

Workgroup: Internet Engineering Task Force
Internet-Draft: draft-kuhn-quic-0rtt-bdp-10
Published: 19 October 2021
Intended Status: Informational
Expires: 22 April 2022
Authors: N. Kuhn E. Stephan G. Fairhurst
 CNES Orange University of Aberdeen
 T. Jones C. Huitema
 University of Aberdeen Private Octopus Inc.
Transport parameters for 0-RTT connections

Abstract

QUIC 0-RTT transport features currently focuses on egress traffic optimization. This draft proposes a QUIC extension that improves the performance of ingress traffic.

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 22 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](#)
 - [1.1. Notations and terms](#)
 - [1.2. Requirements Language](#)
- [2. Safe jump start](#)
 - [2.1. Rationale behind the safety guidelines](#)
 - [2.2. Rationale #1: Variable network conditions](#)
 - [2.3. Rationale #2: Malicious clients](#)
 - [2.4. Trade-off between the different solutions](#)
 - [2.4.1. Security aspects](#)
 - [2.4.2. Interoperability and use-cases](#)
 - [2.4.3. Summary](#)
- [3. Safety guidelines](#)
- [4. Implementation considerations](#)
 - [4.1. Rationale behind the different implementation options](#)
 - [4.2. Independent local storage of values](#)
 - [4.3. Using NEW TOKEN frames](#)
 - [4.4. BDP Frame](#)
 - [4.4.1. BDP Frame Format](#)
 - [4.4.2. Extension activation](#)
- [5. Discussion](#)
 - [5.1. BDP extension protected as much as initial_max_data](#)
 - [5.2. Other use-cases](#)
 - [5.2.1. Optimizing client's requests](#)
 - [5.2.2. Sharing transport information across multiple connections](#)
- [6. Acknowledgments](#)
- [7. IANA Considerations](#)
- [8. Security Considerations](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

QUIC 0-RTT transport features currently focus on egress traffic optimization. This draft proposes a QUIC extension to improve the performance of ingress traffic.

[RFC9000] mentions that "Generally, implementations are advised to be cautious when using previous values on a new path." This draft proposes a discussion on how using previous values can be achieved in a interoperable manner and how it can be done safely.

When clients resume a session to download a large document, the congestion control algorithms will require time to ramp-up the packet rate. This document specifies a method that can improve

traffic delivery and that allows a QUIC connection to avoid a slow Round-Trip Time (RTT)-based process to grow connection parameters such as the congestion window (CWND):

1. During a previous session, current RTT (`current_rtt`), bottleneck bandwidth (`current_bb`) and client's current IP (`current_client_ip`) are stored as `saved_rtt`, `saved_bb` and `saved_client_ip`;
2. When resuming a session, the server might set the `current_rtt` and the `current_bb` to the `saved_rtt` and `saved_bb` of a previous connection.

This method applies to any QUIC resumed sessions: both `saved_session` and `recon_session` can be a 0-RTT QUIC connection or a 1-RTT QUIC connection.

This draft consider different solutions: (1) the saved parameters are not sent to the client; (2) the saved parameters are sent to the client and the client can not read them; (3) the saved parameters are sent to the client and the client can read them. There is no solution where the client can modify the parameters.

Sometimes the parameters of a previous session are not relevant, e.g.: (1) network conditions can change where using a previously estimated bottleneck bandwidth could increase congestion; (2) a client could convince a server to use a CWND much larger than required.

This draft:

1. proposes guidelines for how to safely apply the previously computed parameters to new sessions;
2. describes different implementation considerations in QUIC for the proposed method;
3. discusses the trade-off associated to the different implementation solutions.

1.1. Notations and terms

*IW: Initial window (e.g. from [[RFC6928](#)]);

*current_iw: Current Initial window;

*recom_iw: Recommended Initial window - it seems important to note that some Content Delivery Networks (CDNs) currently exploit a very high Initial Window (IW) [[TMA18](#)] for a local path;

*BDP: defined below;

*CWND: congestion window used by server (bytes allowed in flight by CC);

*current_bb : Current estimated bottleneck bandwidth;

*saved_bb: Estimated bottleneck bandwidth preserved from a previous connection;

*RTT: Round-Trip Time;

*current_rtt: Current RTT;

*saved_rtt: RTT preserved from a previous connection.

