

webtrans
Internet-Draft
Intended status: Standards Track
Expires: 8 September 2022

A. Frindell
Facebook Inc.
E. Kinnear
T. Pauly
Apple Inc.
M. Thomson
Mozilla
V. Vasiliev
Google
G. Xie
Facebook Inc.
7 March 2022

WebTransport using HTTP/2
draft-ietf-webtrans-http2-03

Abstract

WebTransport defines a set of low-level communications features designed for client-server interactions that are initiated by Web clients. This document describes a protocol that can provide many of the capabilities of WebTransport over HTTP/2. This protocol enables the use of WebTransport when a UDP-based protocol is not available.

Note to Readers

Discussion of this draft takes place on the WebTransport mailing list (webtransport@ietf.org (<mailto:webtransport@ietf.org>)), which is archived at https://mailarchive.ietf.org/arch/search/?email_list=webtransport.

The repository tracking the issues for this draft can be found at <https://github.com/ietf-wg-webtrans/draft-webtransport-http2>. The web API draft corresponding to this document can be found at <https://w3c.github.io/webtransport/>.

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 8 September 2022.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	3
2. Protocol Overview	3
3. Session Establishment	4
3.1. Establishing a Transport-Capable HTTP/2 Connection	5
3.2. Extended CONNECT in HTTP/2	5
3.3. Creating a New Session	5
3.4. Limiting the Number of Simultaneous Sessions	6
4. WebTransport Features	6
4.1. Transport Considerations	6
4.2. WebTransport Stream States	7
5. WebTransport Frames	7
5.1. WT_PADDING Frames	8
5.2. WT_RESET_STREAM Frames	8
5.3. WT_STOP_SENDING Frames	9
5.4. WT_STREAM Frames	9
5.5. WT_MAX_DATA Frames	10
5.6. WT_MAX_STREAM_DATA Frames	10
5.7. WT_MAX_STREAMS Frames	11
5.8. WT_DATA_BLOCKED Frames	12
5.9. WT_STREAM_DATA_BLOCKED Frames	12
5.10. WT_STREAMS_BLOCKED Frames	13
5.11. WT_DATAGRAM Frames	13
6. Examples	14
7. Session Termination	16

8. Transport Properties	16
9. Security Considerations	17
10. IANA Considerations	17
10.1. HTTP/2 SETTINGS Parameter Registration	18
11. References	18
11.1. Normative References	18
11.2. Informative References	19
Acknowledgments	20
Index	20
Authors' Addresses	21

1. Introduction

WebTransport [OVERVIEW] is designed to provide generic communication capabilities to Web clients that use HTTP/3 [HTTP3]. The HTTP/3 WebTransport protocol [WEBTRANSPORT-H3] allows Web clients to use QUIC [QUIC] features such as streams or datagrams [DATAGRAM]. However, there are some environments where QUIC cannot be deployed.

This document defines a protocol that provides all of the core functions of WebTransport using HTTP semantics. This includes unidirectional streams, bidirectional streams, and datagrams.

By relying only on generic HTTP semantics, this protocol might allow deployment using any HTTP version. However, this document only defines negotiation for HTTP/2 [H2] as the current most common TCP-based fallback to HTTP/3.

1.1. Terminology

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

This document follows terminology defined in Section 1.2 of [OVERVIEW]. Note that this document distinguishes between a WebTransport server and an HTTP/2 server. An HTTP/2 server is the server that terminates HTTP/2 connections; a WebTransport server is an application that accepts WebTransport sessions, which can be accessed using HTTP/2 and this protocol.

2. Protocol Overview

WebTransport servers are identified by an HTTPS URI as defined in Section 4.2.2 of [HTTP].

When an HTTP/2 connection is established, both the client and server have to send a `SETTINGS_ENABLE_WEBTRANSPORT` setting in order to indicate that they both support WebTransport over HTTP/2.

A client initiates a WebTransport session by sending an extended `CONNECT` request [RFC8441]. If the server accepts the request, a WebTransport session is established. The stream that carries the `CONNECT` request is used to exchange bidirectional data for the session. This stream will be referred to as a `_CONNECT stream_`. The stream ID of a `CONNECT` stream, which will be referred to as a `_Session ID_`, is used to uniquely identify a given WebTransport session within the connection.

After the session is established, endpoints exchange `_WebTransport frames_` using the bidirectional `CONNECT` stream. Within this stream, `_WebTransport streams_` and `_WebTransport datagrams_` are multiplexed. In HTTP/2, WebTransport frames are carried in HTTP/2 DATA frames. Multiple independent WebTransport sessions can share a connection if the HTTP version supports that, as HTTP/2 does.

WebTransport frames closely mirror a subset of QUIC frames and provide the essential WebTransport features. Within a WebTransport session, endpoints can

- * create and use bidirectional or unidirectional streams with no additional round trips using `WT_STREAM` frames

Stream creation and data flow on streams uses flow control mechanisms modeled on those in QUIC. Flow control is managed using the WebTransport frames: `WT_MAX_DATA`, `WT_MAX_STREAM_DATA`, `WT_MAX_STREAMS`, `WT_DATA_BLOCKED`, `WT_STREAM_DATA_BLOCKED`, and `WT_STREAMS_BLOCKED`. Flow control for the `CONNECT` stream as a whole, as provided by the HTTP version in use, applies in addition to any WebTransport-session-level flow control.

WebTransport streams can be aborted using a `WT_RESET_STREAM` frame and a receiver can request that a sender stop sending with a `WT_STOP_SENDING` frame.

A WebTransport session is terminated when the `CONNECT` stream that created it is closed. This implicitly closes all WebTransport streams that were multiplexed over that `CONNECT` stream.

3. Session Establishment

3.1. Establishing a Transport-Capable HTTP/2 Connection

In order to indicate support for WebTransport, both the client and the server MUST send a `SETTINGS_ENABLE_WEBTRANSPORT` value set to "1" in their `SETTINGS` frame. Endpoints MUST NOT use any WebTransport-related functionality unless the parameter has been negotiated.

3.2. Extended CONNECT in HTTP/2

[RFC8441] defines an extended `CONNECT` method in Section 4, enabled by the `SETTINGS_ENABLE_CONNECT_PROTOCOL` parameter. An endpoint does not need to send both `SETTINGS_ENABLE_CONNECT_PROTOCOL` and `SETTINGS_ENABLE_WEBTRANSPORT`; the `SETTINGS_ENABLE_WEBTRANSPORT` setting implies that an endpoint supports extended `CONNECT`.

3.3. Creating a New Session

As WebTransport sessions are established over HTTP, they are identified using the `https` URI scheme [RFC7230].

