

QUIC  
Internet-Draft  
Intended status: Standards Track  
Expires: June 21, 2019

J. Iyengar, Ed.  
Fastly  
I. Swett, Ed.  
Google  
December 18, 2018

**QUIC Loss Detection and Congestion Control**  
**draft-ietf-quic-recovery-17**

Abstract

This document describes loss detection and congestion control mechanisms for QUIC.

Note to Readers

Discussion of this draft takes place on the QUIC working group mailing list ([quic@ietf.org](mailto:quic@ietf.org)), which is archived at [https://mailarchive.ietf.org/arch/search/?email\\_list=quic](https://mailarchive.ietf.org/arch/search/?email_list=quic) [1].

Working Group information can be found at <https://github.com/quicwg> [2]; source code and issues list for this draft can be found at <https://github.com/quicwg/base-drafts/labels/-recovery> [3].

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 June 21, 2019.

Copyright Notice

Copyright (c) 2018 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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">4</a>
<a href="#">2.</a>	Conventions and Definitions . . . . .	<a href="#">4</a>
<a href="#">3.</a>	Design of the QUIC Transmission Machinery . . . . .	<a href="#">4</a>
<a href="#">3.1.</a>	Relevant Differences Between QUIC and TCP . . . . .	<a href="#">5</a>
<a href="#">3.1.1.</a>	Separate Packet Number Spaces . . . . .	<a href="#">5</a>
<a href="#">3.1.2.</a>	Monotonically Increasing Packet Numbers . . . . .	<a href="#">6</a>
<a href="#">3.1.3.</a>	No Reneging . . . . .	<a href="#">6</a>
<a href="#">3.1.4.</a>	More ACK Ranges . . . . .	<a href="#">6</a>
<a href="#">3.1.5.</a>	Explicit Correction For Delayed ACKs . . . . .	<a href="#">6</a>
<a href="#">4.</a>	Generating Acknowledgements . . . . .	<a href="#">7</a>
<a href="#">4.1.</a>	Crypto Handshake Data . . . . .	<a href="#">7</a>
<a href="#">4.2.</a>	ACK Ranges . . . . .	<a href="#">7</a>
<a href="#">4.3.</a>	Receiver Tracking of ACK Frames . . . . .	<a href="#">8</a>
<a href="#">5.</a>	Computing the RTT estimate . . . . .	<a href="#">8</a>
<a href="#">6.</a>	Loss Detection . . . . .	<a href="#">9</a>
<a href="#">6.1.</a>	Acknowledgement-based Detection . . . . .	<a href="#">9</a>
<a href="#">6.1.1.</a>	Packet Threshold . . . . .	<a href="#">9</a>
<a href="#">6.1.2.</a>	Time Threshold . . . . .	<a href="#">10</a>
<a href="#">6.2.</a>	Timeout Loss Detection . . . . .	<a href="#">10</a>
<a href="#">6.2.1.</a>	Crypto Retransmission Timeout . . . . .	<a href="#">10</a>
<a href="#">6.2.2.</a>	Probe Timeout . . . . .	<a href="#">12</a>
<a href="#">6.3.</a>	Tracking Sent Packets . . . . .	<a href="#">13</a>
<a href="#">6.3.1.</a>	Sent Packet Fields . . . . .	<a href="#">14</a>
<a href="#">6.4.</a>	Pseudocode . . . . .	<a href="#">14</a>
<a href="#">6.4.1.</a>	Constants of interest . . . . .	<a href="#">14</a>
<a href="#">6.4.2.</a>	Variables of interest . . . . .	<a href="#">15</a>
<a href="#">6.4.3.</a>	Initialization . . . . .	<a href="#">16</a>
<a href="#">6.4.4.</a>	On Sending a Packet . . . . .	<a href="#">16</a>
<a href="#">6.4.5.</a>	On Receiving an Acknowledgment . . . . .	<a href="#">16</a>
<a href="#">6.4.6.</a>	On Packet Acknowledgment . . . . .	<a href="#">18</a>
<a href="#">6.4.7.</a>	Setting the Loss Detection Timer . . . . .	<a href="#">18</a>
<a href="#">6.4.8.</a>	On Timeout . . . . .	<a href="#">19</a>
<a href="#">6.4.9.</a>	Detecting Lost Packets . . . . .	<a href="#">20</a>
<a href="#">6.5.</a>	Discussion . . . . .	<a href="#">21</a>
<a href="#">7.</a>	Congestion Control . . . . .	<a href="#">22</a>
<a href="#">7.1.</a>	Explicit Congestion Notification . . . . .	<a href="#">22</a>



<a href="#">7.2.</a>	Slow Start . . . . .	<a href="#">22</a>
<a href="#">7.3.</a>	Congestion Avoidance . . . . .	<a href="#">22</a>
<a href="#">7.4.</a>	Recovery Period . . . . .	<a href="#">23</a>
<a href="#">7.5.</a>	Probe Timeout . . . . .	<a href="#">23</a>
<a href="#">7.6.</a>	Pacing . . . . .	<a href="#">23</a>
<a href="#">7.7.</a>	Sending data after an idle period . . . . .	<a href="#">24</a>
<a href="#">7.8.</a>	Discarding Packet Number Space State . . . . .	<a href="#">24</a>
<a href="#">7.9.</a>	Pseudocode . . . . .	<a href="#">24</a>
<a href="#">7.9.1.</a>	Constants of interest . . . . .	<a href="#">24</a>
<a href="#">7.9.2.</a>	Variables of interest . . . . .	<a href="#">25</a>
<a href="#">7.9.3.</a>	Initialization . . . . .	<a href="#">26</a>
<a href="#">7.9.4.</a>	On Packet Sent . . . . .	<a href="#">26</a>
<a href="#">7.9.5.</a>	On Packet Acknowledgement . . . . .	<a href="#">26</a>
<a href="#">7.9.6.</a>	On New Congestion Event . . . . .	<a href="#">26</a>
<a href="#">7.9.7.</a>	Process ECN Information . . . . .	<a href="#">27</a>
<a href="#">7.9.8.</a>	On Packets Lost . . . . .	<a href="#">27</a>
<a href="#">8.</a>	Security Considerations . . . . .	<a href="#">27</a>
<a href="#">8.1.</a>	Congestion Signals . . . . .	<a href="#">28</a>
<a href="#">8.2.</a>	Traffic Analysis . . . . .	<a href="#">28</a>
<a href="#">8.3.</a>	Misreporting ECN Markings . . . . .	<a href="#">28</a>
<a href="#">9.</a>	IANA Considerations . . . . .	<a href="#">28</a>
<a href="#">10.</a>	References . . . . .	<a href="#">28</a>
<a href="#">10.1.</a>	Normative References . . . . .	<a href="#">29</a>
<a href="#">10.2.</a>	Informative References . . . . .	<a href="#">29</a>
<a href="#">10.3.</a>	URIs . . . . .	<a href="#">31</a>
<a href="#">Appendix A.</a>	Change Log . . . . .	<a href="#">31</a>
<a href="#">A.1.</a>	Since <a href="#">draft-ietf-quic-recovery-16</a> . . . . .	<a href="#">31</a>
<a href="#">A.2.</a>	Since <a href="#">draft-ietf-quic-recovery-14</a> . . . . .	<a href="#">32</a>
<a href="#">A.3.</a>	Since <a href="#">draft-ietf-quic-recovery-13</a> . . . . .	<a href="#">32</a>
<a href="#">A.4.</a>	Since <a href="#">draft-ietf-quic-recovery-12</a> . . . . .	<a href="#">32</a>
<a href="#">A.5.</a>	Since <a href="#">draft-ietf-quic-recovery-11</a> . . . . .	<a href="#">32</a>
<a href="#">A.6.</a>	Since <a href="#">draft-ietf-quic-recovery-10</a> . . . . .	<a href="#">32</a>
<a href="#">A.7.</a>	Since <a href="#">draft-ietf-quic-recovery-09</a> . . . . .	<a href="#">33</a>
<a href="#">A.8.</a>	Since <a href="#">draft-ietf-quic-recovery-08</a> . . . . .	<a href="#">33</a>
<a href="#">A.9.</a>	Since <a href="#">draft-ietf-quic-recovery-07</a> . . . . .	<a href="#">33</a>
<a href="#">A.10.</a>	Since <a href="#">draft-ietf-quic-recovery-06</a> . . . . .	<a href="#">33</a>
<a href="#">A.11.</a>	Since <a href="#">draft-ietf-quic-recovery-05</a> . . . . .	<a href="#">33</a>
<a href="#">A.12.</a>	Since <a href="#">draft-ietf-quic-recovery-04</a> . . . . .	<a href="#">33</a>
<a href="#">A.13.</a>	Since <a href="#">draft-ietf-quic-recovery-03</a> . . . . .	<a href="#">33</a>
<a href="#">A.14.</a>	Since <a href="#">draft-ietf-quic-recovery-02</a> . . . . .	<a href="#">33</a>
<a href="#">A.15.</a>	Since <a href="#">draft-ietf-quic-recovery-01</a> . . . . .	<a href="#">34</a>
<a href="#">A.16.</a>	Since <a href="#">draft-ietf-quic-recovery-00</a> . . . . .	<a href="#">34</a>
<a href="#">A.17.</a>	Since <a href="#">draft-iyengar-quic-loss-recovery-01</a> . . . . .	<a href="#">34</a>
	Acknowledgments . . . . .	<a href="#">34</a>
	Authors' Addresses . . . . .	<a href="#">34</a>