*client_ip : IP address of the client

*current_client_ip : Current IP address of the client

*saved_client_ip : IP address of the client preserved from a previous connection;

*remembered BDP parameters: combination of saved_rtt and saved_bb.

*ITT : Interpacket Transmission Time

*MSS : Maximum Message Size

*AEAD : Authenticated Encryption with Associated Data

*LRU : Least Recently Used

[[RFC6349](#)] defines the BDP as follows: "Derived from Round-Trip Time (RTT) and network Bottleneck Bandwidth (BB), the Bandwidth-Delay Product (BDP) determines the Send and Received Socket buffer sizes required to achieve the maximum TCP Throughput." This draft considers the Bandwidth-Delay Product (BDP) as estimated by the server which includes all buffering along the network path. A QUIC connection might not exactly reproduce the procedure detailed in [[RFC6349](#)] to measure the BDP. The server can exploit internal evaluations of the Bottleneck Bandwidth and the to assess the BDP.

This document refers to the saved_bb and current_bb for the previously estimated bottleneck bandwidth. This value may be easily reached for protocols such as BBR. Other protocols, such as CUBIC or RENO, may estimate this value by exploiting the congestion window and the RTT. This approach may result in over estimating the bottleneck bandwidth and should be considered with caution.

1.2. Requirements Language

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.

2. Safe jump start

2.1. Rationale behind the safety guidelines

The previously measured `saved_rtt` and `saved_bb` should not be used as-is to avoid potential congestion collapse:

- *Rationale #1: An Internet method needs to be robust to network conditions that can differ between sessions.

- *Rationale #2: Information sent by a malicious client would not be relevant since it might try to convince servers to use a `CWND` higher than required. This could increase congestion.

2.2. Rationale #1: Variable network conditions

The server **MUST** check the validity of the `saved_rtt` and `saved_bb` parameters, whether they are sent by a client or stored at the server. Indeed, the following events make use of these parameters inappropriate:

- *IP address changes: If the client changes its IP address (i.e. `saved_client_ip` is different from `current_client_ip`), the different address indicates a different network path. This new path does not necessarily exhibit the same characteristics as the old one. If the server changes its IP address after a migration, it may not be safe to exploit previously estimated parameters.

- *Lifetime of the extension: If the network conditions change, e.g., the path was not congested when BDP parameters were evaluated, but later the path experiences congestion for the next connection, the previously estimated parameters would not be valid.

- *Relevance of the BB estimation: There are cases where the congestion window is exploited to assess the bottleneck bandwidth. However, at the end of slow start its value may be significantly larger than the value on which the connection converges after a few rounds. Directly exploiting such value for the bottleneck bandwidth estimation may be inappropriate.

There are different solutions for the variable network conditions:

*Rationale #1 - Solution #1 : When resuming a session, set the current_bb and current_rtt to the saved_bb and saved_rtt parameters estimated from a previous connection.

*Rationale #1 - Solution #2 : When resuming a session, implement a safety check to measure whether using the saved_bb and saved_rtt parameters would not cause congestion over the path. In this case, the current_bb and current_rtt might not be set directly to the saved_bb and saved_rtt: the server might wait for the completion of the safety check before doing so.

[Section 3](#) describes various approaches for Rationale #1 - Solution #2.

2.3. Rationale #2: Malicious clients

The server MUST check the integrity of the saved_rtt and saved_bb parameters received from a client.

There are different solutions to avoid attacks by malicious clients:

*Rationale #2 - Solution #1 : The server stores a local estimation of the bottleneck bandwidth and RTT parameters as saved_bb and saved_rtt.

*Rationale #2 - Solution #2 : The server sends the estimation of the bottleneck bandwidth and RTT parameters to the client as saved_bb and saved_rtt. The information is encrypted by the server. The client resends the information when resuming a connection. The client can neither read nor modify the saved_rtt and saved_bb parameters.

*Rationale #2 - Solution #3 : The server sends the estimation of the saved_rtt and saved_bb parameters to the client. The information includes integrity protection. The client resends the information when resuming a connection. The client can read, but can not modify, the saved_rtt and saved_bb parameters.

[Section 4](#) describes various implementation approaches for each of these solutions using local storage ([Section 4.2](#) for Rationale #2 - Solution #1), NEW_TOKEN Frame ([Section 4.3](#) for Rationale #2 - Solution #2), BDP extension Frame ([Section 4.4](#) for Rationale #2 - Solution #3).