In order to create a new WebTransport session, a client can send an HTTP `CONNECT` request. The `:protocol` pseudo-header field ([RFC8441]) MUST be set to `webtransport` (Section 7.1 of [WEBTRANSPORT-H3]). The `:scheme` field MUST be `https`. Both the `:authority` and the `:path` value MUST be set; those fields indicate the desired WebTransport server. In a Web context, the request MUST include an `Origin` header field [ORIGIN] that includes the origin of the site that requested the creation of the session.

Upon receiving an extended `CONNECT` request with a `:protocol` field set to `webtransport`, the HTTP server checks if the identified resource supports WebTransport sessions. If the resource does not, the server SHOULD reply with status code 404 (Section 6.5.4 of [RFC7231]). To accept a WebTransport session the server replies with 2xx status code. Before accepting a session, a server MUST ensure that it authorizes use of the session by the site identified in the `Origin` header.

From the client's perspective, a WebTransport session is established when the client receives a 200 response. From the server's perspective, a session is established once it sends a 200 response. Both endpoints MUST NOT send any WebTransport frames on a given session before that session is established.

3.4. Limiting the Number of Simultaneous Sessions

From a flow control perspective, WebTransport sessions count against HTTP/2 session flow control limits just like regular HTTP requests, since they are established via an HTTP CONNECT request. This document does not make any effort to introduce a separate flow control mechanism for WebTransport sessions. If the server needs to limit the rate of incoming requests, it has alternative mechanisms at its disposal:

- * HTTP_STREAM_REFUSED error code defined in [RFC7540] indicates to the receiving HTTP/2 stack that the request was not processed in any way.
- * HTTP status code 429 indicates that the request was rejected due to rate limiting [RFC6585]. Unlike the previous method, this signal is directly propagated to the application.

4. WebTransport Features

WebTransport over TCP-based HTTP semantics provides the following features described in [OVERVIEW]: unidirectional streams, bidirectional streams, and datagrams, initiated by either endpoint.

WebTransport streams and datagrams that belong to different WebTransport sessions are identified by the CONNECT stream on which they are transmitted, with one WebTransport session consuming one CONNECT stream.

4.1. Transport Considerations

Because WebTransport over TCP-based HTTP semantics relies on the underlying protocols to provide in order and reliable delivery, there are some notable differences from the way in which QUIC handles application data.

Endpoints MUST send stream data in order. As there is no ordering mechanism available for the receiver to reassemble incoming data, receivers assume that all data arriving in STREAM frames is contiguous and in order.

DATAGRAM frames are delivered to the remote WebTransport endpoint reliably, however this does not require that the receiving implementation deliver that data to the application in a reliable manner.

4.2. WebTransport Stream States

WebTransport streams have states that mirror the states of QUIC streams (Section 3 of [RFC9000]) as closely as possible to aid in ease of implementation.

Because WebTransport does not provide an acknowledgement mechanism for WebTransport frames, it relies on the underlying transport's in order delivery to inform stream state transitions. Wherever QUIC relies on receiving an ack for a packet to transition between stream states, WebTransport performs that transition immediately.

5. WebTransport Frames

WebTransport frames mirror their QUIC counterparts as closely as possible to enable maximal reuse of any applicable QUIC infrastructure by implementors.

A WebTransport frame begins with a Frame Type and Frame Length which are followed by zero or more fields that are type-dependent.

```
Frame {  
    Frame Type (i),  
    Frame Length (i),  
    Type-Dependent Fields (...),  
}
```

Figure 1: WebTransport Frame Format

The Frame Type field indicates the type of the frame, defining what type-dependent fields will be present.

The Frame Length field indicates the length of the WebTransport frame, including all type-dependent fields and other information. It does not include the size of the Frame Type or Frame Length fields themselves.

Both of these fields use a variable-length integer encoding (see Section 16 of [RFC9000]), with one exception. To ensure simple and efficient implementations of frame parsing, the frame type and length MUST use the shortest possible encoding. For example, for the frame types defined in this document, this means a single-byte encoding, even though it is possible to encode these values as a two-, four-, or eight-byte variable-length integer.

5.1. WT_PADDING Frames

A WT_PADDING frame (type=0x00) has no semantic value. PADDING frames can be used to introduce additional data between other WebTransport frames and can also be used to provide protection against traffic analysis or for other reasons.

```
WT_PADDING Frame {  
    Type (i) = 0x00,  
    Length (i),  
    Padding (...),  
}
```

Figure 2: WT_PADDING Frame Format

The Padding field MUST be set to an all-zero sequence of bytes of any length as specified by the Length field.

5.2. WT_RESET_STREAM Frames

A WebTransport frame called WT_RESET_STREAM is introduced for either endpoint to abruptly terminate the sending part of a WebTransport stream.

An endpoint uses a WT_RESET_STREAM frame (type=0x04) to abruptly terminate the sending part of a stream.

After sending a WT_RESET_STREAM, an endpoint ceases transmission and retransmission of WT_STREAM frames on the identified stream. A receiver of WT_RESET_STREAM can discard any data that it already received on that stream.

```
WT_RESET_STREAM Frame {  
    Type (i) = 0x04,  
    Length (i),  
    Stream ID (i),  
    Application Protocol Error Code (i),  
}
```

Figure 3: WT_RESET_STREAM Frame Format

The WT_RESET_STREAM frame defines the following fields:

Stream ID: A variable-length integer encoding of the WebTransport stream ID of the stream being terminated.

Application Protocol Error Code: A variable-length integer

containing the application protocol error code that indicates why the stream is being closed.

Unlike the equivalent QUIC frame, this frame does not include a Final Size field. In-order delivery of WT_STREAM frames ensures that the amount of session-level flow control consumed by a stream is always known by both endpoints.

5.3. WT_STOP_SENDING Frames

A WebTransport frame called WT_STOP_SENDING is introduced to communicate that incoming data is being discarded on receipt per application request. WT_STOP_SENDING requests that a peer cease transmission on a stream.

```
WT_STOP_SENDING Frame {  
  Type (i) = 0x05,  
  Length (i),  
  Stream ID (i),  
  Application Protocol Error Code (i),  
}
```

Figure 4: WT_STOP_SENDING Frame Format

The WT_STOP_SENDING frame defines the following fields:

Stream ID: A variable-length integer carrying the WebTransport stream ID of the stream being ignored.

Application Protocol Error Code: A variable-length integer containing the application-specified reason the sender is ignoring the stream.

5.4. WT_STREAM Frames

WT_STREAM frames implicitly create a stream and carry stream data.

The Type field in the WT_STREAM frame is either 0x0a or 0x0b. This uses the same frame types as a QUIC STREAM frame with the OFF bit clear and the LEN bit set. The FIN bit (0x01) in the frame type indicates that the frame marks the end of the stream in one direction. Stream data consists of any number of 0x0a frames followed by a terminal 0x0b frame.

```
WT_STREAM Frame {  
    Type (i) = 0x0a..0x0b,  
    Length (i),  
    Stream ID (i),  
    Stream Data (..),  
}
```

Figure 5: WT_STREAM Frame Format

WT_STREAM frames contain the following fields:

Stream ID: The stream ID for the stream.