## **1. Introduction**

QUIC is a new multiplexed and secure transport atop UDP. QUIC builds on decades of transport and security experience, and implements mechanisms that make it attractive as a modern general-purpose transport. The QUIC protocol is described in [[QUIC-TRANSPORT](#)].

QUIC implements the spirit of known TCP loss recovery mechanisms, described in RFCs, various Internet-drafts, and also those prevalent in the Linux TCP implementation. This document describes QUIC congestion control and loss recovery, and where applicable, attributes the TCP equivalent in RFCs, Internet-drafts, academic papers, and/or TCP implementations.

## **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.

Definitions of terms that are used in this document:

ACK-only: Any packet containing only one or more ACK frame(s).

In-flight: Packets are considered in-flight when they have been sent and neither acknowledged nor declared lost, and they are not ACK-only.

Ack-eliciting Frames: All frames besides ACK or PADDING are considered ack-eliciting.

Ack-eliciting Packets: Packets that contain ack-eliciting frames elicit an ACK from the receiver within the maximum ack delay and are called ack-eliciting packets.

Crypto Packets: Packets containing CRYPTO data sent in Initial or Handshake packets.

## **3. Design of the QUIC Transmission Machinery**

All transmissions in QUIC are sent with a packet-level header, which indicates the encryption level and includes a packet sequence number (referred to below as a packet number). The encryption level indicates the packet number space, as described in [[QUIC-TRANSPORT](#)]. Packet numbers never repeat within a packet number space for the



lifetime of a connection. Packet numbers monotonically increase within a space, preventing ambiguity.

This design obviates the need for disambiguating between transmissions and retransmissions and eliminates significant complexity from QUIC's interpretation of TCP loss detection mechanisms.

QUIC packets can contain multiple frames of different types. The recovery mechanisms ensure that data and frames that need reliable delivery are acknowledged or declared lost and sent in new packets as necessary. The types of frames contained in a packet affect recovery and congestion control logic:

- o All packets are acknowledged, though packets that contain no ack-eliciting frames are only acknowledged along with ack-eliciting packets.
- o Long header packets that contain CRYPTO frames are critical to the performance of the QUIC handshake and use shorter timers for acknowledgement and retransmission.
- o Packets that contain only ACK frames do not count toward congestion control limits and are not considered in-flight. Note that this means PADDING frames cause packets to contribute toward bytes in flight without directly causing an acknowledgment to be sent.

### **3.1. Relevant Differences Between QUIC and TCP**

Readers familiar with TCP's loss detection and congestion control will find algorithms here that parallel well-known TCP ones. Protocol differences between QUIC and TCP however contribute to algorithmic differences. We briefly describe these protocol differences below.

#### **3.1.1. Separate Packet Number Spaces**

QUIC uses separate packet number spaces for each encryption level, except 0-RTT and all generations of 1-RTT keys use the same packet number space. Separate packet number spaces ensures acknowledgement of packets sent with one level of encryption will not cause spurious retransmission of packets sent with a different encryption level. Congestion control and RTT measurement are unified across packet number spaces.





### **3.1.2. Monotonically Increasing Packet Numbers**

TCP conflates transmission order at the sender with delivery order at the receiver, which results in retransmissions of the same data carrying the same sequence number, and consequently leads to "retransmission ambiguity". QUIC separates the two: QUIC uses a packet number to indicate transmission order, and any application data is sent in one or more streams, with delivery order determined by stream offsets encoded within STREAM frames.

QUIC's packet number is strictly increasing within a packet number space, and directly encodes transmission order. A higher packet number signifies that the packet was sent later, and a lower packet number signifies that the packet was sent earlier. When a packet containing ack-eliciting frames is detected lost, QUIC rebundles necessary frames in a new packet with a new packet number, removing ambiguity about which packet is acknowledged when an ACK is received. Consequently, more accurate RTT measurements can be made, spurious retransmissions are trivially detected, and mechanisms such as Fast Retransmit can be applied universally, based only on packet number.

This design point significantly simplifies loss detection mechanisms for QUIC. Most TCP mechanisms implicitly attempt to infer transmission ordering based on TCP sequence numbers - a non-trivial task, especially when TCP timestamps are not available.

### **3.1.3. No Reneging**

QUIC ACKs contain information that is similar to TCP SACK, but QUIC does not allow any acked packet to be reneged, greatly simplifying implementations on both sides and reducing memory pressure on the sender.

### **3.1.4. More ACK Ranges**

QUIC supports many ACK ranges, opposed to TCP's 3 SACK ranges. In high loss environments, this speeds recovery, reduces spurious retransmits, and ensures forward progress without relying on timeouts.

### **3.1.5. Explicit Correction For Delayed ACKs**

QUIC ACKs explicitly encode the delay incurred at the receiver between when a packet is received and when the corresponding ACK is sent. This allows the receiver of the ACK to adjust for receiver delays, specifically the delayed ack timer, when estimating the path RTT. This mechanism also allows a receiver to measure and report the delay from when a packet was received by the OS kernel, which is



useful in receivers which may incur delays such as context-switch latency before a userspace QUIC receiver processes a received packet.