2.4. Trade-off between the different solutions

This section provides a description of different implementation options and discusses their respective advantages and drawbacks.

While there are some discussions for the solutions regarding Rationale #2, the server MUST consider Rationale #1 - Solution #2 and avoid Rationale #1 - Solution #1: the server MUST implement a safety check to measure whether the remembered BDP parameters (i.e. saved_rtt and saved_bb) are relevant or check that their usage would not cause congestion over the path.

2.4.1. Security aspects

The client may send information related to the saved_rtt and saved_bb to the server with the BDP Frame extension using either Rationale #2 - Solution #2 or Rationale #2 - Solution #3. However, the server may not trust the client. Indeed, even if 0-RTT packets containing the BDP Frame are encrypted, a client could modify the values within the extension and encrypt the 0-RTT packet. Authentication mechanisms might not guarantee that the values are safe. The server could then need to also store the saved_rtt and saved_bb parameters.

A malicious client might modify the saved_bb parameter to convince the server to use a CWND much larger than required. Using the algorithms proposed in [Section 3](#), the server may reduce any intended harm and can check that part of the information provided by the client are valid. A supplementary check could decide not to use values that would be higher than those currently used by CDNs [[TMA18](#)].

Storing the BDP parameters locally at the server reduces the associated risks by allowing the client to transmit information related to the BDP of the path.

2.4.2. Interoperability and use-cases

If the server stores a resumption ticket for each client to protect against replay on a third party IP, it could also store the IP address (i.e. saved_client_ip) and BDP parameters (i.e. saved_rtt and saved_bb) of the previous session of the client.

In cases where the BDP Frame extension is exploited, the approach of storing the BDP parameters locally at the server can provide a cross-check of the BDP parameters sent by a client. The server can anyway enable a safe jump start, but without the BDP Frame extension, the client does not have the choice of accepting it or not.

While storing local values related to the BDP would help in improving the ingress for 0-RTT connections, not using a BDP Frame extension may reduce the interest of the approach where (1) the client knows the BDP estimations done at the server, (2) the client

decides to accept or reject ingress optimization, (3) the client tunes application level requests.

2.4.3. Summary

As a summary, the approach of local storage of values is more secure and the BDP Frame extension provides more information to the client and more interoperability. The [Figure 1](#) provides a summary of the advantages and drawbacks of each approach.

Rationale	Solution	Advantage	Drawback	Comment
#1 Variable Network	#1 set current_* to saved_*	Ingress optim.	Risks of adding congestion	MUST NOT implement
	#2 Implement safety check	Reduce risks of adding congestion	Negative impact on ingress optim.	MUST implement Section 3
#2 Malicious client	#1 Local storage	Enforced security	Client can not decide to reject Malicious server could fill client's buffer Limited use-cases	Section 4.2
	#2 NEW_TOKEN	Save resource at server Opaque token protected	Malicious client may change token even if protected Malicious server could fill client's buffer Server may not trust client	Section 4.3
	#3 BDP extension	Extended use-cases Save resource at server Client can read and decide to reject BDP extension protected	Malicious client may change BDP even if protected Server may not trust client	Section 4.4

Figure 1: Comparing solutions

3. Safety guidelines

The safety guidelines are designed to avoid a server adding excessive congestion to an already congested path. The following mechanisms should help in fulfilling this objective:

- *The server SHOULD compare the measured transport parameters (in particular `current_rtt`) of the 0-RTT connection with those of the 1-RTT connection (in particular `saved_rtt`);
- *The server SHOULD NOT consider the `saved_bb` parameter if there is any loss of packet during the first transmission of data;
- *The server MUST NOT send more than a recommended maximum IW (`recom_iw`) in the first transmission of data. This value could be based on a local understanding of the path characteristics and what is deployed in CDNs [TMA18].
- *The server SHOULD NOT store and/or send information related to the previously estimated bottleneck bandwidth (`saved_bb`) if this estimation has not been computed after some rounds.

The proposed mechanisms SHOULD be limited by any rate-limitation mechanisms of QUIC, such as flow control mechanisms or amplification attacks prevention. In particular, it may be necessary to issue proactive `MAX_DATA` frames to increase the flow control limits of the connection. In particular, the maximum number of packets that can be sent without acknowledgment needs to be chosen to avoid the creation and the increase of congestion for the path. Moreover, this extension should not be an opportunity for the current connection to be a vector of an amplification attack. The address validation process, used to prevent amplification attacks, SHOULD be performed [RFC9000].

