This will often produce a shorter encoding, but can produce a longer encoding for settings which use the full 32-bit space. Settings ported from HTTP/2 might choose to redefine their value to limit it to 30 bits for more efficient encoding, or to make use of the 62-bit space if more than 30 bits are required.

Settings need to be defined separately for HTTP/2 and HTTP/3. The IDs of settings defined in [\[HTTP2\]](#) have been reserved for simplicity. Note that the settings identifier space in HTTP/3 is substantially larger (62 bits versus 16 bits), so many HTTP/3 settings have no equivalent HTTP/2 code point. See [Section 11.2.2](#).

As QUIC streams might arrive out-of-order, endpoints are advised to not wait for the peers' settings to arrive before responding to other streams. See [Section 7.2.4.2](#).

A.4. HTTP/2 Error Codes

QUIC has the same concepts of "stream" and "connection" errors that HTTP/2 provides. However, the differences between HTTP/2 and HTTP/3 mean that error codes are not directly portable between versions.

The HTTP/2 error codes defined in Section 7 of [\[HTTP2\]](#) logically map to the HTTP/3 error codes as follows:

NO_ERROR (0x0): H3_NO_ERROR in [Section 8.1](#).

PROTOCOL_ERROR (0x1):

This is mapped to H3_GENERAL_PROTOCOL_ERROR except in cases where more specific error codes have been defined. This includes H3_FRAME_UNEXPECTED and H3_CLOSED_CRITICAL_STREAM defined in [Section 8.1](#).

INTERNAL_ERROR (0x2): H3_INTERNAL_ERROR in [Section 8.1](#).

FLOW_CONTROL_ERROR (0x3): Not applicable, since QUIC handles flow control.

SETTINGS_TIMEOUT (0x4): Not applicable, since no acknowledgement of SETTINGS is defined.

STREAM_CLOSED (0x5): Not applicable, since QUIC handles stream management.

FRAME_SIZE_ERROR (0x6): H3_FRAME_ERROR error code defined in [Section 8.1](#).

REFUSED_STREAM (0x7): H3_REQUEST_REJECTED (in [Section 8.1](#)) is used to indicate that a request was not processed. Otherwise, not applicable because QUIC handles stream management.

CANCEL (0x8): H3_REQUEST_CANCELLED in [Section 8.1](#).

COMPRESSION_ERROR (0x9): Multiple error codes are defined in [\[QPACK\]](#).

CONNECT_ERROR (0xa): H3_CONNECT_ERROR in [Section 8.1](#).

ENHANCE_YOUR_CALM (0xb): H3_EXCESSIVE_LOAD in [Section 8.1](#).

INADEQUATE_SECURITY (0xc): Not applicable, since QUIC is assumed to provide sufficient security on all connections.

H3_1_1_REQUIRED (0xd): H3_VERSION_FALLBACK in [Section 8.1](#).

Error codes need to be defined for HTTP/2 and HTTP/3 separately. See [Section 11.2.3](#).

A.4.1. Mapping Between HTTP/2 and HTTP/3 Errors

An intermediary that converts between HTTP/2 and HTTP/3 may encounter error conditions from either upstream. It is useful to communicate the occurrence of error to the downstream but error codes largely reflect connection-local problems that generally do not make sense to propagate.

An intermediary that encounters an error from an upstream origin can indicate this by sending an HTTP status code such as 502, which is suitable for a broad class of errors.

There are some rare cases where it is beneficial to propagate the error by mapping it to the closest matching error type to the receiver. For example, an intermediary that receives an HTTP/2 stream error of type REFUSED_STREAM from the origin has a clear signal that the request was not processed and that the request is safe to retry. Propagating this error condition to the client as an HTTP/3 stream error of type H3_REQUEST_REJECTED allows the client to take the action it deems most appropriate. In the reverse direction the intermediary might deem it beneficial to pass on client request cancellations that are indicated by terminating a stream with H3_REQUEST_CANCELLED.

