

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: 24 November 2025

A. Custura
G. Fairhurst
University of Aberdeen
23 May 2025

Quic Logging for Convergence of Congestion Control from Retained State
draft-ietf-tsvwg-careful-resume-qlog-01

Abstract

This document specifies a logging format for a cautious method for Careful Resume when using the IETF quic transport protocol. It defines the logging format for qlog.

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 24 November 2025.

Copyright Notice

Copyright (c) 2025 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
2. Language, Notation and Terms	2
2.1. Requirements Language	3
2.2. Notation and Terms	3
3. The Phases of CC using Careful Resume	3
3.1. Observing	4
3.2. Mapping of triggers to phase transitions	4
3.3. QLOG support for QUIC	6
3.3.1. The CarefulResumePhaseUpdated Event	6
4. Acknowledgments	8
5. IANA Considerations	8
6. Security Considerations	8
7. References	8
7.1. Normative References	8
7.2. Informative References	9
Appendix A. Internet Draft Revision details	9
Authors' Addresses	9

1. Introduction

This document defines a way to provide logging for Careful Resume[I-D.ietf-tsvwg-careful-resume] when used with QUIC [RFC9000][RFC9002].

[I-D.ietf-tsvwg-careful-resume] defines a Congestion Control (CC) mechanism called Careful Resume, which is expected to reduce the time to complete a transfer when the transfer sends significantly more data than allowed by the Initial congestion Window (IW), and where the BDP of the path is also significantly more than the IW. It introduces an alternative mechanism to select initial CC parameters that seeks to more rapidly and safely grow the sending rate controlled by the congestion window (CWND). CC algorithms that are rate-based can make similar adjustments to their target sending rate.

When used with the QUIC transport, Careful Resume provides transport services that resemble those that could be implemented in TCP, using methods such as TCP Control Block (TCB) [RFC9040] caching.

2. Language, Notation and Terms

This subsection provides a brief summary of key terms and the requirements language.

2.1. Requirements Language

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.

2.2. Notation and Terms

The document uses language drawn from a range of IETF RFCs. The following terms are defined:

Careful Resume (CR): The method specified in [I-D.ietf-tsvwg-careful-resume]. All other terms in this section, except CWND, are defined in the same document.

CWND: The congestion window, or equivalent CC variable limiting the maximum sending rate;

current_rtt: A sample measurement of the current RTT;

PipeSize: A measure of the validated available capacity based on the acknowledged data;

saved_cwnd: The preserved capacity derived from observation of a previous connection;

saved_rtt: The preserved minimum RTT;

Unvalidated Packet: A packet sent when the CWND has been increased beyond the size normally permitted by the CC algorithm; if such a packet is acknowledged, it contributes to the PipeSize, if congestion is detected, it triggers entry to the Safe Retreat Phase.

3. The Phases of CC using Careful Resume

Careful Resume[I-D.ietf-tsvwg-careful-resume] defines a series of phases that a congestion controller moves through as a connection uses the mechanism. The key phases are illustrated in Figure 1.

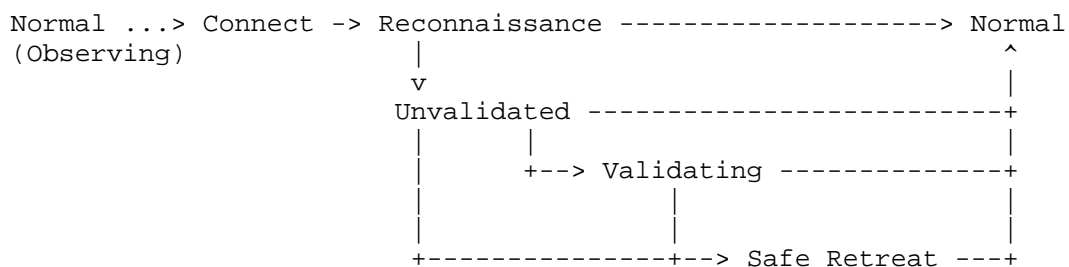


Figure 1: Key transitions between Phases in Careful Resume.