Stream Data: Zero or more bytes of data for the stream. Empty WT_STREAM frames MUST NOT be used unless they open or close a stream; an endpoint MAY treat an empty WT_STREAM frame that neither starts nor ends a stream as a session error.

5.5. WT_MAX_DATA Frames

A WebTransport frame called WT_MAX_DATA is introduced to inform the peer of the maximum amount of data that can be sent on the WebTransport session as a whole.

```
WT_MAX_DATA Frame {  
    Type (i) = 0x10,  
    Length (i),  
    Maximum Data (i),  
}
```

Figure 6: WT_MAX_DATA Frame Format

WT_MAX_DATA frames contain the following field:

Maximum Data: A variable-length integer indicating the maximum amount of data that can be sent on the entire connection, in units of bytes.

All data sent in WT_STREAM frames counts toward this limit. The sum of the lengths of Stream Data fields in WT_STREAM frames MUST NOT exceed the value advertised by a receiver.

5.6. WT_MAX_STREAM_DATA Frames

A WebTransport frame called WT_MAX_STREAM_DATA is introduced to inform a peer of the maximum amount of data that can be sent on a stream.

```
WT_MAX_STREAM_DATA Frame {  
    Type (i) = 0x11,  
    Length (i),  
    Stream ID (i),  
    Maximum Stream Data (i),  
}
```

Figure 7: WT_MAX_STREAM_DATA Frame Format

WT_MAX_STREAM_DATA frames contain the following fields:

Stream ID: The stream ID of the affected WebTransport stream, encoded as a variable-length integer.

Maximum Stream Data: A variable-length integer indicating the maximum amount of data that can be sent on the identified stream, in units of bytes.

All data sent in WT_STREAM frames for the identified stream counts toward this limit. The sum of the lengths of Stream Data fields in WT_STREAM frames on the identified stream MUST NOT exceed the value advertised by a receiver.

5.7. WT_MAX_STREAMS Frames

A WebTransport frame called WT_MAX_STREAMS is introduced to inform the peer of the cumulative number of streams of a given type it is permitted to open. A WT_MAX_STREAMS frame with a type of 0x12 applies to bidirectional streams, and a WT_MAX_STREAMS frame with a type of 0x13 applies to unidirectional streams.

```
WT_MAX_STREAMS Frame {  
    Type (i) = 0x12..0x13,  
    Length (i),  
    Maximum Streams (i),  
}
```

Figure 8: WT_MAX_STREAMS Frame Format

WT_MAX_STREAMS frames contain the following field:

Maximum Streams: A count of the cumulative number of streams of the corresponding type that can be opened over the lifetime of the connection. This value cannot exceed 2^{60} , as it is not possible to encode stream IDs larger than $2^{62}-1$.

An endpoint MUST NOT open more streams than permitted by the current stream limit set by its peer. For instance, a server that receives a unidirectional stream limit of 3 is permitted to open streams 3, 7, and 11, but not stream 15.

Note that this limit includes streams that have been closed as well as those that are open.

5.8. WT_DATA_BLOCKED Frames

A sender SHOULD send a WT_DATA_BLOCKED frame (type=0x14) when it wishes to send data but is unable to do so due to WebTransport session-level flow control. WT_DATA_BLOCKED frames can be used as input to tuning of flow control algorithms.

```
WT_DATA_BLOCKED Frame {  
    Type (i) = 0x14,  
    Length (i),  
    Maximum Data (i),  
}
```

Figure 9: WT_DATA_BLOCKED Frame Format

WT_DATA_BLOCKED frames contain the following field:

Maximum Data: A variable-length integer indicating the session-level limit at which blocking occurred.

5.9. WT_STREAM_DATA_BLOCKED Frames

A sender SHOULD send a WT_STREAM_DATA_BLOCKED frame (type=0x15) when it wishes to send data but is unable to do so due to stream-level flow control. This frame is analogous to WT_DATA_BLOCKED.

```
WT_STREAM_DATA_BLOCKED Frame {  
    Type (i) = 0x15,  
    Length (i),  
    Stream ID (i),  
    Maximum Stream Data (i),  
}
```

Figure 10: WT_STREAM_DATA_BLOCKED Frame Format

WT_STREAM_DATA_BLOCKED frames contain the following fields:

Stream ID: A variable-length integer indicating the WebTransport stream that is blocked due to flow control.

Maximum Stream Data: A variable-length integer indicating the offset of the stream at which the blocking occurred.

5.10. WT_STREAMS_BLOCKED Frames

A sender SHOULD send a WT_STREAMS_BLOCKED frame (type=0x16 or 0x17) when it wishes to open a stream but is unable to do so due to the maximum stream limit set by its peer. A WT_STREAMS_BLOCKED frame of type 0x16 is used to indicate reaching the bidirectional stream limit, and a STREAMS_BLOCKED frame of type 0x17 is used to indicate reaching the unidirectional stream limit.

A WT_STREAMS_BLOCKED frame does not open the stream, but informs the peer that a new stream was needed and the stream limit prevented the creation of the stream.

```
WT_STREAMS_BLOCKED Frame {  
  Type (i) = 0x16..0x17,  
  Length (i),  
  Maximum Streams (i),  
}
```

Figure 11: WT_STREAMS_BLOCKED Frame Format

WT_STREAMS_BLOCKED frames contain the following field:

Maximum Streams: A variable-length integer indicating the maximum number of streams allowed at the time the frame was sent. This value cannot exceed 2^{60} , as it is not possible to encode stream IDs larger than $2^{62}-1$.

5.11. WT_DATAGRAM Frames

The WT_DATAGRAM frame type (0x31) is used to carry datagram traffic. Frame type 0x30 is also reserved to maintain parity with QUIC, but unused, as all WebTransport frames MUST contain a length field.

```
WT_DATAGRAM Frame {  
  Type (i) = 0x31,  
  Length (i),  
  Datagram Data (..),  
}
```

Figure 12: WT_DATAGRAM Frame Format

WT_DATAGRAM frames contain the following fields:

Datagram Data: The content of the datagram to be delivered.

The data in WT_DATAGRAM frames is not subject to flow control. The receiver MAY discard this data if it does not have sufficient space to buffer it.

An intermediary could forward the data in a WT_DATAGRAM frame over another protocol, such as WebTransport over HTTP/3. In QUIC, a datagram frame can span at most one packet. Because of that, the applications have to know the maximum size of the datagram they can send. However, when proxying the datagrams, the hop-by-hop MTUs can vary.

