The WebTransport Protocol Framework

Abstract

The WebTransport Protocol Framework enables clients constrained by
the Web security model to communicate with a remote server using a
secure multiplexed transport. It consists of a set of individual
protocols that are safe to expose to untrusted applications,
combined with a model that allows them to be used interchangeably.

This document defines the overall requirements on the protocols used
in WebTransport, as well as the common features of the protocols,
support for some of which may be optional.

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-overview/
issues>. The web API draft corresponding to this document can be
found at <https://wicg.github.io/web-transport/>.

Copyright Notice

Table of Contents

1. Introduction

The WebTransport Protocol Framework enables clients constrained by the Web security model to communicate with a remote server using a secure multiplexed transport. It consists of a set of individual protocols that are safe to expose to untrusted applications, combined with a model that allows them to be used interchangeably.

This document defines the overall requirements on the protocols used in WebTransport, as well as the common features of the protocols, support for some of which may be optional.

1.1. Background

Historically, web applications that needed a bidirectional data stream between a client and a server could rely on WebSockets [RFC6455], a message-based protocol compatible with the Web security model. However, since the abstraction it provides is a single ordered stream of messages, it suffers from head-of-line blocking (HOLB), meaning that all messages must be sent and received in order even if they are independent and some of them are no longer needed. This makes it a poor fit for latency-sensitive applications which rely on partial reliability and stream independence for performance.

One existing option available to Web developers are WebRTC data channels [I-D.ietf-rtcweb-data-channel], which provide a WebSocket-like API for a peer-to-peer SCTP channel protected by DTLS. In theory, it is possible to use it for the use cases addressed by this specification. However, in practice, its use in non-browser-to-browser settings has been quite low due to its dependency on ICE (which fits poorly with the Web model) and userspace SCTP (which has very few implementations available).

An alternative design would be to layer WebSockets over HTTP/3 [I-D.ietf-quic-http] in a manner similar to how they are currently layered over HTTP/2 [RFC8441]. That would avoid head-of-line blocking and provide an ability to cancel a stream by closing the corresponding WebSocket object. However, this approach has a number of drawbacks, which all stem primarily from the fact that semantically each WebSocket is a completely independent entity:

  *Each new stream would require a WebSocket handshake to agree on application protocol used, meaning that it would take at least one RTT to establish each new stream before the client can write to it.

  *Only clients can initiate streams. Server-initiated streams and other alternative modes of communication (such as the QUIC DATAGRAM frame [I-D.pauly-quic-datagram]) are not available.

  *While the streams would normally be pooled by the user agent, this is not guaranteed, and the general process of mapping a WebSocket to a server is opaque to the client. This introduces unpredictable performance properties into the system, and prevents optimizations which rely on the streams being on the same connection (for instance, it might be possible for the client to request different retransmission priorities for different streams, but that would be much more complex unless they are all on the same connection).

The WebTransport protocol framework avoids all of those issues by letting applications create a single transport object that can contain multiple streams multiplexed together in a single context (similar to SCTP, HTTP/2, QUIC and others), and can be also used to send unreliable datagrams (similar to UDP).

1.2. Conventions and Definitions

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.

WebTransport is a framework that aims to abstract away the underlying transport protocol while still exposing a few key transport-layer aspects to application developers. It is structured around the following concepts:

Transport session: A transport session is a single communication context established between a client and a server. It may correspond to a specific transport-layer connection, or it may be a logical entity within an existing multiplexed transport-layer connection. Transport sessions are logically independent from one another even if some sessions can share an underlying transport-layer connection.

Transport protocol: A transport protocol (WebTransport protocol in contexts where this might be ambiguous) is an instantiation of WebTransport over a given transport-layer protocol.

Datagram: A datagram is a unit of transmission that is treated atomically.

Stream: A stream is a sequence of bytes that is reliably delivered to the receiving application in the same order as it was transmitted by the sender. Streams can be of arbitrary length, and therefore cannot always be buffered entirely in memory. It is expected for transport protocols and APIs to provide partial stream data to the application before the stream has been entirely received.