#### **4. Generating Acknowledgements**

QUIC SHOULD delay sending acknowledgements in response to packets, but MUST NOT excessively delay acknowledgements of ack-eliciting packets. Specifically, implementations MUST attempt to enforce a maximum ack delay to avoid causing the peer spurious timeouts. The maximum ack delay is communicated in the "max\_ack\_delay" transport parameter and the default value is 25ms.

An acknowledgement SHOULD be sent immediately upon receipt of a second packet but the delay SHOULD NOT exceed the maximum ack delay. QUIC recovery algorithms do not assume the peer generates an acknowledgement immediately when receiving a second full-packet.

Out-of-order packets SHOULD be acknowledged more quickly, in order to accelerate loss recovery. The receiver SHOULD send an immediate ACK when it receives a new packet which is not one greater than the largest received packet number.

Similarly, packets marked with the ECN Congestion Experienced (CE) codepoint in the IP header SHOULD be acknowledged immediately, to reduce the peer's response time to congestion events.

As an optimization, a receiver MAY process multiple packets before sending any ACK frames in response. In this case they can determine whether an immediate or delayed acknowledgement should be generated after processing incoming packets.

##### **4.1. Crypto Handshake Data**

In order to quickly complete the handshake and avoid spurious retransmissions due to crypto retransmission timeouts, crypto packets SHOULD use a very short ack delay, such as 1ms. ACK frames MAY be sent immediately when the crypto stack indicates all data for that packet number space has been received.

##### **4.2. ACK Ranges**

When an ACK frame is sent, one or more ranges of acknowledged packets are included. Including older packets reduces the chance of spurious retransmits caused by losing previously sent ACK frames, at the cost of larger ACK frames.

ACK frames SHOULD always acknowledge the most recently received packets, and the more out-of-order the packets are, the more



important it is to send an updated ACK frame quickly, to prevent the peer from declaring a packet as lost and spuriously retransmitting the frames it contains.

Below is one recommended approach for determining what packets to include in an ACK frame.

#### **4.3. Receiver Tracking of ACK Frames**

When a packet containing an ACK frame is sent, the largest acknowledged in that frame may be saved. When a packet containing an ACK frame is acknowledged, the receiver can stop acknowledging packets less than or equal to the largest acknowledged in the sent ACK frame.

In cases without ACK frame loss, this algorithm allows for a minimum of 1 RTT of reordering. In cases with ACK frame loss and reordering, this approach does not guarantee that every acknowledgement is seen by the sender before it is no longer included in the ACK frame. Packets could be received out of order and all subsequent ACK frames containing them could be lost. In this case, the loss recovery algorithm may cause spurious retransmits, but the sender will continue making forward progress.

#### **5. Computing the RTT estimate**

RTT is calculated when an ACK frame arrives by computing the difference between the current time and the time the largest acked packet was sent. An RTT sample **MUST NOT** be taken for a packet that is not newly acknowledged or not ack-eliciting.

When RTT is calculated, the ack delay field from the ACK frame **SHOULD** be limited to the `max_ack_delay` specified by the peer. Limiting `ack_delay` to `max_ack_delay` ensures a peer specifying an extremely small `max_ack_delay` doesn't cause more spurious timeouts than a peer that correctly specifies `max_ack_delay`. It **SHOULD** be subtracted from the RTT as long as the result is larger than the `min_rtt`. If the result is smaller than the `min_rtt`, the RTT should be used, but the ack delay field should be ignored.

Like TCP, QUIC calculates both smoothed RTT and RTT variance similar to those specified in [[RFC6298](#)].

`min_rtt` is the minimum RTT measured over the connection, prior to adjusting by ack delay. Ignoring ack delay for min RTT prevents intentional or unintentional underestimation of min RTT, which in turn prevents underestimating smoothed RTT.



## 6. Loss Detection

QUIC senders use both ack information and timeouts to detect lost packets, and this section provides a description of these algorithms. Estimating the network round-trip time (RTT) is critical to these algorithms and is described first.

If a packet is lost, the QUIC transport needs to recover from that loss, such as by retransmitting the data, sending an updated frame, or abandoning the frame. For more information, see Section 13.2 of [\[QUIC-TRANSPORT\]](#).

### 6.1. Acknowledgement-based Detection

Acknowledgement-based loss detection implements the spirit of TCP's Fast Retransmit [\[RFC5681\]](#), Early Retransmit [\[RFC5827\]](#), FACK [\[FACK\]](#), SACK loss recovery [\[RFC6675\]](#), and RACK [\[RACK\]](#). This section provides an overview of how these algorithms are implemented in QUIC.

A packet is declared lost if it meets all the following conditions:

- o The packet is unacknowledged, in-flight, and was sent prior to an acknowledged packet.
- o Either its packet number is `kPacketThreshold` smaller than an acknowledged packet ([Section 6.1.1](#)), or it was sent long enough in the past ([Section 6.1.2](#)).

The acknowledgement indicates that a packet sent later was delivered, while the packet and time thresholds provide some tolerance for packet reordering.

Spuriously declaring packets as lost leads to unnecessary retransmissions and may result in degraded performance due to the actions of the congestion controller upon detecting loss. Implementations that detect spurious retransmissions and increase the reordering threshold in packets or time MAY choose to start with smaller initial reordering thresholds to minimize recovery latency.

#### 6.1.1. Packet Threshold

The RECOMMENDED initial value for the packet reordering threshold (`kPacketThreshold`) is 3, based on best practices for TCP loss detection [\[RFC5681\]](#) [\[RFC6675\]](#).

Some networks may exhibit higher degrees of reordering, causing a sender to detect spurious losses. Implementers MAY use algorithms





developed for TCP, such as TCP-NCR [[RFC4653](#)], to improve QUIC's reordering resilience.

#### **6.1.2. Time Threshold**

Once a later packet has been acknowledged, an endpoint SHOULD declare an earlier packet lost if it was sent a threshold amount of time in the past. The time threshold is computed as  $kTimeThreshold * \max(SRTT, latest\_RTT)$ . If packets sent prior to the largest acknowledged packet cannot yet be declared lost, then a timer SHOULD be set for the remaining time.

The RECOMMENDED time threshold ( $kTimeThreshold$ ), expressed as a round-trip time multiplier, is 9/8.

Using  $\max(SRTT, latest\_RTT)$  protects from the two following cases:

- o the latest RTT sample is lower than the SRTT, perhaps due to reordering where packet whose ack triggered the Early Retransmit process encountered a shorter path;
- o the latest RTT sample is higher than the SRTT, perhaps due to a sustained increase in the actual RTT, but the smoothed SRTT has not yet caught up.

Implementers MAY experiment with using other reordering thresholds, including absolute thresholds, bearing in mind that a lower multiplier reduces reordering resilience and increases spurious retransmissions, and a higher multiplier increases loss detection delay.

### **6.2. Timeout Loss Detection**

Timeout loss detection recovers from losses that cannot be handled by acknowledgement-based loss detection. It uses a single timer which switches between a crypto retransmission timer and a probe timer.

#### **6.2.1. Crypto Retransmission Timeout**

Data in CRYPTO frames is critical to QUIC transport and crypto negotiation, so a more aggressive timeout is used to retransmit it.

The initial crypto retransmission timeout SHOULD be set to twice the initial RTT.

At the beginning, there are no prior RTT samples within a connection. Resumed connections over the same network SHOULD use the previous connection's final smoothed RTT value as the resumed connection's



initial RTT. If no previous RTT is available, or if the network changes, the initial RTT SHOULD be set to 100ms. When an acknowledgement is received, a new RTT is computed and the timer SHOULD be set for twice the newly computed smoothed RTT.