6. Examples

An example of negotiating a WebTransport Stream on an HTTP/2 connection follows. This example is intended to closely follow the example in Section 5.1 of [RFC8441] to help illustrate the differences defined in this document.

```
[[ From Client ]]
```

```
SETTINGS
```

```
SETTINGS_ENABLE_WEBTRANSPORT = 1
```

```
HEADERS + END_HEADERS
```

```
Stream ID = 3
```

```
:method = CONNECT
```

```
:protocol = webtransport
```

```
:scheme = https
```

```
:path = /
```

```
:authority = server.example.com
```

```
origin: server.example.com
```

```
[[ From Server ]]
```

```
SETTINGS
```

```
SETTINGS_ENABLE_WEBTRANSPORT = 1
```

```
HEADERS + END_HEADERS
```

```
Stream ID = 3
```

```
:status = 200
```

```
WT_STREAM
```

```
Stream ID = 5
```

```
WebTransport Data
```

```
WT_STREAM + FIN
```

```
Stream ID = 5
```

```
WebTransport Data
```

```
WT_STREAM + FIN
```

```
Stream ID = 5
```

```
WebTransport Data
```

An example of the server initiating a WebTransport Stream follows.
The only difference here is the endpoint that sends the first
WT_STREAM frame.

<pre> [[From Client]] SETTINGS SETTINGS_ENABLE_WEBTRANSPORT = 1 HEADERS + END_HEADERS Stream ID = 3 :method = CONNECT :protocol = webtransport :scheme = https :path = / :authority = server.example.com origin: server.example.com WT_STREAM + FIN Stream ID = 2 WebTransport Data </pre>	<pre> [[From Server]] SETTINGS SETTINGS_ENABLE_WEBTRANSPORT = 1 HEADERS + END_HEADERS Stream ID = 3 :status = 200 WT_STREAM Stream ID = 2 WebTransport Data WT_STREAM + FIN Stream ID = 2 WebTransport Data </pre>
---	--

7. Session Termination

An WebTransport session over HTTP/2 is terminated when either endpoint closes the stream associated with the CONNECT request that initiated the session. Upon learning about the session being terminated, the endpoint MUST stop sending new datagrams and reset all of the streams associated with the session.

8. Transport Properties

The WebTransport framework [OVERVIEW] defines a set of optional transport properties that clients can use to determine the presence of features which might allow additional optimizations beyond the common set of properties available via all WebTransport protocols. Below are details about support in Http2Transport for those properties.

Stream Independence: Http2Transport does not support stream independence, as HTTP/2 inherently has head of line blocking.

Partial Reliability: Http2Transport does not support partial reliability, as HTTP/2 retransmits any lost data. This means that any datagrams sent via Http2Transport will be retransmitted regardless of the preference of the application. The receiver is permitted to drop them, however, if it is unable to buffer them.

Pooling Support: Http2Transport supports pooling, as multiple transports using Http2Transport may share the same underlying HTTP/2 connection and therefore share a congestion controller and other transport context.

Connection Mobility: Http2Transport does not support connection mobility, unless an underlying transport protocol that supports multipath or migration, such as MPTCP [MPTCP], is used underneath HTTP/2 and TLS. Without such support, Http2Transport connections cannot survive network transitions.

9. Security Considerations

WebTransport over HTTP/2 satisfies all of the security requirements imposed by [OVERVIEW] on WebTransport protocols, thus providing a secure framework for client-server communication in cases when the client is potentially untrusted.

WebTransport over HTTP/2 requires explicit opt-in through the use of HTTP SETTINGS; this avoids potential protocol confusion attacks by ensuring the HTTP/2 server explicitly supports it. It also requires the use of the Origin header, providing the server with the ability to deny access to Web-based clients that do not originate from a trusted origin.

Just like HTTP traffic going over HTTP/2, WebTransport pools traffic to different origins within a single connection. Different origins imply different trust domains, meaning that the implementations have to treat each transport as potentially hostile towards others on the same connection. One potential attack is a resource exhaustion attack: since all of the transports share both congestion control and flow control context, a single client aggressively using up those resources can cause other transports to stall. The user agent thus SHOULD implement a fairness scheme that ensures that each transport within connection gets a reasonable share of controlled resources; this applies both to sending data and to opening new streams.

10. IANA Considerations

10.1. HTTP/2 SETTINGS Parameter Registration

The following entry is added to the "HTTP/2 Settings" registry established by [RFC7540]:

The SETTINGS_ENABLE_WEBTRANSPORT parameter indicates that the specified HTTP/2 connection is WebTransport-capable.

Setting Name: ENABLE_WEBTRANSPORT

Value: 0x2b603742

Default: 0

Specification: This document

11. References

11.1. Normative References

- [H2] Thomson, M. and C. Benfield, "HTTP/2", Work in Progress, Internet-Draft, draft-ietf-httpbis-http2bis-07, 24 January 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-httpbis-http2bis-07>>.
- [HTTP] Fielding, R. T., Nottingham, M., and J. Reschke, "HTTP Semantics", Work in Progress, Internet-Draft, draft-ietf-httpbis-semantics-19, 12 September 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-httpbis-semantics-19>>.
- [ORIGIN] Barth, A., "The Web Origin Concept", RFC 6454, DOI 10.17487/RFC6454, December 2011, <<https://www.rfc-editor.org/rfc/rfc6454>>.
- [OVERVIEW] Vasiliev, V., "The WebTransport Protocol Framework", Work in Progress, Internet-Draft, draft-ietf-webtrans-overview-03, 7 March 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-webtrans-overview-03>>.
- [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/rfc/rfc2119>>.
- [RFC6585] Nottingham, M. and R. Fielding, "Additional HTTP Status Codes", RFC 6585, DOI 10.17487/RFC6585, April 2012, <<https://www.rfc-editor.org/rfc/rfc6585>>.

- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<https://www.rfc-editor.org/rfc/rfc7230>>.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <<https://www.rfc-editor.org/rfc/rfc7231>>.
- [RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", RFC 7540, DOI 10.17487/RFC7540, May 2015, <<https://www.rfc-editor.org/rfc/rfc7540>>.
- [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/rfc/rfc8174>>.
- [RFC8441] McManus, P., "Bootstrapping WebSockets with HTTP/2", RFC 8441, DOI 10.17487/RFC8441, September 2018, <<https://www.rfc-editor.org/rfc/rfc8441>>.
- [RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <<https://www.rfc-editor.org/rfc/rfc9000>>.
- [WEBTRANSPORT-H3]
Frindell, A., Kinnear, E., and V. Vasiliev, "WebTransport over HTTP/3", Work in Progress, Internet-Draft, draft-ietf-webtrans-http3-02, 25 October 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-webtrans-http3-02>>.

11.2. Informative References

- [DATAGRAM] Pauly, T., Kinnear, E., and D. Schinazi, "An Unreliable Datagram Extension to QUIC", Work in Progress, Internet-Draft, draft-ietf-quic-datagram-10, 4 February 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-quic-datagram-10>>.
- [HTTP3] Bishop, M., "Hypertext Transfer Protocol Version 3 (HTTP/3)", Work in Progress, Internet-Draft, draft-ietf-quic-http-34, 2 February 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-quic-http-34>>.