Conversion between errors is described in the logical mapping. The error codes are defined in non-overlapping spaces in order to protect against accidental conversion that could result in the use of inappropriate or unknown error codes for the target version. An intermediary is permitted to promote stream errors to connection errors but they should be aware of the cost to the connection for what might be a temporary or intermittent error.

Appendix B. Change Log

RFC Editor's Note: Please remove this section prior to publication of a final version of this document.

B.1. Since draft-ietf-quic-http-27

- *Updated text to refer to latest HTTP revisions
- *Use the HTTP definition of authority for establishing and coalescing connections (#253, #2223, #3558)
- *Define use of GOAWAY from both endpoints (#2632, #3129)
- *Require either :authority or Host if the URI scheme has a mandatory authority component (#3408, #3475)

B.2. Since draft-ietf-quic-http-26

- *No changes

B.3. Since draft-ietf-quic-http-25

- *Require QUICv1 for HTTP/3 (#3117, #3323)

*Remove DUPLICATE_PUSH and allow duplicate PUSH_PROMISE (#3275, #3309)

*Clarify the definition of "malformed" (#3352, #3345)

B.4. Since draft-ietf-quick-http-24

*Removed H3_EARLY_RESPONSE error code; H3_NO_ERROR is recommended instead (#3130, #3208)

*Unknown error codes are equivalent to H3_NO_ERROR (#3276, #3331)

*Some error codes are reserved for greasing (#3325, #3360)

B.5. Since draft-ietf-quick-http-23

*Removed quick Alt-Svc parameter (#3061, #3118)

*Clients need not persist unknown settings for use in 0-RTT (#3110, #3113)

*Clarify error cases around CANCEL_PUSH (#2819, #3083)

B.6. Since draft-ietf-quick-http-22

*Removed priority signaling (#2922, #2924)

*Further changes to error codes (#2662, #2551):

- Error codes renumbered

- HTTP_MALFORMED_FRAME replaced by HTTP_FRAME_ERROR, HTTP_ID_ERROR, and others

*Clarify how unknown frame types interact with required frame sequence (#2867, #2858)

*Describe interactions with the transport in terms of defined interface terms (#2857, #2805)

*Require the use of the http-opportunistic resource (RFC 8164) when scheme is http (#2439, #2973)

*Settings identifiers cannot be duplicated (#2979)

*Changes to SETTINGS frames in 0-RTT (#2972, #2790, #2945):

- Servers must send all settings with non-default values in their SETTINGS frame, even when resuming

- If a client doesn't have settings associated with a 0-RTT ticket, it uses the defaults
- Servers can't accept early data if they cannot recover the settings the client will have remembered
- *Clarify that Upgrade and the 101 status code are prohibited (#2898, #2889)
- *Clarify that frame types reserved for greasing can occur on any stream, but frame types reserved due to HTTP/2 correspondence are prohibited (#2997, #2692, #2693)
- *Unknown error codes cannot be treated as errors (#2998, #2816)

B.7. Since draft-ietf-quick-http-21

No changes

B.8. Since draft-ietf-quick-http-20

- *Prohibit closing the control stream (#2509, #2666)
- *Change default priority to use an orphan node (#2502, #2690)
- *Exclusive priorities are restored (#2754, #2781)
- *Restrict use of frames when using CONNECT (#2229, #2702)
- *Close and maybe reset streams if a connection error occurs for CONNECT (#2228, #2703)
- *Encourage provision of sufficient unidirectional streams for QPACK (#2100, #2529, #2762)
- *Allow extensions to use server-initiated bidirectional streams (#2711, #2773)
- *Clarify use of maximum header list size setting (#2516, #2774)
- *Extensive changes to error codes and conditions of their sending
 - Require connection errors for more error conditions (#2511, #2510)
 - Updated the error codes for illegal GOAWAY frames (#2714, #2707)
 - Specified error code for HEADERS on control stream (#2708)