When crypto packets are sent, the sender MUST set a timer for the crypto timeout period. Upon timeout, the sender MUST retransmit all unacknowledged CRYPTO data if possible.

Until the server has validated the client's address on the path, the amount of data it can send is limited, as specified in [\[QUIC-TRANSPORT\]](#). If not all unacknowledged CRYPTO data can be sent, then all unacknowledged CRYPTO data sent in Initial packets should be retransmitted. If no data can be sent, then no alarm should be armed until data has been received from the client.

Because the server could be blocked until more packets are received, the client MUST start the crypto retransmission timer even if there is no unacknowledged CRYPTO data. If the timer expires and the client has no CRYPTO data to retransmit and does not have Handshake keys, it SHOULD send an Initial packet in a UDP datagram of at least 1200 bytes. If the client has Handshake keys, it SHOULD send a Handshake packet.

On each consecutive expiration of the crypto timer without receiving an acknowledgement for a new packet, the sender SHOULD double the crypto retransmission timeout and set a timer for this period.

When crypto packets are in flight, the probe timer ([Section 6.2.2](#)) is not active.

#### **[6.2.1.1](#). Retry and Version Negotiation**

A Retry or Version Negotiation packet causes a client to send another Initial packet, effectively restarting the connection process and resetting congestion control and loss recovery state, including resetting any pending timers. Either packet indicates that the Initial was received but not processed. Neither packet can be treated as an acknowledgment for the Initial.

#### **[6.2.1.2](#). Discarding Initial State**

As described in Section 17.5.1 of [\[QUIC-TRANSPORT\]](#), endpoints stop sending and receiving Initial packets once they start exchanging Handshake packets. At this point, all loss recovery state for the Initial packet number space is also discarded. Packets that are in flight for the packet number space are not declared as either



acknowledged or lost. After discarding state, new Initial packets will not be sent.

The client MAY however compute an RTT estimate to the server as the time period from when the first Initial was sent to when a Retry or a Version Negotiation packet is received. The client MAY use this value to seed the RTT estimator for a subsequent connection attempt to the server.

### **6.2.2. Probe Timeout**

A Probe Timeout (PTO) triggers a probe packet when ack-eliciting data is in flight but an acknowledgement is not received within the expected period of time. A PTO enables a connection to recover from loss of tail packets or acks. The PTO algorithm used in QUIC implements the reliability functions of Tail Loss Probe [TLP] [RACK], RTT [RFC5681] and F-RTT algorithms for TCP [RFC5682], and the timeout computation is based on TCP's retransmission timeout period [RFC6298].

#### **6.2.2.1. Computing PTO**

When an ack-eliciting packet is transmitted, the sender schedules a timer for the PTO period as follows:

$$\text{PTO} = \max(\text{smoothed\_rtt} + 4 * \text{rttvar} + \text{max\_ack\_delay}, \text{kGranularity})$$

kGranularity, smoothed\_rtt, rttvar, and max\_ack\_delay are defined in [Section 6.4.1](#) and [Section 6.4.2](#).

The PTO period is the amount of time that a sender ought to wait for an acknowledgement of a sent packet. This time period includes the estimated network roundtrip-time (smoothed\_rtt), the variance in the estimate (4\*rttvar), and max\_ack\_delay, to account for the maximum time by which a receiver might delay sending an acknowledgement.

The PTO value MUST be set to at least kGranularity, to avoid the timer expiring immediately.

When a PTO timer expires, the PTO period MUST be set to twice its current value. This exponential reduction in the sender's rate is important because the PTOs might be caused by loss of packets or acknowledgements due to severe congestion.

A sender computes its PTO timer every time an ack-eliciting packet is sent. A sender might choose to optimize this by setting the timer fewer times if it knows that more ack-eliciting packets will be sent within a short period of time.



#### **6.2.2.2. Sending Probe Packets**

When a PTO timer expires, the sender **MUST** send one ack-eliciting packet as a probe. A sender **MAY** send up to two ack-eliciting packets, to avoid an expensive consecutive PTO expiration due to a single packet loss.

Consecutive PTO periods increase exponentially, and as a result, connection recovery latency increases exponentially as packets continue to be dropped in the network. Sending two packets on PTO expiration increases resilience to packet drops, thus reducing the probability of consecutive PTO events.

Probe packets sent on a PTO **MUST** be ack-eliciting. A probe packet **SHOULD** carry new data when possible. A probe packet **MAY** carry retransmitted unacknowledged data when new data is unavailable, when flow control does not permit new data to be sent, or to opportunistically reduce loss recovery delay. Implementations **MAY** use alternate strategies for determining the content of probe packets, including sending new or retransmitted data based on the application's priorities.

#### **6.2.2.3. Loss Detection**

Delivery or loss of packets in flight is established when an ACK frame is received that newly acknowledges one or more packets.

A PTO timer expiration event does not indicate packet loss and **MUST NOT** cause prior unacknowledged packets to be marked as lost. After a PTO timer has expired, an endpoint uses the following rules to mark packets as lost when an acknowledgement is received that newly acknowledges packets.

When an acknowledgement is received that newly acknowledges packets, loss detection proceeds as dictated by packet and time threshold mechanisms, see [Section 6.1](#).

### **6.3. Tracking Sent Packets**

To correctly implement congestion control, a QUIC sender tracks every ack-eliciting packet until the packet is acknowledged or lost. It is expected that implementations will be able to access this information by packet number and crypto context and store the per-packet fields ([Section 6.3.1](#)) for loss recovery and congestion control.

After a packet is declared lost, it **SHOULD** be tracked for an amount of time comparable to the maximum expected packet reordering, such as 1 RTT. This allows for detection of spurious retransmissions.





Sent packets are tracked for each packet number space, and ACK processing only applies to a single space.

#### **6.3.1. Sent Packet Fields**

`packet_number`: The packet number of the sent packet.

`ack_eliciting`: A boolean that indicates whether a packet is ack-eliciting. If true, it is expected that an acknowledgement will be received, though the peer could delay sending the ACK frame containing it by up to the `MaxAckDelay`.

`in_flight`: A boolean that indicates whether the packet counts towards bytes in flight.

`is_crypto_packet`: A boolean that indicates whether the packet contains cryptographic handshake messages critical to the completion of the QUIC handshake. In this version of QUIC, this includes any packet with the long header that includes a CRYPTO frame.

`sent_bytes`: The number of bytes sent in the packet, not including UDP or IP overhead, but including QUIC framing overhead.

`time_sent`: The time the packet was sent.

#### **6.4. Pseudocode**

##### **6.4.1. Constants of interest**

Constants used in loss recovery are based on a combination of RFCs, papers, and common practice. Some may need to be changed or negotiated in order to better suit a variety of environments.

`kPacketThreshold`: Maximum reordering in packets before packet threshold loss detection considers a packet lost. The RECOMMENDED value is 3.

`kTimeThreshold`: Maximum reordering in time before time threshold loss detection considers a packet lost. Specified as an RTT multiplier. The RECOMMENDED value is 9/8.

`kGranularity`: Timer granularity. This is a system-dependent value. However, implementations SHOULD use a value no smaller than 1ms.

`kInitialRtt`: The RTT used before an RTT sample is taken. The RECOMMENDED value is 100ms.



#### **6.4.2. Variables of interest**

Variables required to implement the congestion control mechanisms are described in this section.