- [MPTCP] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, DOI 10.17487/RFC6824, January 2013, <<https://www.rfc-editor.org/rfc/rfc6824>>.
- [QUIC] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <<https://www.rfc-editor.org/rfc/rfc9000>>.

Acknowledgments

Thanks to Anthony Chivetta, Joshua Otto, and Valentin Pistol for their contributions in the design and implementation of this work.

Index

W

W

WT_DATAGRAM Section 5.11, Paragraph 1; Section 5.11, Paragraph 3; Section 5.11, Paragraph 5; Section 5.11, Paragraph 6

WT_DATA_BLOCKED Section 2, Paragraph 7; Section 5.8, Paragraph 1; Section 5.8, Paragraph 1; Section 5.8, Paragraph 3; Section 5.9, Paragraph 1

WT_MAX_DATA Section 2, Paragraph 7; Section 5.5, Paragraph 1; Section 5.5, Paragraph 3

WT_MAX_STREAMS Section 2, Paragraph 7; Section 5.7, Paragraph 1; Section 5.7, Paragraph 1; Section 5.7, Paragraph 1; Section 5.7, Paragraph 3

WT_MAX_STREAM_DATA Section 2, Paragraph 7; Section 5.6, Paragraph 1; Section 5.6, Paragraph 3

WT_PADDING Section 5.1, Paragraph 1

WT_RESET_STREAM Section 2, Paragraph 8; Section 5.2, Paragraph 1; Section 5.2, Paragraph 2; Section 5.2, Paragraph 3; Section 5.2, Paragraph 3; Section 5.2, Paragraph 5

WT_STOP_SENDING Section 2, Paragraph 8; Section 5.3, Paragraph 1; Section 5.3, Paragraph 1; Section 5.3, Paragraph 3

WT_STREAM Section 2, Paragraph 6, Item 1; Section 5.2, Paragraph 3; Section 5.2, Paragraph 7; Section 5.4, Paragraph 1; Section 5.4, Paragraph 2; Section 5.4, Paragraph 4; Section 5.4, Paragraph 5.4.1; Section 5.4, Paragraph 5.4.1; Section 5.5, Paragraph 5; Section 5.5, Paragraph 5; Section 5.6, Paragraph 5; Section 5.6, Paragraph 5; Section 6, Paragraph 3

WT_STREAMS_BLOCKED Section 2, Paragraph 7; Section 5.10,

Paragraph 1; Section 5.10, Paragraph 1; Section 5.10,
Paragraph 2; Section 5.10, Paragraph 4
WT_STREAM_DATA_BLOCKED Section 2, Paragraph 7; Section 5.9,
Paragraph 1; Section 5.9, Paragraph 3

Authors' Addresses

Alan Frindell
Facebook Inc.
Email: afrind@fb.com

Eric Kinnear
Apple Inc.
One Apple Park Way
Cupertino, California 95014,
United States of America
Email: ekinnear@apple.com

Tommy Pauly
Apple Inc.
One Apple Park Way
Cupertino, California 95014,
United States of America
Email: tpaully@apple.com

Martin Thomson
Mozilla
Email: mt@lowentropy.net

Victor Vasiliev
Google
Email: vasilvv@google.com

Guowu Xie
Facebook Inc.
Email: woo@fb.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 28 April 2022

A. Frindell
Facebook
E. Kinnear
Apple Inc.
V. Vasiliev
Google
25 October 2021

WebTransport over HTTP/3
draft-ietf-webtrans-http3-02

Abstract

WebTransport [OVERVIEW] is a protocol framework that enables clients constrained by the Web security model to communicate with a remote server using a secure multiplexed transport. This document describes a WebTransport protocol that is based on HTTP/3 [HTTP3] and provides support for unidirectional streams, bidirectional streams and datagrams, all multiplexed within the same HTTP/3 connection.

Note to Readers

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

The repository tracking the issues for this draft can be found at <https://github.com/ietf-wg-webtrans/draft-ietf-webtrans-http3/issues>. The web API draft corresponding to this document can be found at <https://w3c.github.io/webtransport/>.

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 28 April 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	3
1.1. Terminology	3
2. Protocol Overview	3
3. Session Establishment	4
3.1. Establishing a Transport-Capable HTTP/3 Connection . . .	4
3.2. Extended CONNECT in HTTP/3	4
3.3. Creating a New Session	4
3.4. Limiting the Number of Simultaneous Sessions	5
4. WebTransport Features	5
4.1. Unidirectional streams	6
4.2. Bidirectional Streams	6
4.3. Resetting Data Streams	7
4.4. Datagrams	8
4.5. Buffering Incoming Streams and Datagrams	8
5. Session Termination	9
6. Negotiating the Draft Version	10
7. Security Considerations	10
8. IANA Considerations	11
8.1. Upgrade Token Registration	11
8.2. HTTP/3 SETTINGS Parameter Registration	11
8.3. Frame Type Registration	12
8.4. Stream Type Registration	12
8.5. HTTP/3 Error Code Registration	12
8.6. Datagram Format Type	13
9. References	13
9.1. Normative References	13
9.2. Informative References	15
Authors' Addresses	15

1. Introduction

HTTP/3 [HTTP3] is a protocol defined on top of QUIC [RFC9000] that can multiplex HTTP requests over a QUIC connection. This document defines a mechanism for multiplexing non-HTTP data with HTTP/3 in a manner that conforms with the WebTransport protocol requirements and semantics [OVERVIEW]. Using the mechanism described here, multiple WebTransport instances can be multiplexed simultaneously with regular HTTP traffic on the same HTTP/3 connection.

1.1. Terminology

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

This document follows terminology defined in Section 1.2 of [OVERVIEW]. Note that this document distinguishes between a WebTransport server and an HTTP/3 server. An HTTP/3 server is the server that terminates HTTP/3 connections; a WebTransport server is an application that accepts WebTransport sessions, which can be accessed via an HTTP/3 server.

2. Protocol Overview

WebTransport servers in general are identified by a pair of authority value and path value (defined in [RFC3986] Sections 3.2 and 3.3 correspondingly).

When an HTTP/3 connection is established, both the client and server have to send a `SETTINGS_ENABLE_WEBTRANSPORT` setting in order to indicate that they both support WebTransport over HTTP/3.

WebTransport sessions are initiated inside a given HTTP/3 connection by the client, who sends an extended `CONNECT` request [RFC8441]. If the server accepts the request, an WebTransport session is established. The resulting stream will be further referred to as a `_CONNECT stream_`, and its stream ID is used to uniquely identify a given WebTransport session within the connection. The ID of the `CONNECT` stream that established a given WebTransport session will be further referred to as a `_Session ID_`.