The following mechanisms could be implemented:

- *Exploit a standard IW:
 1. The server sends the first data packet using the IW - this can be considered a safe starting point for an unknown path, which avoids adding congestion to the path;
 2. If the reception of IW exhibits characteristics that resemble those of a recent previous session from the client (i.e. $\text{current_rtt} < 1.2 * \text{saved_rtt}$ and all the data was acknowledged), the method permits the sender to consider the `saved_bb` as an input to adapt `current_bb` and help rapidly determine a new safe rate;

3. The sender needs to avoid a burst of packets being sent as a result of a step-increase in the congestion window [[RFC9000](#)]. Pacing the packets as a function of the current_rtt can provide this additional safety during the period in which the CWND is increased by the method.

*Identify a relevant pacing rhythm:

- The server estimates the pacing rhythm using saved_rtt and saved_bb. The Interpacket Transmission Time (ITT) is determined by the ratio between the current Maximum Message Size (MSS) for packets and the ratio between the saved_bb and saved_rtt. A tunable safety margin might be introduced to avoid sending more than a recommended maximum IW (recom_iw):

ocurrent_iw = min(recom_iw, saved_bb)

oITT = MSS/(current_iw/saved_rtt)

- When the IW is acknowledged, the server falls back to a standard slow-start mechanism.

*Tune slow-start mechanisms. After transport parameters are set to previously estimated bottleneck bandwidth, if slow-start mechanisms continue, important packets overshoot may be transmitted from the server. This can occur even if the safety check described in this section is implemented.

- For NewReno and CUBIC, it is recommended to exit slow-start and enter in congestion avoidance phase.

- For BBR, it is recommended to move to the "probe bandwidth" state.

This follows the idea of [[RFC4782](#)], [[I-D.irtf-iccrq-sallantin-initial-spreading](#)] and [[CONEXT15](#)].

While safety recommendations are necessary, it seems important to note that some Content Delivery Networks (CDNs) currently exploit a very high Initial Window (IW) [[TMA18](#)] for a local path.

4. Implementation considerations

4.1. Rationale behind the different implementation options

Using NewSessionTickets messages of TLS is a solution that could have been envisioned. The idea would have been to add a 'bdp_metada' field in the NewSessionTickets that the client could read. The sole extension currently defined in TLS1.3 that can be seen by the client is max_early_data_size (see section 4.6.1 of [[RFC8446](#)]). However, in

the general design of QUIC, TLS sessions are managed by the TLS stacks.

Three distinct approaches are presented: sending an opaque blob to the client that it may retransmit for future connection (see [Section 4.3](#)), enable a local storage of BDP related values (see [Section 4.2](#)) and a BDP Frame extension (see [Section 4.4](#)).

4.2. Independent local storage of values

This approach independently lets both a client and a server remember their BDP parameters:

- *During a 1-RTT session, the endpoint stores the RTT (as the `saved_rtt`) and bottleneck bandwidth (as the `saved_bb`) together with the session resume ticket. The client can also store the IP address of the server.

- *The server maintains a table of previously issued tickets, indexed by the random ticket identifier that is used to guarantee uniqueness of the Authenticated Encryption with Associated Data (AEAD) encryption. Old tokens are removed from the table using the Least Recently Used (LRU) logic. For each ticket identifier, the table holds the RTT and bottleneck bandwidth (i.e. `saved_rtt` and `saved_bb`), and also the IP address of the client (i.e. `saved_client_ip`).

During the 0-RTT session, the endpoint wait for the first RTT measurement from the peer's IP address. This is used to verify that the `current_rtt` has not significantly changed from the `saved_rtt`, and hence is an indication that the BDP information applies to the path that is currently being used.

If this RTT is confirmed (e.g. $\text{current_rtt} < 1.2 * \text{saved_rtt}$), the endpoint also verifies that an initial window of data has been acknowledged without requiring retransmission. This second check is used to detect a path with significant incipient congestion (i.e. where it would not be safe to update the CWND based on the `saved_bb`). In practice, this could be realized by a proportional increase in the CWND, where the increase is $(\text{saved_bb} / \text{IW}) * \text{proportion_of_IW_currently_ACKed}$.