Message: A message is a stream that is sufficiently small that it can be fully buffered before being passed to the application. WebTransport does not define messages as a primitive, since from the transport perspective they can be simulated by fully buffering a stream before passing it to the application. However, this distinction is important to highlight since some of the similar protocols and APIs (notably WebSocket [RFC6455]) use messages as a core abstraction.

Transport property:
A transport property is a specific behavior that may or may not be exhibited by a transport. Some of those are inherent for all instances of a given transport protocol (TCP-based transport cannot support unreliable delivery), while others can vary even within the same protocol (QUIC connections may or may not support connection migration).

Server: A WebTransport server is an application that accepts incoming transport sessions.

Client: A WebTransport client is an application that initiates the transport session and may be running in a constrained security context, for instance, a JavaScript application running inside a browser.

User agent: A WebTransport user agent is a software system that has an unrestricted access to the host network stack and can create transports on behalf of the client.

2. Common Transport Requirements

Since clients are not necessarily trusted and have to be constrained by the Web security model, WebTransport imposes certain requirements on any specific transport protocol used.

Any transport protocol used MUST use TLS [RFC8446] or a semantically equivalent security protocol (for instance, DTLS [I-D.ietf-tls-dtls13]). The protocols SHOULD use TLS version 1.3 or later, unless they aim for backwards compatibility with legacy systems.

Any transport protocol used MUST require the user agent to obtain and maintain explicit consent from the server to send data. For connection-oriented protocols (such as TCP or QUIC), the connection establishment and keep-alive mechanisms suffice. STUN Consent Freshness [RFC7675] is another example of the mechanism satisfying this requirement.

Any transport protocol used MUST limit the rate at which the client sends data. This SHOULD be accomplished via a feedback-based congestion control mechanism (such as [RFC5681] or [I-D.ietf-quic-recovery]).

Any transport protocol used MUST support simultaneously establishing multiple sessions between the same client and server.

Any transport protocol used MUST prevent the clients from establishing transport sessions to network endpoints that are not WebTransport servers.

Any transport protocol used MUST provide a way for servers to filter clients that can access it by checking the initiating origin [RFC6454].

Any transport protocol used MUST provide a way for a server endpoint location to be described using a URI [RFC3986]. This enables integration with various Web platform features that represent resources as URIs, such as Content Security Policy [CSP].

3. Session Establishment

WebTransport session establishment is most often asynchronous, although in some transports it can succeed instantaneously (for instance, if a transport is immediately pooled with an existing connection). A session MUST NOT be considered established until it is secure against replay attacks. For instance, in protocols creating a new TLS 1.3 session [RFC8446], this would mean that the user agent MUST NOT treat the session as established until it received a Finished message from the server.

In some cases, the transport protocol might allow transmitting data before the session is established; an example is TLS 0-RTT data. Since this data can be replayed by attackers, it MUST NOT be used unless the client has explicitly requested 0-RTT for specific streams or datagrams it knows to be safely replayable.

4. Transport Features

The following transport features are defined in this document. This list is not meant to be comprehensive; future documents may define new features for both new and already existing transports.

All transport protocols MUST provide datagrams, unidirectional and bidirectional streams in order to make the transport protocols easily interchangeable.

4.1. Datagrams

A datagram is a sequence of bytes that is limited in size (generally to the path MTU) and is not expected to be transmitted reliably. The general goal for WebTransport datagrams is to be similar in behavior to UDP while being subject to common requirements expressed in Section 2.

The WebTransport sender is not expected to retransmit datagrams, though it may if it is using a TCP-based protocol or some other underlying protocol that requires reliable delivery. WebTransport datagrams are not expected to be flow controlled, meaning that the receiver might drop datagrams if the application is not consuming them fast enough.

The application MUST be provided with the maxiumum datagram size
that it can send. The size SHOULD be derived from the result of
performing path MTU discovery.

4.2. Streams

A unidirectional stream is a one-way reliable in-order stream of
bytes where the initiator is the only endpoint that can send data. A
bidirectional stream allows both endpoints to send data and can be
conceptually represented as a pair of unidirectional streams.