After the session is established, the peers can exchange data using the following mechanisms:

- * A client can create a bidirectional stream using a special indefinite-length HTTP/3 frame that transfers ownership of the stream to WebTransport.
- * A server can create a bidirectional stream, which is possible since HTTP/3 does not define any semantics for server-initiated bidirectional streams.
- * Both client and server can create a unidirectional stream using a special stream type.
- * A datagram can be sent using HTTP Datagrams [HTTP-DATAGRAM].

An WebTransport session is terminated when the CONNECT stream that created it is closed.

3. Session Establishment

3.1. Establishing a Transport-Capable HTTP/3 Connection

In order to indicate support for WebTransport, both the client and the server MUST send a `SETTINGS_ENABLE_WEBTRANSPORT` value set to "1" in their `SETTINGS` frame. The `SETTINGS_ENABLE_WEBTRANSPORT` parameter value SHALL be either "0" or "1", with "0" being the default; an endpoint that receives a value other than "0" or "1" MUST close the connection with the `H3_SETTINGS_ERROR` error code.

The client MUST NOT send a WebTransport request until it has received the setting indicating WebTransport support from the server. Similarly, the server MUST NOT process any incoming WebTransport requests until the client settings have been received, as the client may be using a version of WebTransport extension that is different from the one used by the server.

3.2. Extended CONNECT in HTTP/3

[RFC8441] defines an extended `CONNECT` method in Section 4, enabled by the `SETTINGS_ENABLE_CONNECT_PROTOCOL` parameter. That parameter is only defined for HTTP/2. This document does not create a new multi-purpose parameter to indicate support for extended `CONNECT` in HTTP/3; instead, the `SETTINGS_ENABLE_WEBTRANSPORT` setting implies that an endpoint supports extended `CONNECT`.

3.3. Creating a New Session

As WebTransport sessions are established over HTTP/3, they are identified using the `https` URI scheme [RFC7230].

In order to create a new WebTransport session, a client can send an HTTP CONNECT request. The `:protocol` pseudo-header field ([RFC8441]) MUST be set to `webtransport`. The `:scheme` field MUST be `https`. Both the `:authority` and the `:path` value MUST be set; those fields indicate the desired WebTransport server. An Origin header [RFC6454] MUST be provided within the request.

Upon receiving an extended CONNECT request with a `:protocol` field set to `webtransport`, the HTTP/3 server can check if it has a WebTransport server associated with the specified `:authority` and `:path` values. If it does not, it SHOULD reply with status code 404 (Section 6.5.4, [RFC7231]). If it does, it MAY accept the session by replying with a 2xx series status code, as defined in Section 15.3 of [SEMANTICS]. The WebTransport server MUST verify the Origin header to ensure that the specified origin is allowed to access the server in question.

From the client's perspective, a WebTransport session is established when the client receives a 2xx response. From the server's perspective, a session is established once it sends a 2xx response. WebTransport over HTTP/3 does not support 0-RTT.

The `webtransport` HTTP Upgrade Token uses the Capsule Protocol as defined in [HTTP-DATAGRAM].

3.4. Limiting the Number of Simultaneous Sessions

From the flow control perspective, WebTransport sessions count against the stream flow control just like regular HTTP requests, since they are established via an HTTP CONNECT request. This document does not make any effort to introduce a separate flow control mechanism for sessions, nor to separate HTTP requests from WebTransport data streams. If the server needs to limit the rate of incoming requests, it has alternative mechanisms at its disposal:

- * HTTP_REQUEST_REJECTED error code defined in [HTTP3] indicates to the receiving HTTP/3 stack that the request was not processed in any way.
- * HTTP status code 429 indicates that the request was rejected due to rate limiting [RFC6585]. Unlike the previous method, this signal is directly propagated to the application.

4. WebTransport Features

WebTransport over HTTP/3 provides the following features described in [OVERVIEW]: unidirectional streams, bidirectional streams and datagrams, initiated by either endpoint.

Session IDs are used to demultiplex streams and datagrams belonging to different WebTransport sessions. On the wire, session IDs are encoded using the QUIC variable length integer scheme described in [RFC9000].

If at any point a session ID is received that cannot a valid ID for a client-initiated bidirectional stream, the receipient MUST close the connection with an H3_ID_ERROR error code.

4.1. Unidirectional streams

Once established, both endpoints can open unidirectional streams. The HTTP/3 unidirectional stream type SHALL be 0x54. The body of the stream SHALL be the stream type, followed by the session ID, encoded as a variable-length integer, followed by the user-specified stream data (Figure 1).

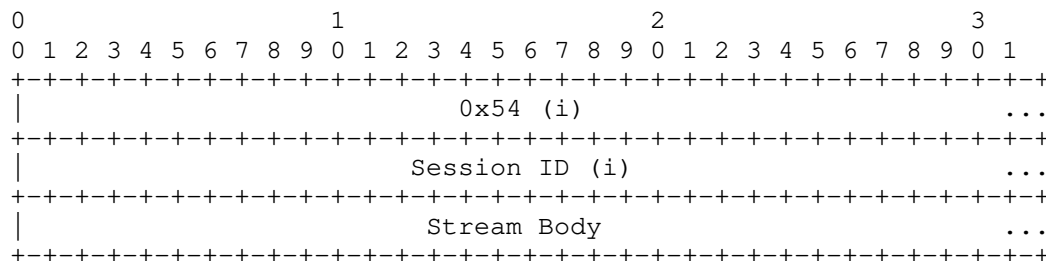


Figure 1: Unidirectional WebTransport stream format

4.2. Bidirectional Streams

WebTransport endpoints can initiate bidirectional streams by opening an HTTP/3 bidirectional stream and sending an HTTP/3 frame with type WEBTRANSPORT_STREAM (type=0x41). The format of the frame SHALL be the frame type, followed by the session ID, encoded as a variable-length integer, followed by the user-specified stream data (Figure 2). The frame SHALL last until the end of the stream.

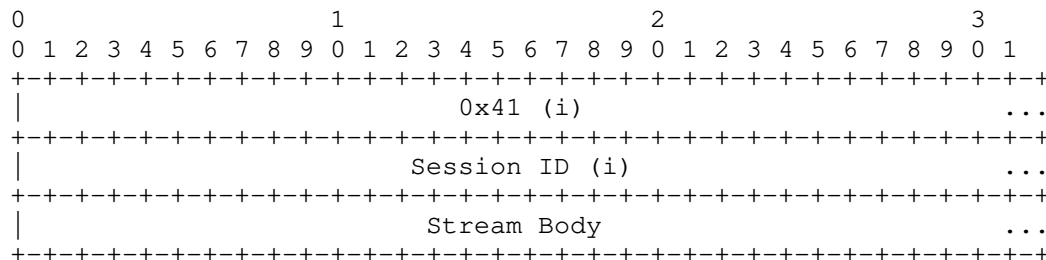


Figure 2: WEBTRANSPORT_STREAM frame format

HTTP/3 does not by itself define any semantics for server-initiated bidirectional streams. If WebTransport setting is negotiated by both endpoints, the syntax of the server-initiated bidirectional streams SHALL be the same as the syntax of client-initiated bidirectional streams, that is, a sequence of HTTP/3 frames. The only frame defined by this document for use within server-initiated bidirectional streams is WEBTRANSPORT_STREAM.