`loss_detection_timer`: Multi-modal timer used for loss detection.

`crypto_count`: The number of times all unacknowledged CRYPTO data has been retransmitted without receiving an ack.

`pto_count`: The number of times a PTO has been sent without receiving an ack.

`time_of_last_sent_ack_eliciting_packet`: The time the most recent ack-eliciting packet was sent.

`time_of_last_sent_crypto_packet`: The time the most recent crypto packet was sent.

`largest_sent_packet`: The packet number of the most recently sent packet.

`largest_acked_packet`: The largest packet number acknowledged in the packet number space so far.

`latest_rtt`: The most recent RTT measurement made when receiving an ack for a previously unacked packet.

`smoothed_rtt`: The smoothed RTT of the connection, computed as described in [[RFC6298](#)]

`rttvar`: The RTT variance, computed as described in [[RFC6298](#)]

`min_rtt`: The minimum RTT seen in the connection, ignoring ack delay.

`max_ack_delay`: The maximum amount of time by which the receiver intends to delay acknowledgments, in milliseconds. The actual `ack_delay` in a received ACK frame may be larger due to late timers, reordering, or lost ACKs.

`loss_time`: The time at which the next packet will be considered lost based on early transmit or exceeding the reordering window in time.

`sent_packets`: An association of packet numbers to information about them. Described in detail above in [Section 6.3](#).



### **6.4.3. Initialization**

At the beginning of the connection, initialize the loss detection variables as follows:

```
loss_detection_timer.reset()
crypto_count = 0
pto_count = 0
loss_time = 0
smoothed_rtt = 0
rttvar = 0
min_rtt = infinite
time_of_last_sent_ack_eliciting_packet = 0
time_of_last_sent_crypto_packet = 0
largest_sent_packet = 0
largest_acked_packet = 0
```

### **6.4.4. On Sending a Packet**

After a packet is sent, information about the packet is stored. The parameters to OnPacketSent are described in detail above in [Section 6.3.1](#).

Pseudocode for OnPacketSent follows:

```
OnPacketSent(packet_number, ack_eliciting, in_flight,
             is_crypto_packet, sent_bytes):
    largest_sent_packet = packet_number
    sent_packets[packet_number].packet_number = packet_number
    sent_packets[packet_number].time_sent = now
    sent_packets[packet_number].ack_eliciting = ack_eliciting
    sent_packets[packet_number].in_flight = in_flight
    if (ack_eliciting):
        if (is_crypto_packet):
            time_of_last_sent_crypto_packet = now
            time_of_last_sent_ack_eliciting_packet = now
        OnPacketSentCC(sent_bytes)
        sent_packets[packet_number].size = sent_bytes
        SetLossDetectionTimer()
```

### **6.4.5. On Receiving an Acknowledgment**

When an ACK frame is received, it may newly acknowledge any number of packets.

Pseudocode for OnAckReceived and UpdateRtt follow:



```
OnAckReceived(ack):
    largest_acked_packet = max(largest_acked_packet,
                               ack.largest_acked)

    // If the largest acknowledged is newly acked and
    // ack-eliciting, update the RTT.
    if (sent_packets[ack.largest_acked] &&
        sent_packets[ack.largest_acked].ack_eliciting):
        latest_rtt =
            now - sent_packets[ack.largest_acked].time_sent
        UpdateRtt(latest_rtt, ack.ack_delay)

    // Process ECN information if present.
    if (ACK frame contains ECN information):
        ProcessECN(ack)

    // Find all newly acked packets in this ACK frame
    newly_acked_packets = DetermineNewlyAckedPackets(ack)
    if (newly_acked_packets.empty()):
        return

    for acked_packet in newly_acked_packets:
        OnPacketAcked(acked_packet.packet_number)

    crypto_count = 0
    pto_count = 0

    DetectLostPackets()
    SetLossDetectionTimer()

UpdateRtt(latest_rtt, ack_delay):
    // min_rtt ignores ack delay.
    min_rtt = min(min_rtt, latest_rtt)
    // Limit ack_delay by max_ack_delay
    ack_delay = min(ack_delay, max_ack_delay)
    // Adjust for ack delay if it's plausible.
    if (latest_rtt - min_rtt > ack_delay):
        latest_rtt -= ack_delay
    // Based on RFC6298.
    if (smoothed_rtt == 0):
        smoothed_rtt = latest_rtt
        rttvar = latest_rtt / 2
    else:
        rttvar_sample = abs(smoothed_rtt - latest_rtt)
        rttvar = 3/4 * rttvar + 1/4 * rttvar_sample
        smoothed_rtt = 7/8 * smoothed_rtt + 1/8 * latest_rtt
```





#### **6.4.6. On Packet Acknowledgment**

When a packet is acknowledged for the first time, the following `OnPacketAked` function is called. Note that a single ACK frame may newly acknowledge several packets. `OnPacketAked` must be called once for each of these newly acknowledged packets.

`OnPacketAked` takes one parameter, `acked_packet`, which is the struct detailed in [Section 6.3.1](#).

Pseudocode for `OnPacketAked` follows:

```
OnPacketAked(acked_packet):  
    if (acked_packet.ack_eliciting):  
        OnPacketAkedCC(acked_packet)  
        sent_packets.remove(acked_packet.packet_number)
```

#### **6.4.7. Setting the Loss Detection Timer**

QUIC loss detection uses a single timer for all timeout loss detection. The duration of the timer is based on the timer's mode, which is set in the packet and timer events further below. The function `SetLossDetectionTimer` defined below shows how the single timer is set.

This algorithm may result in the timer being set in the past, particularly if timers wake up late. Timers set in the past SHOULD fire immediately.

Pseudocode for `SetLossDetectionTimer` follows:



```
SetLossDetectionTimer():
    // Don't arm timer if there are no ack-eliciting packets
    // in flight.
    if (no ack-eliciting packets in flight):
        loss_detection_timer.cancel()
        return

    if (crypto packets are in flight):
        // Crypto retransmission timer.
        if (smoothed_rtt == 0):
            timeout = 2 * kInitialRtt
        else:
            timeout = 2 * smoothed_rtt
        timeout = max(timeout, kGranularity)
        timeout = timeout * (2 ^ crypto_count)
        loss_detection_timer.set(
            time_of_last_sent_crypto_packet + timeout)
        return
    if (loss_time != 0):
        // Time threshold loss detection.
        loss_detection_timer.set(loss_time)
        return

    // Calculate PTO duration
    timeout =
        smoothed_rtt + 4 * rttvar + max_ack_delay
    timeout = max(timeout, kGranularity)
    timeout = timeout * (2 ^ pto_count)

    loss_detection_timer.set(
        time_of_last_sent_ack_eliciting_packet + timeout)
```

#### **6.4.8. On Timeout**

When the loss detection timer expires, the timer's mode determines the action to be performed.

Pseudocode for OnLossDetectionTimeout follows:



```
OnLossDetectionTimeout():
  if (crypto packets are in flight):
    // Crypto retransmission timeout.
    RetransmitUnackedCryptoData()
    crypto_count++
  else if (loss_time != 0):
    // Time threshold loss Detection
    DetectLostPackets()
  else:
    // PTO
    SendTwoPackets()
    pto_count++

  SetLossDetectionTimer()
```

#### **6.4.9. Detecting Lost Packets**

DetectLostPackets is called every time an ACK is received and operates on the sent\_packets for that packet number space. If the loss detection timer expires and the loss\_time is set, the previous largest acknowledged packet is supplied.