- Specified error code for servers receiving PUSH_PROMISE (#2709)
- Specified error code for receiving DATA before HEADERS (#2715)
- Describe malformed messages and their handling (#2410, #2764)
- Remove HTTP_PUSH_ALREADY_IN_CACHE error (#2812, #2813)
- Refactor Push ID related errors (#2818, #2820)
- Rationalize HTTP/3 stream creation errors (#2821, #2822)

B.9. Since draft-ietf-quic-http-19

- *SETTINGS_NUM_PLACEHOLDERS is 0x9 (#2443, #2530)
- *Non-zero bits in the Empty field of the PRIORITY frame MAY be treated as an error (#2501)

B.10. Since draft-ietf-quic-http-18

- *Resetting streams following a GOAWAY is recommended, but not required (#2256, #2457)
- *Use variable-length integers throughout (#2437, #2233, #2253, #2275)
 - Variable-length frame types, stream types, and settings identifiers
 - Renumbered stream type assignments
 - Modified associated reserved values
- *Frame layout switched from Length-Type-Value to Type-Length-Value (#2395, #2235)
- *Specified error code for servers receiving DUPLICATE_PUSH (#2497)
- *Use connection error for invalid PRIORITY (#2507, #2508)

B.11. Since draft-ietf-quic-http-17

- *HTTP_REQUEST_REJECTED is used to indicate a request can be retried (#2106, #2325)
- *Changed error code for GOAWAY on the wrong stream (#2231, #2343)

B.12. Since draft-ietf-quic-http-16

- *Rename "HTTP/QUIC" to "HTTP/3" (#1973)

*Changes to PRIORITY frame (#1865, #2075)

- Permitted as first frame of request streams

- Remove exclusive reprioritization

- Changes to Prioritized Element Type bits

*Define DUPLICATE_PUSH frame to refer to another PUSH_PROMISE (#2072)

*Set defaults for settings, allow request before receiving SETTINGS (#1809, #1846, #2038)

*Clarify message processing rules for streams that aren't closed (#1972, #2003)

*Removed reservation of error code 0 and moved HTTP_NO_ERROR to this value (#1922)

*Removed prohibition of zero-length DATA frames (#2098)

B.13. Since draft-ietf-quick-http-15

Substantial editorial reorganization; no technical changes.

B.14. Since draft-ietf-quick-http-14

*Recommend sensible values for QUIC transport parameters (#1720, #1806)

*Define error for missing SETTINGS frame (#1697, #1808)

*Setting values are variable-length integers (#1556, #1807) and do not have separate maximum values (#1820)

*Expanded discussion of connection closure (#1599, #1717, #1712)

*HTTP_VERSION_FALLBACK falls back to HTTP/1.1 (#1677, #1685)

B.15. Since draft-ietf-quick-http-13

*Reserved some frame types for grease (#1333, #1446)

*Unknown unidirectional stream types are tolerated, not errors; some reserved for grease (#1490, #1525)

*Require settings to be remembered for 0-RTT, prohibit reductions (#1541, #1641)

*Specify behavior for truncated requests (#1596, #1643)

B.16. Since draft-ietf-quic-http-12

- *TLS SNI extension isn't mandatory if an alternative method is used (#1459, #1462, #1466)
- *Removed flags from HTTP/3 frames (#1388, #1398)
- *Reserved frame types and settings for use in preserving extensibility (#1333, #1446)
- *Added general error code (#1391, #1397)
- *Unidirectional streams carry a type byte and are extensible (#910, #1359)
- *Priority mechanism now uses explicit placeholders to enable persistent structure in the tree (#441, #1421, #1422)

B.17. Since draft-ietf-quic-http-11

- *Moved QPACK table updates and acknowledgments to dedicated streams (#1121, #1122, #1238)

B.18. Since draft-ietf-quic-http-10

- *Settings need to be remembered when attempting and accepting 0-RTT (#1157, #1207)