TODO: move the paragraph above into a separate draft; define what happens with already existing HTTP/3 frames on server-initiated bidirectional streams.

4.3. Resetting Data Streams

A WebTransport endpoint may send a RESET_STREAM or a STOP_SENDING frame for a WebTransport data stream. Those signals are propagated by the WebTransport implementation to the application.

A WebTransport application SHALL provide an error code for those operations. Since WebTransport shares the error code space with HTTP/3, WebTransport application errors for streams are limited to an unsigned 8-bit integer, assuming values between 0x00 and 0xff. WebTransport implementations SHALL remap those error codes into an error range where 0x00 corresponds to 0x52e4a40fa8db, and 0xff corresponds to 0x52e4a40fa9e2. Note that there are code points inside that range of form "0x1f * N + 0x21" that are reserved by Section 8.1 of [HTTP3]; those have to be accounted for when mapping the error codes by skipping them (i.e. the two HTTP/3 error codepoints adjacent to a GREASE codepoint would map to two adjacent WebTransport application error codepoints). An example pseudocode can be seen in Figure 3.

```
first = 0x52e4a40fa8db
last = 0x52e4a40fa9e2

def webtransport_code_to_http_code(n):
    return first + n + floor(n / 0x1e)

def http_code_to_webtransport_code(h):
    assert(first <= h <= last)
    assert((h - 0x21) % 0x1f != 0)
    shifted = h - first
    return shifted - shifted // 0x1f
```

Figure 3: Pseudocode for converting between WebTransport application errors and HTTP/3 error codes; here, `//` is integer division

WebTransport data streams are associated with sessions through a header at the beginning of the stream; resetting a stream may result in that data being discarded. Because of that, WebTransport application error codes are best effort, as the WebTransport stack is not always capable of associating the reset code with a session. The only exception is the situation where there is only one session on a given HTTP/3 connection, and no intermediaries between the client and the server.

WebTransport implementations SHALL forward the error code for a stream associated with a known session to the application that owns that session; similarly, the intermediaries SHALL reset the streams with corresponding error code when receiving a reset from the peer. If a WebTransport implementation intentionally allows only one session over a given HTTP/3 connection, it SHALL forward the error codes within WebTransport application error code range to the application that owns the only session on that connection.

4.4. Datagrams

Datagrams can be sent using HTTP Datagrams, using the WEB_TRANSPORT HTTP Datagram Format Type (see value in Section 8.6). When using the WEB_TRANSPORT HTTP Datagram Format Type, the WebTransport datagram payload is sent unmodified in the "HTTP Datagram Payload" field of an HTTP Datagram. When sending a registration capsule using the "Datagram Format Type" set to WEB_TRANSPORT, the "Datagram Format Additional Data" field SHALL be empty.

In QUIC, a datagram frame can span at most one packet. Because of that, the applications have to know the maximum size of the datagram they can send. However, when proxying the datagrams, the hop-by-hop MTUs can vary. TODO: Describe how the path MTU can be computed, specifically propagation across HTTP proxies.

4.5. Buffering Incoming Streams and Datagrams

In WebTransport over HTTP/3, the client MAY send its SETTINGS frame, as well as multiple WebTransport CONNECT requests, WebTransport data streams and WebTransport datagrams, all within a single flight. As those can arrive out of order, a WebTransport server could be put into a situation where it receives a stream or a datagram without a corresponding session. Similarly, a client may receive a server-initiated stream or a datagram before receiving the CONNECT response headers from the server.

To handle this case, WebTransport endpoints SHOULD buffer streams and datagrams until those can be associated with an established session. To avoid resource exhaustion, the endpoints MUST limit the number of buffered streams and datagrams. When the number of buffered streams is exceeded, a stream SHALL be closed by sending a RESET_STREAM and/or STOP_SENDING with the H3_WEBTRANSPORT_BUFFERED_STREAM_REJECTED error code. When the number of buffered datagrams is exceeded, a datagram SHALL be dropped. It is up to an implementation to choose what stream or datagram to discard.

5. Session Termination

A WebTransport session over HTTP/3 is considered terminated when either of the following conditions is met:

- * the CONNECT stream is closed, either cleanly or abruptly, on either side; or
- * a CLOSE_WEBTRANSPORT_SESSION capsule is either sent or received.

Upon learning that the session has been terminated, the endpoint MUST reset all of the streams associated with the session; it MUST NOT send any new datagrams or open any new streams.

To terminate a session with a detailed error message, an application MAY send an HTTP capsule [HTTP-DATAGRAM] of type CLOSE_WEBTRANSPORT_SESSION (0x2843). The format of the capsule SHALL be as follows:

```
CLOSE_WEBTRANSPORT_SESSION Capsule {  
  Type (i) = CLOSE_WEBTRANSPORT_SESSION,  
  Length (i),  
  Application Error Code (32),  
  Application Error Message (..8192),  
}
```

CLOSE_WEBTRANSPORT_SESSION has the following fields:

Application Error Code: A 32-bit error code provided by the application closing the connection.

Application Error Message: A UTF-8 encoded error message string provided by the application closing the connection. The message takes up the remainder of the capsule, and its length MUST NOT exceed 1024 bytes.

A CLOSE_WEBTRANSPORT_SESSION capsule MUST be followed by a FIN on the sender side. If any additional stream data is received on the CONNECT stream after CLOSE_WEBTRANSPORT_SESSION, the stream MUST be reset with code H3_MESSAGE_ERROR. The recipient MUST close the stream upon receiving a FIN. If the sender of CLOSE_WEBTRANSPORT_SESSION does not receive a FIN after some time, it SHOULD send STOP_SENDING on the CONNECT stream.

Cleanly terminating a CONNECT stream without a CLOSE_WEBTRANSPORT_SESSION capsule SHALL be semantically equivalent to terminating it with a CLOSE_WEBTRANSPORT_SESSION capsule that has an error code of 0 and an empty error string.

6. Negotiating the Draft Version

[[RFC editor: please remove this section before publication.]]

WebTransport over HTTP/3 uses two different mechanisms to negotiate versions for the different parts of the draft.

The hop-by-hop wire format aspects of the protocol are negotiated by changing the codepoint used for the SETTINGS_ENABLE_WEBTRANSPORT parameter. Because of that, any WebTransport endpoint MUST wait for the peer's SETTINGS frame before sending or processing any WebTransport traffic. When multiple versions are supported by both of the peers, the most recent version supported by both is selected.