Pseudocode for DetectLostPackets follows:



```
DetectLostPackets():
    loss_time = 0
    lost_packets = {}
    loss_delay = kTimeThreshold * max(latest_rtt, smoothed_rtt)

    // Packets sent before this time are deemed lost.
    lost_send_time = now() - loss_delay

    // Packets with packet numbers before this are deemed lost.
    lost_pn = largest_acked_packet - kPacketThreshold

    foreach unacked in sent_packets:
        if (unacked.packet_number > largest_acked_packet):
            continue

        // Mark packet as lost, or set time when it should be marked.
        if (unacked.time_sent <= lost_send_time ||
            unacked.packet_number <= lost_pn):
            sent_packets.remove(unacked.packet_number)
            if (unacked.in_flight):
                lost_packets.insert(unacked)
        else if (loss_time == 0):
            loss_time = unacked.time_sent + loss_delay
        else:
            loss_time = min(loss_time, unacked.time_sent + loss_delay)

    // Inform the congestion controller of lost packets and
    // let it decide whether to retransmit immediately.
    if (!lost_packets.empty()):
        OnPacketsLost(lost_packets)
```

### 6.5. Discussion

The majority of constants were derived from best common practices among widely deployed TCP implementations on the internet. Exceptions follow.

A shorter delayed ack time of 25ms was chosen because longer delayed acks can delay loss recovery and for the small number of connections where less than packet per 25ms is delivered, acking every packet is beneficial to congestion control and loss recovery.

The default initial RTT of 100ms was chosen because it is slightly higher than both the median and mean min\_rtt typically observed on the public internet.





## **7. Congestion Control**

QUIC's congestion control is based on TCP NewReno [[RFC6582](#)]. NewReno is a congestion window based congestion control. QUIC specifies the congestion window in bytes rather than packets due to finer control and the ease of appropriate byte counting [[RFC3465](#)].

QUIC hosts MUST NOT send packets if they would increase `bytes_in_flight` (defined in [Section 7.9.2](#)) beyond the available congestion window, unless the packet is a probe packet sent after a PTO timer expires, as described in [Section 6.2.2](#).

Implementations MAY use other congestion control algorithms, such as Cubic [[RFC8312](#)], and endpoints MAY use different algorithms from one another. The signals QUIC provides for congestion control are generic and are designed to support different algorithms.

### **7.1. Explicit Congestion Notification**

If a path has been verified to support ECN, QUIC treats a Congestion Experienced codepoint in the IP header as a signal of congestion. This document specifies an endpoint's response when its peer receives packets with the Congestion Experienced codepoint. As discussed in [[RFC8311](#)], endpoints are permitted to experiment with other response functions.

### **7.2. Slow Start**

QUIC begins every connection in slow start and exits slow start upon loss or upon increase in the ECN-CE counter. QUIC re-enters slow start anytime the congestion window is less than `ssthresh`, which typically only occurs after an PTO. While in slow start, QUIC increases the congestion window by the number of bytes acknowledged when each acknowledgment is processed.

### **7.3. Congestion Avoidance**

Slow start exits to congestion avoidance. Congestion avoidance in NewReno uses an additive increase multiplicative decrease (AIMD) approach that increases the congestion window by one maximum packet size per congestion window acknowledged. When a loss is detected, NewReno halves the congestion window and sets the slow start threshold to the new congestion window.



#### **[7.4.](#) Recovery Period**

Recovery is a period of time beginning with detection of a lost packet or an increase in the ECN-CE counter. Because QUIC does not retransmit packets, it defines the end of recovery as a packet sent after the start of recovery being acknowledged. This is slightly different from TCP's definition of recovery, which ends when the lost packet that started recovery is acknowledged.

The recovery period limits congestion window reduction to once per round trip. During recovery, the congestion window remains unchanged irrespective of new losses or increases in the ECN-CE counter.

#### **[7.5.](#) Probe Timeout**

Probe packets MUST NOT be blocked by the congestion controller. A sender MUST however count these packets as being additionally in flight, since these packets add network load without establishing packet loss. Note that sending probe packets might cause the sender's bytes in flight to exceed the congestion window until an acknowledgement is received that establishes loss or delivery of packets.

If a threshold number of consecutive PTOs have occurred (`pto_count` is more than `kPersistentCongestionThreshold`, see [Section 7.9.1](#)), the network is considered to be experiencing persistent congestion, and the sender's congestion window MUST be reduced to the minimum congestion window.

#### **[7.6.](#) Pacing**

This document does not specify a pacer, but it is RECOMMENDED that a sender pace sending of all in-flight packets based on input from the congestion controller. For example, a pacer might distribute the congestion window over the SRTT when used with a window-based controller, and a pacer might use the rate estimate of a rate-based controller.

An implementation should take care to architect its congestion controller to work well with a pacer. For instance, a pacer might wrap the congestion controller and control the availability of the congestion window, or a pacer might pace out packets handed to it by the congestion controller. Timely delivery of ACK frames is important for efficient loss recovery. Packets containing only ACK frames should therefore not be paced, to avoid delaying their delivery to the peer.



As an example of a well-known and publicly available implementation of a flow pacer, implementers are referred to the Fair Queue packet scheduler (fq qdisc) in Linux (3.11 onwards).

### **7.7. Sending data after an idle period**

A sender becomes idle if it ceases to send data and has no bytes in flight. A sender's congestion window **MUST** not increase while it is idle.

When sending data after becoming idle, a sender **MUST** reset its congestion window to the initial congestion window (see [Section 4.1 of \[RFC5681\]](#)), unless it paces the sending of packets. A sender **MAY** retain its congestion window if it paces the sending of any packets in excess of the initial congestion window.

A sender **MAY** implement alternate mechanisms to update its congestion window after idle periods, such as those proposed for TCP in [\[RFC7661\]](#).

### **7.8. Discarding Packet Number Space State**

When keys for an packet number space are discarded, any packets sent with those keys are removed from the count of bytes in flight. No loss events will occur any in-flight packets from that space, as a result of discarding loss recovery state (see [Section 6.2.1.2](#)). Note that it is expected that keys are discarded after those packets would be declared lost, but Initial secrets are destroyed earlier.

### **7.9. Pseudocode**

#### **7.9.1. Constants of interest**

Constants used in congestion control are based on a combination of RFCs, papers, and common practice. Some may need to be changed or negotiated in order to better suit a variety of environments.

**kMaxDatagramSize:** The sender's maximum payload size. Does not include UDP or IP overhead. The max packet size is used for calculating initial and minimum congestion windows. The **RECOMMENDED** value is 1200 bytes.

**kInitialWindow:** Default limit on the initial amount of data in flight, in bytes. Taken from [\[RFC6928\]](#). The **RECOMMENDED** value is the minimum of  $10 * kMaxDatagramSize$  and  $\max(2 * kMaxDatagramSize, 14600)$ .



`kMinimumWindow`: Minimum congestion window in bytes. The RECOMMENDED value is  $2 * kMaxDatagramSize$ .

`kLossReductionFactor`: Reduction in congestion window when a new loss event is detected. The RECOMMENDED value is 0.5.