The streams are in general expected to follow the semantics and the
state machine of QUIC streams ([I-D.ietf-quic-transport], Sections 2
and 3). TODO: describe the stream state machine explicitly.

A WebTransport stream can be reset, indicating that the endpoint is
not interested in either sending or receiving any data related to
the stream. In that case, the sender is expected to not retransmit
any data that was already sent on that stream.

Streams SHOULD be sufficiently lightweight that they can be used as
messages.

Data sent on a stream is flow controlled by the transport protocol.
In addition to flow controlling stream data, the creation of new
streams is flow controlled as well: an endpoint may only open a
limited number of streams until the peer explicitly allows creating
more streams.

Every stream within a transport has a unique 64-bit number
identifying it. Both unidirectional and bidirectional streams share
the number space. The client and the server have to agree on the
numbering, so it can be referenced in the application payload.
WebTransport does not impose any other specific restrictions on the
structure of stream IDs, and they should be treated as opaque 64-bit
blobs.

4.3. Protocol-Specific Features

In addition to features described above, there are some capabilities
that may be provided by an individual protocol but are not
universally applicable to all protocols. Those are allowed, but any
protocol is expected to be useful without those features, and
portable clients should not rely on them.

A notable class of protocol-specific features are features available
only in non-pooled transports. Since those transports have a
dedicated connection, a user agent can provide clients with an
extensive amount of transport-level data that would be too noisy and
difficult to interpret when the connection is shared with unrelated

traffic. For instance, a user agent can provide the number of packets lost, or the number of times stream data was delayed due to flow control. It can also expose variables related to congestion control, such as the size of the congestion window or the current pacing rate.

4.4. Bandwidth Prediction

Using congestion control state and transport metrics, the client can predict the rate at which it can send data. That is essential for many WebTransport use cases; for instance, real time media applications adapt the video bitrate to be a fraction of the throughput they expect to be available. While not all transport protocols can provide low-level transport details, all protocols SHOULD provide the client with an estimate of the available bandwidth.

5. Buffering and Prioritization

TODO: expand this outline into a full summary.

  *Datagrams are intended for low-latency communications, so the buffers for them should be small, and prioritized over stream data.

  *In general, the transport should not apply aggregation algorithms (e.g., Nagle's algorithm [RFC0896]) to datagrams.

6. Transport Properties

In addition to common requirements, each transport can have multiple optional properties associated with it. Querying them allows the client to ascertain the presence of features it can use without requiring knowledge of all protocols. This allows introducing new transports as drop-in replacements for existing ones.

The following properties are defined in this specification:

  *Stream independence. This indicates that there is no head of line blocking between different streams.

  *Partial reliability. This indicates that if a stream is reset, none of the data sent on it will be retransmitted. This also indicates that datagrams will not be retransmitted.

  *Pooling support. Indicates that multiple transports using this transport protocol may end up sharing the same transport layer connection, and thus share a congestion controller and other contexts.

*Connection mobility. Indicates that the transport may continue
    existing even if the network path between the client and the
    server changes.

7. Security Considerations

   Providing untrusted clients with a reasonably low-level access to
   the network comes with risks. This document mitigates those risks by
   imposing a set of common requirements described in Section 2.

   WebTransport mandates the use of TLS for all protocols implementing
   it. This has a dual purpose. On one hand, it protects the transport
   from the network, including both potential attackers and
   ossification by middleboxes. On the other hand, it protects the
   network elements from potential confusion attacks such as the one
   discussed in Section 10.3 of [RFC6455].

   One potential concern is that even when a transport cannot be
   created, the connection error would reveal enough information to
   allow an attacker to scan the network addresses that would normally
   be inaccessible. Because of that, the user agent that runs untrusted
   clients MUST NOT provide any detailed error information until the
   server has confirmed that it is a WebTransport endpoint. For
   example, the client must not be able to distinguish between a
   network address that is unreachable and one that is reachable but is
   not a WebTransport server.