Examples of these transitions between phases are provided in the annexe to [I-D.ietf-tsvwg-careful-resume].

3.1. Observing

There are currently no qlog events associated with this phase.

3.2. Mapping of triggers to phase transitions

Careful Resume phases are abbreviated as follows:

- * *Unval* = Unvalidated Phase
- * *Val* = Validating Phase
- * *Recon* = Reconnaissance Phase
- * *SR* = Safe Retreat Phase
- * *Norm* = Normal Phase

Trigger	Description	From	To
cwnd_limited	When sender has confirmed the RTT, has received an ACK for the initial data without reported congestion and has more data to send than the CWND would allow.	Recon	Unval
rtt_not_validated	If the current_rtt is not confirmed the sender MUST enter the Normal Phase.	Recon Unval	Norm

last_unvalidated_packet_sent	Completed sending all unvalidated packets, e.g., when flight_size is equal to the CWND after the jump	Unval	Val
first_unvalidated_packet_acknowledged	The sender enters the Validating Phase when an ACK is received for the first packet number (or higher) sent in the Unvalidated Phase.	Unval	Val
rtt_exceeded	When greater than 1 RTT has passed in Unvalidated Phase.	Unval	Val
rate_limited	If the flight_size is less than or equal to the PipeSize sender enters Normal Phase.	Unval	Norm
last_unvalidated_packet_acknowledged	The sender enters the Normal Phase when an ACK is received for the last packet number (or higher) that was sent in the Unvalidated Phase.	Val	Norm
path_changed	If a sender determines that it is not valid to use the previous CC parameters due to a detected path change (e.g., a change in RTT or an explicit signal indicating a path change.	Unval	SR
packet_loss	If a sender determines that congestion was experienced, e.g., packet loss, the sender enters the Safe Retreat Phase.	Unval Val	SR
	If the sender determines congestion was experienced in the Recon Phase, the sender enters the Normal Phase.	Recon	Norm
ECN_CE	If a sender determines that congestion was experienced, e.g., ECN-CE marking, sender enters the Safe Retreat Phase.	Unval Val	SR

	If the sender determines congestion was experienced in the Recon Phase, the sender enters the Normal Phase.	Recon	Norm
exit_recovery	The sender enters the Normal Phase when the last packet sent in the Unvalidated Phase is ACKed.	SR	Norm

Figure 2: Mapping of triggers to phase transitions

3.3. QLOG support for QUIC

This section provides definitions that enable a Careful Resume implementation to generate qlog events when using QUIC. It introduces an event to report a phase transition at the sender, and its associated description. The event's state parameters (which include mandatory CR-specific metrics and optionally CWND or ssthresh) are intended to reflect the state after the phase transition.

The event and data structure definitions in this section are expressed in the Concise Data Definition Language (CDDL) [RFC8610] and its extensions described in [I-D.ietf-quic-qlog-quic-events]. The current convention is to use long names for variables. For example, "CWND" is expanded as "congestion_window" and "saved_cwnd" is expanded as "saved_congestion_window".

3.3.1. The CarefulResumePhaseUpdated Event

Importance: Extra

When the CC algorithm changes the Careful Resume Phase described in Section 3 of this specification.