`kPersistentCongestionThreshold`: Number of consecutive PTOs after which network is considered to be experiencing persistent congestion. The rationale for this threshold is to enable a sender to use initial PTOs for aggressive probing, similar to Tail Loss Probe (TLP) in TCP [[TLP](#)] [[RACK](#)]. Once the number of consecutive PTOs reaches this threshold - that is, persistent congestion is established - the sender responds by collapsing its congestion window to `kMinimumWindow`, similar to a Retransmission Timeout (RTO) in TCP [[RFC5681](#)]. The RECOMMENDED value for `kPersistentCongestionThreshold` is 2, which is equivalent to having two TLPs before an RTO in TCP.

#### [7.9.2](#). Variables of interest

Variables required to implement the congestion control mechanisms are described in this section.

`ecn_ce_counter`: The highest value reported for the ECN-CE counter by the peer in an ACK frame. This variable is used to detect increases in the reported ECN-CE counter.

`bytes_in_flight`: The sum of the size in bytes of all sent packets that contain at least one ack-eliciting or PADDING frame, and have not been acked or declared lost. The size does not include IP or UDP overhead, but does include the QUIC header and AEAD overhead. Packets only containing ACK frames do not count towards `bytes_in_flight` to ensure congestion control does not impede congestion feedback.

`congestion_window`: Maximum number of bytes-in-flight that may be sent.

`recovery_start_time`: The time when QUIC first detects a loss, causing it to enter recovery. When a packet sent after this time is acknowledged, QUIC exits recovery.

`ssthresh`: Slow start threshold in bytes. When the congestion window is below `ssthresh`, the mode is slow start and the window grows by the number of bytes acknowledged.





### **7.9.3. Initialization**

At the beginning of the connection, initialize the congestion control variables as follows:

```
congestion_window = kInitialWindow
bytes_in_flight = 0
recovery_start_time = 0
ssthresh = infinite
ecn_ce_counter = 0
```

### **7.9.4. On Packet Sent**

Whenever a packet is sent, and it contains non-ACK frames, the packet increases bytes\_in\_flight.

```
OnPacketSentCC(bytes_sent):
    bytes_in_flight += bytes_sent
```

### **7.9.5. On Packet Acknowledgement**

Invoked from loss detection's OnPacketAked and is supplied with the acked\_packet from sent\_packets.

```
InRecovery(sent_time):
    return sent_time <= recovery_start_time

OnPacketAkedCC(acked_packet):
    // Remove from bytes_in_flight.
    bytes_in_flight -= acked_packet.size
    if (InRecovery(acked_packet.time_sent)):
        // Do not increase congestion window in recovery period.
        return
    if (congestion_window < ssthresh):
        // Slow start.
        congestion_window += acked_packet.size
    else:
        // Congestion avoidance.
        congestion_window += kMaxDatagramSize * acked_packet.size
        / congestion_window
```

### **7.9.6. On New Congestion Event**

Invoked from ProcessECN and OnPacketsLost when a new congestion event is detected. May start a new recovery period and reduces the congestion window.



```
CongestionEvent(sent_time):
    // Start a new congestion event if the sent time is larger
    // than the start time of the previous recovery epoch.
    if (!InRecovery(sent_time)):
        recovery_start_time = Now()
        congestion_window *= kLossReductionFactor
        congestion_window = max(congestion_window, kMinimumWindow)
        ssthresh = congestion_window
        // Collapse congestion window if persistent congestion
        if (pto_count > kPersistentCongestionThreshold):
            congestion_window = kMinimumWindow
```

#### [7.9.7.](#) Process ECN Information

Invoked when an ACK frame with an ECN section is received from the peer.

```
ProcessECN(ack):
    // If the ECN-CE counter reported by the peer has increased,
    // this could be a new congestion event.
    if (ack.ce_counter > ecn_ce_counter):
        ecn_ce_counter = ack.ce_counter
        // Start a new congestion event if the last acknowledged
        // packet was sent after the start of the previous
        // recovery epoch.
        CongestionEvent(sent_packets[ack.largest_acked].time_sent)
```

#### [7.9.8.](#) On Packets Lost

Invoked by loss detection from DetectLostPackets when new packets are detected lost.

```
OnPacketsLost(lost_packets):
    // Remove lost packets from bytes_in_flight.
    for (lost_packet : lost_packets):
        bytes_in_flight -= lost_packet.size
    largest_lost_packet = lost_packets.last()

    // Start a new congestion epoch if the last lost packet
    // is past the end of the previous recovery epoch.
    CongestionEvent(largest_lost_packet.time_sent)
```

### [8.](#) Security Considerations



### **8.1. Congestion Signals**

Congestion control fundamentally involves the consumption of signals - both loss and ECN codepoints - from unauthenticated entities. On-path attackers can spoof or alter these signals. An attacker can cause endpoints to reduce their sending rate by dropping packets, or alter send rate by changing ECN codepoints.

### **8.2. Traffic Analysis**

Packets that carry only ACK frames can be heuristically identified by observing packet size. Acknowledgement patterns may expose information about link characteristics or application behavior. Endpoints can use PADDING frames or bundle acknowledgments with other frames to reduce leaked information.

### **8.3. Misreporting ECN Markings**

A receiver can misreport ECN markings to alter the congestion response of a sender. Suppressing reports of ECN-CE markings could cause a sender to increase their send rate. This increase could result in congestion and loss.

A sender MAY attempt to detect suppression of reports by marking occasional packets that they send with ECN-CE. If a packet marked with ECN-CE is not reported as having been marked when the packet is acknowledged, the sender SHOULD then disable ECN for that path.

Reporting additional ECN-CE markings will cause a sender to reduce their sending rate, which is similar in effect to advertising reduced connection flow control limits and so no advantage is gained by doing so.

Endpoints choose the congestion controller that they use. Though congestion controllers generally treat reports of ECN-CE markings as equivalent to loss [[RFC8311](#)], the exact response for each controller could be different. Failure to correctly respond to information about ECN markings is therefore difficult to detect.

## **9. IANA Considerations**

This document has no IANA actions. Yet.

## **10. References**



### **10.1. Normative References**

[QUIC-TRANSPORT]

Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", [draft-ietf-quic-transport-17](#) (work in progress), December 2018.

[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>>.

[RFC8311] Black, D., "Relaxing Restrictions on Explicit Congestion Notification (ECN) Experimentation", [RFC 8311](#), DOI 10.17487/RFC8311, January 2018, <<https://www.rfc-editor.org/info/rfc8311>>.

### **10.2. Informative References**

[FACK] Mathis, M. and J. Mahdavi, "Forward Acknowledgement: Refining TCP Congestion Control", ACM SIGCOMM , August 1996.

[RACK] Cheng, Y., Cardwell, N., Dukkupati, N., and P. Jha, "RACK: a time-based fast loss detection algorithm for TCP", [draft-ietf-tcpm-rack-04](#) (work in progress), July 2018.

[RFC3465] Allman, M., "TCP Congestion Control with Appropriate Byte Counting (ABC)", [RFC 3465](#), DOI 10.17487/RFC3465, February 2003, <<https://www.rfc-editor.org/info/rfc3465>>.

[RFC4653] Bhandarkar, S., Reddy, A., Allman, M., and E. Blanton, "Improving the Robustness of TCP to Non-Congestion Events", [RFC 4653](#), DOI 10.17487/RFC4653, August 2006, <<https://www.rfc-editor.org/info/rfc4653>>.

[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>>.