This solution does not allow the client to refuse the exploitation of the BDP parameters. If the server does not want to store the metrics from previous connections, an equivalent of the `tcp_no_metrics_save` for QUIC may be necessary. This option could be negotiable so that a client can refuse the exploitation of previous sessions.

4.3. Using NEW_TOKEN frames

Using NEW_TOKEN Frames, the server could send a token to the client through a NEW_TOKEN Frame. The token is an opaque blob and the client can not read its content (see section 19.7 of [\[RFC9000\]](#)). The client sends the received token in the header of an Initial packet for future connection.

4.4. BDP Frame

This section proposes the exploitation of a new Frame, the BDP Frame. The BDP Frame MUST be contained in 0-RTT packets if sent by the client. The BDP Frame MUST be contained in 1-RTT packets if sent by the server. The BDP Frame MUST be considered in the congestion control and its data may not be limited by flow control limits. The server MAY send multiple BDP Frames in both 1-RTT and 0-RTT connections. The client may send BDP Frames during 1-RTT and 0-RTT connections.

4.4.1. BDP Frame Format

A BDP Frame is formatted as shown in [Figure 2](#).

```
BDP Frame {  
  Type (i) = 0xFFFF,  
  Lifetime (i),  
  Saved BB (i),  
  Saved RTT (i),  
  Saved IP length (i),  
  Saved IP (...)  
}
```

Figure 2: BDP Frame Format

A BDP Frame contains the following fields:

- *Lifetime (extension_lifetime): The extension_lifetime is a value in milliseconds, encoded as a variable length integer. This follows the idea of NewSessionTicket of TLS [\[RFC8446\]](#). This represents the validity in time of this extension.
- *Saved BB (saved_bb): The saved_bb is a value in bytes, encoded as a variable length integer. The bottleneck bandwidth estimated on the previous connection by the server. Using the previous values of bytes_in_flight defined in [\[RFC9002\]](#) can result in overshoot and is not advised.
- *Saved RTT (saved_rtt): The saved_rtt is a value in milliseconds, encoded as a variable length integer. This could be set to the

min_rtt defined in [\[RFC9002\]](#), saved_rtt can be set to min_rtt. The min_rtt parameter might not track a decreasing RTT: the min_rtt that is reported here might not be the actual minimum RTT measured during the 1-RTT connection, but usually reflects the characteristics of the path latency.

*Saved IP length (saved_ip_lenght) : The length of the IP address set to either 4 (IPv4) or 16 (IPv6).

*Saved IP (saved_client_ip) : The saved_client_ip could be set to the IP address of the client.

4.4.2. Extension activation

The client can accept the transmission of BDP Frames from the server by using the following enable_bdp transport extension.

enable_bdp (0XTBD): in the 1-RTT connection, the client indicates to the server that it wishes to receive BDP extension Frames for improving ingress of 0-RTT connection. The default value is 0. Values strictly above 3 are invalid, and receipt of these values MUST be treated as a connection error of type TRANSPORT_PARAMETER_ERROR.

*0: Default value. If the client does not send this parameter, the server considers that the client does not support or does not wish to activate the BDP extension.

*1: The client indicates to the server that it wishes to receive BDP Frame and activates the ingress optimization for the 0-RTT connection.

*2: The client indicates that it does not wish to receive BDP Frames but activates ingress optimization.

*3: The client indicates that it wishes to receive BDP Frames but does not activate ingress optimization.

This Transport Parameter is encoded as per Section 18 of [\[RFC9000\]](#).

5. Discussion

5.1. BDP extension protected as much as initial_max_data

The BDP metadata parameters are measured by the server during a previous connection. The BDP extension is protected by the mechanism that protects the exchange of the 0-RTT transport parameters. For version 1 of QUIC, the BDP extension is protected using the mechanism that already protects the "initial_max_data" parameter. This is defined in sections 4.5 to 4.7 of [\[RFC9001\]](#). This provides

the server with a way to verify that the parameters proposed by the client are the same as those that the server sent to the client during the previous connection.

5.2. Other use-cases

5.2.1. Optimizing client's requests

In a case with Dynamic Adaptive Streaming over HTTPS (DASH), clients might encounter issues in knowing the available path capacity or DASH can encounter issues in reaching the best available video playback quality. The client requests could then be adapted and specific traffic could utilize information from the path characteristics (such as encouraging the client to increase the quality of video chunks, to fill the buffers and avoid video blocking or to send high quality adds).