Definition:

```
QUICCarefulResumePhaseUpdated = {
  ? old: CarefulResumePhase,
  new: CarefulResumePhase,
  state_data: CarefulResumeStateParameters,
  ? restored_data: CarefulResumeRestoredParameters,
  ? trigger:
```

```

; for the Unvalidated Phase,
; when there is more data to send than available CWND.
"congestion_window_limited" /
; for the Validating Phase
; when more than one RTT has passed in the Unvalidated Phase.
"rtt_exceeded" /
; for the Validating Phase
"first_unvalidated_packet_acknowledged" /
; for the Validating Phase
; when all unvalidated packets have been sent.
"last_unvalidated_packet_sent" /
; for the Normal Phase
; when no remaining unvalidated packets to be acknowledged.
"last_unvalidated_packet_acknowledged" /
; for the Normal Phase
; when CR not allowed due to a detected RTT change.
"rtt_not_validated" /
; for the Normal Phase,
; when sending fewer unvalidated packets than CWND permits.
"rate_limited" /
; for the Safe Retreat Phase, when path change detected.
"path_changed" /
; for the Safe Retreat Phase, when loss detected.
; for the Normal Phase,
; when loss detected in the Reconnaissance Phase.
"packet_loss" /
; for the Safe Retreat Phase,
; when ECN congestion experienced reported.
; for the Normal Phase,
; when ECN congestion experienced in the Reconnaissance Phase.
"ECN_CE" /
; for the Normal Phase 1 RTT after a congestion event
"exit_recovery"

* $quic-carefulresumephaseupdated-extension
}

CarefulResumePhase =
    "reconnaissance" /
    "unvalidated" /
    "validating" /
    "normal" /
    "safe_retreat"

CarefulResumeStateParameters = {
    pipesize: uint32,
    first_unvalidated_packet: uint32,
    last_unvalidated_packet: uint32,

```

```
? congestion_window: uint32,  
? ssthresh: uint32  
}  
  
CarefulResumeRestoredParameters = {  
    saved_congestion_window: uint32,  
    saved_rtt: float32  
}  
  
$ProtocolEventData /= QUICCarefulResumePhaseUpdated
```

Figure 3: CDDL description of the CarefulResumePhaseUpdated event

4. Acknowledgments

The authors would like to thank Robin Marx and Lucas Pardue for their fruitful comments in developing this specification.

5. IANA Considerations

Note by RM: This event type is to be added to the "qlog event schema URIs" registered at IANA, but the qlog event schema URIs registry has not yet been actually created at IANA because the qlog registry has not yet been created.

6. Security Considerations

This document does not exhibit specific security considerations.

7. References

7.1. Normative References

- [I-D.ietf-quic-qlog-quic-events]
Marx, R., Niccolini, L., Seemann, M., and L. Pardue, "QUIC event definitions for qlog", Work in Progress, Internet-Draft, draft-ietf-quic-qlog-quic-events-10, 17 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-quic-qlog-quic-events-10>>.
- [I-D.ietf-tsvwg-careful-resume]
Kuhn, N., Stephan, E., Fairhurst, G., Secchi, R., and C. Huitema, "Convergence of Congestion Control from Retained State", Work in Progress, Internet-Draft, draft-ietf-tsvwg-careful-resume-18, 23 May 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-careful-resume-18>>.

- [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>>.
- [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>>.
- [RFC8610] Birkholz, H., Vigano, C., and C. Bormann, "Concise Data Definition Language (CDDL): A Notational Convention to Express Concise Binary Object Representation (CBOR) and JSON Data Structures", RFC 8610, DOI 10.17487/RFC8610, June 2019, <<https://www.rfc-editor.org/info/rfc8610>>.
- [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>>.

7.2. Informative References

- [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>>.
- [RFC9040] Touch, J., Welzl, M., and S. Islam, "TCP Control Block Interdependence", RFC 9040, DOI 10.17487/RFC9040, July 2021, <<https://www.rfc-editor.org/info/rfc9040>>.

Appendix A. Internet Draft Revision details

WG -00 Addresses comments from Mike Heard and Joerg Deutschmann.

WG -01 Adds a new trigger, `path_changed`, to align with the main draft, fixes formatting NiTs.

Authors' Addresses

Ana Custura
University of Aberdeen
Department of Engineering
Fraser Noble Building
Aberdeen
AB24 3UE
United Kingdom
Email: ana@erg.abdn.ac.uk

Godred Fairhurst
University of Aberdeen
Department of Engineering
Fraser Noble Building
Aberdeen
AB24 3UE
United Kingdom
Email: gorry@erg.abdn.ac.uk