- [RFC5682] Sarolahti, P., Kojo, M., Yamamoto, K., and M. Hata, "Forward RTO-Recovery (F-RTO): An Algorithm for Detecting Spurious Retransmission Timeouts with TCP", [RFC 5682](#), DOI 10.17487/RFC5682, September 2009, <<https://www.rfc-editor.org/info/rfc5682>>.
- [RFC5827] Allman, M., Avrachenkov, K., Ayesta, U., Blanton, J., and P. Hurtig, "Early Retransmit for TCP and Stream Control Transmission Protocol (SCTP)", [RFC 5827](#), DOI 10.17487/RFC5827, May 2010, <<https://www.rfc-editor.org/info/rfc5827>>.
- [RFC6298] Paxson, V., Allman, M., Chu, J., and M. Sargent, "Computing TCP's Retransmission Timer", [RFC 6298](#), DOI 10.17487/RFC6298, June 2011, <<https://www.rfc-editor.org/info/rfc6298>>.
- [RFC6582] Henderson, T., Floyd, S., Gurtov, A., and Y. Nishida, "The NewReno Modification to TCP's Fast Recovery Algorithm", [RFC 6582](#), DOI 10.17487/RFC6582, April 2012, <<https://www.rfc-editor.org/info/rfc6582>>.
- [RFC6675] Blanton, E., Allman, M., Wang, L., Jarvinen, I., Kojo, M., and Y. Nishida, "A Conservative Loss Recovery Algorithm Based on Selective Acknowledgment (SACK) for TCP", [RFC 6675](#), DOI 10.17487/RFC6675, August 2012, <<https://www.rfc-editor.org/info/rfc6675>>.
- [RFC6928] Chu, J., Dukkupati, 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>>.
- [RFC7661] Fairhurst, G., Sathiaselan, A., and R. Secchi, "Updating TCP to Support Rate-Limited Traffic", [RFC 7661](#), DOI 10.17487/RFC7661, October 2015, <<https://www.rfc-editor.org/info/rfc7661>>.
- [RFC8312] Rhee, I., Xu, L., Ha, S., Zimmermann, A., Eggert, L., and R. Scheffenegger, "CUBIC for Fast Long-Distance Networks", [RFC 8312](#), DOI 10.17487/RFC8312, February 2018, <<https://www.rfc-editor.org/info/rfc8312>>.
- [TLP] Dukkupati, N., Cardwell, N., Cheng, Y., and M. Mathis, "Tail Loss Probe (TLP): An Algorithm for Fast Recovery of Tail Losses", [draft-dukkupati-tcpm-tcp-loss-probe-01](#) (work in progress), February 2013.



### **10.3. URIs**

- [1] [https://mailarchive.ietf.org/arch/search/?email\\_list=quic](https://mailarchive.ietf.org/arch/search/?email_list=quic)
- [2] <https://github.com/quicwg>
- [3] <https://github.com/quicwg/base-drafts/labels/-recovery>

### **Appendix A. Change Log**

\*RFC Editor's Note:\* Please remove this section prior to publication of a final version of this document.

Issue and pull request numbers are listed with a leading octothorp.

#### **A.1. Since [draft-ietf-quic-recovery-16](#)**

- o Unify TLP and RT0 into a single PT0; eliminate min RT0, min TLP and min crypto timeouts; eliminate timeout validation (#2114, #2166, #2168, #1017)
- o Redefine how congestion avoidance in terms of when the period starts (#1928, #1930)
- o Document what needs to be tracked for packets that are in flight (#765, #1724, #1939)
- o Integrate both time and packet thresholds into loss detection (#1969, #1212, #934, #1974)
- o Reduce congestion window after idle, unless pacing is used (#2007, #2023)
- o Disable RTT calculation for packets that don't elicit acknowledgment (#2060, #2078)
- o Limit ack\_delay by max\_ack\_delay (#2060, #2099)
- o Initial keys are discarded once Handshake are available (#1951, #2045)
- o Reorder ECN and loss detection in pseudocode (#2142)
- o Only cancel loss detection timer if ack-eliciting packets are in flight (#2093, #2117)



**A.2. Since [draft-ietf-quick-recovery-14](#)**

- o Used max\_ack\_delay from transport params (#1796, #1782)
- o Merge ACK and ACK\_ECN (#1783)

**A.3. Since [draft-ietf-quick-recovery-13](#)**

- o Corrected the lack of ssthresh reduction in CongestionEvent pseudocode (#1598)
- o Considerations for ECN spoofing (#1426, #1626)
- o Clarifications for PADDING and congestion control (#837, #838, #1517, #1531, #1540)
- o Reduce early retransmission timer to RTT/8 (#945, #1581)
- o Packets are declared lost after an RT0 is verified (#935, #1582)

**A.4. Since [draft-ietf-quick-recovery-12](#)**

- o Changes to manage separate packet number spaces and encryption levels (#1190, #1242, #1413, #1450)
- o Added ECN feedback mechanisms and handling; new ACK\_ECN frame (#804, #805, #1372)

**A.5. Since [draft-ietf-quick-recovery-11](#)**

No significant changes.

**A.6. Since [draft-ietf-quick-recovery-10](#)**

- o Improved text on ack generation (#1139, #1159)
- o Make references to TCP recovery mechanisms informational (#1195)
- o Define time\_of\_last\_sent\_handshake\_packet (#1171)
- o Added signal from TLS the data it includes needs to be sent in a Retry packet (#1061, #1199)
- o Minimum RTT (min\_rtt) is initialized with an infinite value (#1169)



**[A.7.](#)** Since [draft-ietf-quic-recovery-09](#)

No significant changes.

**[A.8.](#)** Since [draft-ietf-quic-recovery-08](#)

- o Clarified pacing and RT0 (#967, #977)

**[A.9.](#)** Since [draft-ietf-quic-recovery-07](#)

- o Include Ack Delay in RT0(and TLP) computations (#981)
- o Ack Delay in SRTT computation (#961)
- o Default RTT and Slow Start (#590)
- o Many editorial fixes.

**[A.10.](#)** Since [draft-ietf-quic-recovery-06](#)

No significant changes.

**[A.11.](#)** Since [draft-ietf-quic-recovery-05](#)

- o Add more congestion control text (#776)

**[A.12.](#)** Since [draft-ietf-quic-recovery-04](#)

No significant changes.

**[A.13.](#)** Since [draft-ietf-quic-recovery-03](#)

No significant changes.

**[A.14.](#)** Since [draft-ietf-quic-recovery-02](#)

- o Integrate F-RT0 (#544, #409)
- o Add congestion control (#545, #395)
- o Require connection abort if a skipped packet was acknowledged (#415)
- o Simplify RT0 calculations (#142, #417)





**A.15.** Since [draft-ietf-quic-recovery-01](#)

- o Overview added to loss detection
- o Changes initial default RTT to 100ms
- o Added time-based loss detection and fixes early retransmit
- o Clarified loss recovery for handshake packets
- o Fixed references and made TCP references informative

**A.16.** Since [draft-ietf-quic-recovery-00](#)

- o Improved description of constants and ACK behavior

**A.17.** Since [draft-iyengar-quic-loss-recovery-01](#)

- o Adopted as base for [draft-ietf-quic-recovery](#)
- o Updated authors/editors list
- o Added table of contents

## Acknowledgments

## Authors' Addresses

Jana Iyengar (editor)

Fastly

Email: [jri.ietf@gmail.com](mailto:jri.ietf@gmail.com)

Ian Swett (editor)

Google

Email: [ianswett@google.com](mailto:ianswett@google.com)