   WebTransport does not support any traditional means of HTTP-based
   authentication. It is not necessarily based on HTTP, and hence does
   not support HTTP cookies or HTTP authentication. Since it requires
   TLS, individual transport protocols MAY expose TLS-based
   authentication capabilities such as client certificates.

8. IANA Considerations

   There are no requests to IANA in this document.

9. References

9.1. Normative References

   [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
              Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/
              RFC2119, March 1997, <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>.

[RFC6454]   Barth, A., "The Web Origin Concept", RFC 6454, DOI
            10.17487/RFC6454, December 2011, <https://www.rfc-
            editor.org/info/rfc6454>.

[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>.

[RFC8446]   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>.

9.2. Informative References

[CSP]   W3C, "Content Security Policy Level 3", October 2020,
            <https://www.w3.org/TR/CSP/>.

[I-D.ietf-quic-http] Bishop, M., "Hypertext Transfer Protocol
            Version 3 (HTTP/3)", Work in Progress, Internet-Draft,
            draft-ietf-quic-http-31, 24 September 2020, <http://
            www.ietf.org/internet-drafts/draft-ietf-quic-
            http-31.txt>.

[I-D.ietf-quic-recovery] Iyengar, J. and I. Swett, "QUIC Loss
            Detection and Congestion Control", Work in Progress,
            Internet-Draft, draft-ietf-quic-recovery-31, 24 September
            2020, <http://www.ietf.org/internet-drafts/draft-ietf-
            quic-recovery-31.txt>.

[I-D.ietf-quic-transport]
            Iyengar, J. and M. Thomson, "QUIC: A UDP-Based
            Multiplexed and Secure Transport", Work in Progress,
            Internet-Draft, draft-ietf-quic-transport-31, 24
            September 2020, <http://www.ietf.org/internet-drafts/
            draft-ietf-quic-transport-31.txt>.

[I-D.ietf-rtcweb-data-channel]
            Jesup, R., Loreto, S., and M. Tuexen, "WebRTC Data
            Channels", Work in Progress, Internet-Draft, draft-ietf-
            rtcweb-data-channel-13, 4 January 2015, <http://
            www.ietf.org/internet-drafts/draft-ietf-rtcweb-data-
            channel-13.txt>.

[I-D.ietf-tls-dtls13] Rescorla, E., Tschofenig, H., and N. Modadugu,
            "The Datagram Transport Layer Security (DTLS) Protocol
            Version 1.3", Work in Progress, Internet-Draft, draft-
            ietf-tls-dtls13-38, 29 May 2020, <http://www.ietf.org/
            internet-drafts/draft-ietf-tls-dtls13-38.txt>.

[I-D.pauly-quic-datagram]
Pauly, T., Kinnear, E., and D. Schinazi, "An Unreliable Datagram Extension to QUIC", Work in Progress, Internet-Draft, draft-pauly-quic-datagram-05, 4 November 2019, <http://www.ietf.org/internet-drafts/draft-pauly-quic-datagram-05.txt>.

[RFC0896] Nagle, J., "Congestion Control in IP/TCP Internetworks", RFC 896, DOI 10.17487/RFC0896, January 1984, <https://www.rfc-editor.org/info/rfc896>.

[RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", RFC 5681, DOI 10.17487/RFC5681, September 2009, <https://www.rfc-editor.org/info/rfc5681>.

[RFC6455] Fette, I. and A. Melnikov, "The WebSocket Protocol", RFC 6455, DOI 10.17487/RFC6455, December 2011, <https://www.rfc-editor.org/info/rfc6455>.

[RFC7675] Perumal, M., Wing, D., Ravindranath, R., Reddy, T., and M. Thomson, "Session Traversal Utilities for NAT (STUN) Usage for Consent Freshness", RFC 7675, DOI 10.17487/RFC7675, October 2015, <https://www.rfc-editor.org/info/rfc7675>.

[RFC8441] McManus, P., "Bootstrapping WebSockets with HTTP/2", RFC 8441, DOI 10.17487/RFC8441, September 2018, <https://www.rfc-editor.org/info/rfc8441>.

Author's Address

Victor Vasiliev
Google

Email: vasilvv@google.com