B.19. Since draft-ietf-quic-http-09

- *Selected QCRAM for header compression (#228, #1117)
- *The server_name TLS extension is now mandatory (#296, #495)
- *Specified handling of unsupported versions in Alt-Svc (#1093, #1097)

B.20. Since draft-ietf-quic-http-08

- *Clarified connection coalescing rules (#940, #1024)

B.21. Since draft-ietf-quic-http-07

- *Changes for integer encodings in QUIC (#595, #905)
- *Use unidirectional streams as appropriate (#515, #240, #281, #886)
- *Improvement to the description of GOAWAY (#604, #898)
- *Improve description of server push usage (#947, #950, #957)

B.22. Since draft-ietf-quic-http-06

*Track changes in QUIC error code usage (#485)

B.23. Since draft-ietf-quic-http-05

*Made push ID sequential, add MAX_PUSH_ID, remove SETTINGS_ENABLE_PUSH (#709)

*Guidance about keep-alive and QUIC PINGs (#729)

*Expanded text on GOAWAY and cancellation (#757)

B.24. Since draft-ietf-quic-http-04

*Cite RFC 5234 (#404)

*Return to a single stream per request (#245,#557)

*Use separate frame type and settings registries from HTTP/2 (#81)

*SETTINGS_ENABLE_PUSH instead of SETTINGS_DISABLE_PUSH (#477)

*Restored GOAWAY (#696)

*Identify server push using Push ID rather than a stream ID (#702,#281)

*DATA frames cannot be empty (#700)

B.25. Since draft-ietf-quic-http-03

None.

B.26. Since draft-ietf-quic-http-02

*Track changes in transport draft

B.27. Since draft-ietf-quic-http-01

*SETTINGS changes (#181):

-SETTINGS can be sent only once at the start of a connection;
no changes thereafter

-SETTINGS_ACK removed

-Settings can only occur in the SETTINGS frame a single time

-Boolean format updated

- *Alt-Svc parameter changed from "v" to "quic"; format updated (#229)
- *Closing the connection control stream or any message control stream is a fatal error (#176)
- *HPACK Sequence counter can wrap (#173)
- *0-RTT guidance added
- *Guide to differences from HTTP/2 and porting HTTP/2 extensions added (#127,#242)

B.28. Since draft-ietf-quic-http-00

- *Changed "HTTP/2-over-QUIC" to "HTTP/QUIC" throughout (#11,#29)
- *Changed from using HTTP/2 framing within Stream 3 to new framing format and two-stream-per-request model (#71,#72,#73)
- *Adopted SETTINGS format from draft-bishop-httpbis-extended-settings-01
- *Reworked SETTINGS_ACK to account for indeterminate inter-stream order (#75)
- *Described CONNECT pseudo-method (#95)
- *Updated ALPN token and Alt-Svc guidance (#13,#87)
- *Application-layer-defined error codes (#19,#74)

B.29. Since draft-shade-quic-http2-mapping-00

- *Adopted as base for draft-ietf-quic-http
- *Updated authors/editors list

Acknowledgements

The original authors of this specification were Robbie Shade and Mike Warres.

The IETF QUIC Working Group received an enormous amount of support from many people. Among others, the following people provided substantial contributions to this document:

- *Bence Beky
- *Daan De Meyer

*Martin Duke

*Roy Fielding

*Alan Frindell

*Alessandro Ghedini

*Nick Harper

*Ryan Hamilton

*Christian Huitema

*Subodh Iyengar

*Robin Marx

*Patrick McManus

*Luca Nicco

* (Kazuho Oku)

*Lucas Pardue

*Roberto Peon

*Julian Reschke

*Eric Rescorla

*Martin Seemann

*Ben Schwartz

*Ian Swett

*Willy Taureau

*Martin Thomson

*Dmitri Tikhonov

*Tatsuhiro Tsujikawa

A portion of Mike's contribution was supported by Microsoft during his employment there.

Author's Address

Mike Bishop (editor)
Akamai

Email: mbishop@evequefou.be