In other cases, applications may provide additional services if clients can know the server's estimation of the path characteristics.

5.2.2. Sharing transport information across multiple connections

There can be benefit in sharing transport information across multiple connections. [[I-D.ietf-tcpm-2140bis](#)] considers the sharing of transport parameters between TCP connections originating from the same host. The proposal in this document has the advantage of storing server-generated information at the client and not requiring the server to retain additional state for each client.

6. Acknowledgments

The authors would like to thank Gabriel Montenegro, Patrick McManus, Ian Swett, Igor Lubashev, Robin Marx, Roland Bless and Franklin Simo for their fruitful comments on earlier versions of this document.

7. IANA Considerations

TBD: Text is required to register the BDP Frame and the enable_bdp transport parameter. Parameters are registered using the procedure defined in [[RFC9000](#)].

8. Security Considerations

Security considerations are discussed in [Section 5](#) and in [Section 3](#).

9. References

9.1. Normative References

[I-D.ietf-tcpm-2140bis]

Touch, J., Welzl, M., and S. Islam, "TCP Control Block Interdependence", Work in Progress, Internet-Draft, draft-ietf-tcpm-2140bis-11, 12 April 2021, <<https://www.ietf.org/archive/id/draft-ietf-tcpm-2140bis-11.txt>>.

[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>>.

[RFC4782] Floyd, S., Allman, M., Jain, A., and P. Sarolahti, "Quick-Start for TCP and IP", RFC 4782, DOI 10.17487/RFC4782, January 2007, <<https://www.rfc-editor.org/info/rfc4782>>.

[RFC6349] Constantine, B., Forget, G., Geib, R., and R. Schrage, "Framework for TCP Throughput Testing", RFC 6349, DOI 10.17487/RFC6349, August 2011, <<https://www.rfc-editor.org/info/rfc6349>>.

[RFC6928] Chu, J., Dukkkipati, N., Cheng, Y., and M. Mathis, "Increasing TCP's Initial Window", RFC 6928, DOI 10.17487/RFC6928, April 2013, <<https://www.rfc-editor.org/info/rfc6928>>.

[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>>.

[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/info/rfc9000>>.

[RFC9001] Thomson, M., Ed. and S. Turner, Ed., "Using TLS to Secure QUIC", RFC 9001, DOI 10.17487/RFC9001, May 2021, <<https://www.rfc-editor.org/info/rfc9001>>.

[RFC9002] Iyengar, J., Ed. and I. Swett, Ed., "QUIC Loss Detection and Congestion Control", RFC 9002, DOI 10.17487/RFC9002, May 2021, <<https://www.rfc-editor.org/info/rfc9002>>.

9.2. Informative References

[CONEXT15]

Li, Q., Dong, M., and P B. Godfrey, "Halfback: Running Short Flows Quickly and Safely", ACM CoNEXT , 2015.

[I-D.irtf-iccr-g-sallantin-initial-spreading]

Sallantin, R., Baudoin, C., Arnal, F., Dubois, E., Chaput, E., and A. Beylot, "Safe increase of the TCP's Initial Window Using Initial Spreading", Work in Progress, Internet-Draft, draft-irtf-iccr-g-sallantin-initial-spreading-00, 15 January 2014, <<https://www.ietf.org/archive/id/draft-irtf-iccr-g-sallantin-initial-spreading-00.txt>>.

[TMA18]

Ruth, J. and O. Hohlfeld, "Demystifying TCP Initial Window Configurations of Content Distribution Networks", 2018 Network Traffic Measurement and Analysis Conference (TMA) , 2018.

Authors' Addresses

Nicolas Kuhn
CNES

Email: nicolas.kuhn.ietf@gmail.com

Emile Stephan
Orange

Email: emile.stephan@orange.com

Godred Fairhurst
University of Aberdeen
Department of Engineering
Fraser Noble Building
Aberdeen

Email: gorry@erg.abdn.ac.uk

Tom Jones
University of Aberdeen
Department of Engineering
Fraser Noble Building
Aberdeen

Email: tom@erg.abdn.ac.uk

Christian Huitema
Private Octopus Inc.

Email: huitema@huitema.net