The data exchanged over the CONNECT stream is transmitted across intermediaries, and thus cannot be versioned using a SETTINGS parameter. To indicate support for different versions of the protocol defined in this draft, the clients SHALL send a header for each version of the draft supported. The header corresponding to the version described in this draft is Sec-Webtransport-Http3-Draft02; its value SHALL be 1. The server SHALL reply with a Sec-Webtransport-Http3-Draft header indicating the selected version; its value SHALL be draft02 for the version described in this draft.

7. Security Considerations

WebTransport over HTTP/3 satisfies all of the security requirements imposed by [OVERVIEW] on WebTransport protocols, thus providing a secure framework for client-server communication in cases when the client is potentially untrusted.

WebTransport over HTTP/3 requires explicit opt-in through the use of a QUIC transport parameter; this avoids potential protocol confusion attacks by ensuring the HTTP/3 server explicitly supports it. It also requires the use of the Origin header, providing the server with the ability to deny access to Web-based clients that do not originate from a trusted origin.

Just like HTTP traffic going over HTTP/3, WebTransport pools traffic to different origins within a single connection. Different origins imply different trust domains, meaning that the implementations have to treat each transport as potentially hostile towards others on the same connection. One potential attack is a resource exhaustion attack: since all of the transports share both congestion control and flow control context, a single client aggressively using up those resources can cause other transports to stall. The user agent thus **SHOULD** implement a fairness scheme that ensures that each transport within connection gets a reasonable share of controlled resources; this applies both to sending data and to opening new streams.

8. IANA Considerations

8.1. Upgrade Token Registration

The following entry is added to the "Hypertext Transfer Protocol (HTTP) Upgrade Token Registry" registry established by [RFC7230]:

The "webtransport" label identifies HTTP/3 used as a protocol for WebTransport:

Value: webtransport

Description: WebTransport over HTTP/3

Reference: This document and [I-D.ietf-webtrans-http2]

8.2. HTTP/3 SETTINGS Parameter Registration

The following entry is added to the "HTTP/3 Settings" registry established by [HTTP3]:

The `SETTINGS_ENABLE_WEBTRANSPORT` parameter indicates that the specified HTTP/3 connection is WebTransport-capable.

Setting Name: `ENABLE_WEBTRANSPORT`

Value: `0x2b603742`

Default: `0`

Specification: This document

8.3. Frame Type Registration

The following entry is added to the "HTTP/3 Frame Type" registry established by [HTTP3]:

The WEBTRANSPORT_STREAM frame allows HTTP/3 client-initiated bidirectional streams to be used by WebTransport:

Code: 0x41

Frame Type: WEBTRANSPORT_STREAM

Specification: This document

8.4. Stream Type Registration

The following entry is added to the "HTTP/3 Stream Type" registry established by [HTTP3]:

The "WebTransport stream" type allows unidirectional streams to be used by WebTransport:

Code: 0x54

Stream Type: WebTransport stream

Specification: This document

Sender: Both

8.5. HTTP/3 Error Code Registration

The following entry is added to the "HTTP/3 Error Code" registry established by [HTTP3]:

Name: H3_WEBTRANSPORT_BUFFERED_STREAM_REJECTED

Value: 0x3994bd84

Description: WebTransport data stream rejected due to lack of associated session.

Specification: This document.

In addition, the following range of entries is registered:

Name: H3_WEBTRANSPORT_APPLICATION_00 ...
H3_WEBTRANSPORT_APPLICATION_FF

Value: 0x52e4a40fa8db to 0x52e4a40fa9e2 inclusive, with the exception of 0x52e4a40fa8f9, 0x52e4a40fa918, 0x52e4a40fa937, 0x52e4a40fa956, 0x52e4a40fa975, 0x52e4a40fa994, 0x52e4a40fa9b3, and 0x52e4a40fa9d2.

Description: WebTransport application error codes.

Specification: This document.

8.6. Datagram Format Type

This document will request IANA to register WEB_TRANSPORT in the "HTTP Datagram Format Types" registry established by [HTTP-DATAGRAM].

Type	Value	Specification
WEB_TRANSPORT	0xff7c00	This Document

Table 1: Registered Datagram Format Type

9. References

9.1. Normative References

[HTTP-DATAGRAM]

Schinazi, D. and L. Pardue, "Using Datagrams with HTTP", Work in Progress, Internet-Draft, draft-ietf-masque-h3-datagram-05, 25 October 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-masque-h3-datagram-05>>.

[HTTP3]

Bishop, M., Ed., "Hypertext Transfer Protocol Version 3 (HTTP/3)", Work in Progress, Internet-Draft, draft-ietf-quic-http, <<https://datatracker.ietf.org/doc/html/draft-ietf-quic-http>>.

[OVERVIEW]

Vasiliev, V., "The WebTransport Protocol Framework", Work in Progress, Internet-Draft, draft-ietf-webtrans-overview-latest, <<https://datatracker.ietf.org/doc/html/draft-ietf-webtrans-overview-latest>>.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://doi.org/10.17487/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://doi.org/10.17487/RFC3986>>.
- [RFC6454] Barth, A., "The Web Origin Concept", RFC 6454, DOI 10.17487/RFC6454, December 2011, <<https://doi.org/10.17487/RFC6454>>.
- [RFC6585] Nottingham, M. and R. Fielding, "Additional HTTP Status Codes", RFC 6585, DOI 10.17487/RFC6585, April 2012, <<https://doi.org/10.17487/RFC6585>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<https://doi.org/10.17487/RFC7230>>.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <<https://doi.org/10.17487/RFC7231>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://doi.org/10.17487/RFC8174>>.
- [RFC8441] McManus, P., "Bootstrapping WebSockets with HTTP/2", RFC 8441, DOI 10.17487/RFC8441, September 2018, <<https://doi.org/10.17487/RFC8441>>.
- [RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <<https://doi.org/10.17487/RFC9000>>.
- [SEMANTICS] Fielding, R. T., Nottingham, M., and J. Reschke, "HTTP Semantics", Work in Progress, Internet-Draft, draft-ietf-httpbis-semantics-19, 12 September 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-httpbis-semantics-19>>.

9.2. Informative References

[I-D.ietf-webtrans-http2]

Frindell, A., Kinnear, E., Pauly, T., Vasiliev, V., and G. Xie, "WebTransport using HTTP/2", Work in Progress, Internet-Draft, draft-ietf-webtrans-http2-01, 30 July 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-webtrans-http2-01>>.

Authors' Addresses

Alan Frindell
Facebook

Email: afrind@fb.com

Eric Kinnear
Apple Inc.

Email: ekinnear@apple.com

Victor Vasiliev
Google

Email: vasilvv@google.com