

Workgroup: WebTransport

Internet-Draft:

draft-thomson-webtrans-session-limit-00

Published: 1 April 2023

Intended Status: Informational

Expires: 3 October 2023

Authors: M. Thomson E. Kinnear

 Mozilla Apple Inc.

Applying Per-Session Limits for WebTransport

Abstract

Limits to how a WebTransport session uses QUIC resources like streams or data can help limit the effect that one WebTransport session has on other uses of the same HTTP/3 connection. This describes mechanisms for limiting the number of streams and quantity of data that can be consumed by each WebTransport session.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at <https://martinthomson.github.io/wt-session-limits/draft-thomson-webtrans-session-limit.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-thomson-webtrans-session-limit/>.

Discussion of this document takes place on the WebTransport Working Group mailing list (<mailto:webtransport@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/webtransport/>. Subscribe at <https://www.ietf.org/mailman/listinfo/webtransport/>.

Source for this draft and an issue tracker can be found at <https://github.com/martinthomson/wt-session-limits>.

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 3 October 2023.

Copyright Notice

Copyright (c) 2023 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](#)
- [2. Conventions and Definitions](#)
- [3. Protocol Definition](#)
 - [3.1. Stream Limits](#)
 - [3.2. Data Limits](#)
- [4. Negotiation](#)
- [5. Security Considerations](#)
- [6. IANA Considerations](#)
- [7. Normative References](#)
- [Acknowledgments](#)
- [Authors' Addresses](#)

1. Introduction

WebTransport in HTTP/3 [[WTH3](#)] provides applications with all the functionality of QUIC [[QUIC](#)] streams. In the case where a single connection includes a WebTransport session that needs to coexist with other WebTransport sessions or HTTP requests, the core draft does not offer any way to place limits on stream usage.

This document describes an additional layer of session-level flow control that governs the creation of streams and sets a session-level limit on the amount of data that can be exchanged in a session.

This document does not define a framework for prioritizing the streams created for a WebTransport session with other streams.

Note that this document is intended as input for [\[WTH3\]](#). Although it is possible to define this as an extension to that protocol, integration of this design is simpler; see [Section 4](#).

2. Conventions and Definitions

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.

3. Protocol Definition

This document uses the following flow control capsules defined in [\[WTH2\]](#):

*WT_MAX_DATA ([Section 5.5](#) of [\[WTH2\]](#))

*WT_MAX_STREAMS ([Section 5.7](#) of [\[WTH2\]](#))

*WT_DATA_BLOCKED ([Section 5.8](#) of [\[WTH2\]](#))

*WT_STREAMS_BLOCKED ([Section 5.10](#) of [\[WTH2\]](#))

These capsules are unchanged, except that where the WebTransport over HTTP/2 capsules refer to streams that flow over the HTTP/2 stream containing the entire WebTransport session, these capsules refer to separate limits as described in subsequent sections.

These capsules use the codepoints allocated in [\[WTH2\]](#).

3.1. Stream Limits

The WT_MAX_STREAMS capsule establishes a limit on the number of streams within a WebTransport session. Like the QUIC MAX_STREAMS frame ([Section 19.11](#) of [\[QUIC\]](#)), this capsule has two types that provide separate limits for unidirectional and bidirectional streams that are initiated by a peer.

The session-level stream limit applies in addition to the QUIC MAX_STREAMS frame, which provides a connection-level stream limit. New streams can only be created within the session if both the stream- and the connection-level limit permit; see [Section 4.6](#) of [\[QUIC\]](#) for details on how QUIC stream limits are applied.

Unlike the WT_MAX_STREAMS capsule or the QUIC MAX_STREAMS frame, there is no simple relationship between the value in this frame and stream IDs in QUIC STREAM frames. This especially applies if there are other users of streams on the connection.

The `WT_STREAMS_BLOCKED` capsule is sent to indicate that an endpoint was unable to create a stream due to the session-level stream limit.

3.2. Data Limits

The `WT_MAX_DATA` capsule establishes a limit on the amount of data that can be sent within a WebTransport session. This limit counts all data that is sent on streams of the corresponding type, excluding the stream header (see Sections 4.2 and 4.2 of [WTH3]). The stream header is excluded from this limit so that this limit does not prevent the sending of information that is essential in linking new streams to a specific WebTransport session.

Implementing `WT_MAX_DATA` requires that the QUIC stack provide the WebTransport implementation with information about the final size of streams; see Section 4.5 of [QUIC].

The `WT_DATA_BLOCKED` capsule is sent to indicate that an endpoint was unable to send data due to a limit set by the `WT_MAX_DATA` capsule.

Because WebTransport over HTTP/3 uses a native QUIC stream for each WebTransport stream, per-stream data limits are provided by QUIC natively. The `WT_MAX_STREAM_DATA` and `WT_STREAM_DATA_BLOCKED` capsules are not used and so are prohibited. Endpoints **MUST** treat receipt of a `WT_MAX_STREAM_DATA` or a `WT_STREAM_DATA_BLOCKED` capsule as a session error.

4. Negotiation

If the use of flow control capsules are merged into the main specification [WTH3], their use will be negotiated along with the use of WebTransport over HTTP/3. This is the simplest approach.

Alternatively, if this remains as an optional extension, new HTTP/3 settings will be needed to negotiate the use of these features. In the abstract, we could define settings that carry initial values for the three variables that are controlled by the session-level flow control capsules defined here. The presence of any of those settings would indicate that these limits will be respected if the capsule is sent.

Both peers need to indicate the setting before these capsules apply. If only one peer advertises any of these settings, that might indicate that they are willing to receive and respect session-level flow control capsules. However, such an endpoint cannot know when to start applying the limit.

5. Security Considerations

Aside from new exposure to the usual programming errors arising from increased protocol complexity, it is believed that the introduction of these capabilities only improves security as it provides better control over endpoint resource allocation.

6. IANA Considerations

This document has no IANA actions.

7. Normative References

- [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>>.
- [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>>.
- [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>>.
- [WTH2] Frindell, A., Kinnear, E., Pauly, T., Thomson, M., Vasiliev, V., and G. Xie, "WebTransport over HTTP/2", Work in Progress, Internet-Draft, draft-ietf-webtrans-http2-05, 13 March 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-webtrans-http2-05>>.
- [WTH3] Frindell, A., Kinnear, E., and V. Vasiliev, "WebTransport over HTTP/3", Work in Progress, Internet-Draft, draft-ietf-webtrans-http3-05, 13 March 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-webtrans-http3-05>>.

Acknowledgments

TODO acknowledge.

Authors' Addresses

Martin Thomson
Mozilla

Email: mt@lowentropy.net

Eric Kinnear
Apple Inc.

Email: ekinnear@apple.com