

Workgroup: IPPM

Internet-Draft:

draft-mdt-ippm-explicit-flow-measurements-00

Published: 2 November 2020

Intended Status: Informational

Expires: 6 May 2021

Authors: M. Cociglio      A. Ferrieux      G. Fioccola  
Telecom Italia      Orange Labs      Huawei Technologies  
I. Lubashev      F. Bulgarella      I. Hamchaoui  
Akamai Technologies      Telecom Italia      Orange Labs  
M. Nilo      R. Sisto  
Telecom Italia      Politecnico di Torino  
D. Tikhonov  
LiteSpeed Technologies

## **Explicit Flow Measurements Techniques**

### **Abstract**

This document describes protocol independent methods called Explicit Flow Measurement Techniques that employ few marking bits, inside the header of each packet, for loss and delay measurement. The endpoints, marking the traffic, signal these metrics to intermediate observers allowing them to measure connection performance, and to locate the network segment where impairments happen. Different alternatives are considered within this document. These signaling methods apply to all protocols but they are especially valuable when applied to protocols that encrypt transport header and do not allow traditional methods for delay and loss detection.

### **Discussion Venues**

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

Discussion of this document takes place on the IPPM Working Group mailing list ([ippm@ietf.org](mailto:ippm@ietf.org)), which is archived at <https://mailarchive.ietf.org/arch/browse/ippm/>.

Source for this draft and an issue tracker can be found at <https://github.com/igorlord/draft-xxx-ippm-flow-measurements>.

### **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 6 May 2021.

## Copyright Notice

Copyright (c) 2020 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](#)
2. [Notational Conventions](#)
3. [Latency Bits](#)
  - 3.1. [Spin Bit](#)
  - 3.2. [Delay Bit](#)
    - 3.2.1. [Generation Phase](#)
    - 3.2.2. [Reflection Phase](#)
    - 3.2.3. [Two Bits Delay Measurement: Spin Bit + Delay Bit](#)
    - 3.2.4. [Observer's Algorithm and Edge Rejection Interval](#)
4. [Loss Bits](#)
  - 4.1. [T Bit - Round Trip Loss Bit](#)
    - 4.1.1. [Round Trip Packet Loss Measurement](#)
    - 4.1.2. [Setting the Round Trip Loss Bit on Outgoing Packets](#)
    - 4.1.3. [Observer's Logic for Round Trip Loss Signal](#)
    - 4.1.4. [Loss Coverage and Signal Timing](#)
  - 4.2. [Q Bit - Square Bit](#)
    - 4.2.1. [Q Block Length Selection](#)
    - 4.2.2. [Upstream Loss](#)
    - 4.2.3. [Identifying Q Block Boundaries](#)
  - 4.3. [L Bit - Loss Event Bit](#)
    - 4.3.1. [End-To-End Loss](#)
    - 4.3.2. [Loss Profile Characterization](#)
  - 4.4. [L+Q Bits - Upstream, Downstream, and End-to-End Loss Measurements](#)
    - 4.4.1. [Correlating End-to-End and Upstream Loss](#)

- [4.5. R Bit - Reflection Square Bit](#)
  - [4.5.1. R+Q Bits - Using R and Q Bits for Passive Loss Measurement](#)
  - [4.5.2. Enhancement of R Block Length Computation](#)
  - [4.5.3. Improved Resilience to Packet Reordering](#)
- [5. Summary of Delay and Loss Marking Methods](#)
- [6. ECN-Echo Event Bit](#)
  - [6.1. Setting the ECN-Echo Event Bit on Outgoing Packets](#)
  - [6.2. Using E Bit for Passive ECN-Reported Congestion Measurement](#)
- [7. Protocol Ossification Considerations](#)
- [8. Examples of Application](#)
  - [8.1. QUIC](#)
  - [8.2. TCP](#)
- [9. Security Considerations](#)
  - [9.1. Optimistic ACK Attack](#)
- [10. Privacy Considerations](#)
- [11. IANA Considerations](#)
- [12. Change Log](#)
- [13. Contributors](#)
- [14. Acknowledgements](#)
- [15. References](#)
  - [15.1. Normative References](#)
  - [15.2. Informative References](#)
- [Authors' Addresses](#)

## **1. Introduction**

Packet loss and delay are hard and pervasive problems of day-to-day network operation. Proactively detecting, measuring, and locating them is crucial to maintaining high QoS and timely resolution of crippling end-to-end throughput issues. To this effect, in a TCP-dominated world, network operators have been heavily relying on information present in the clear in TCP headers: sequence and acknowledgment numbers and SACKs when enabled (see [RFC8517]). These allow for quantitative estimation of packet loss and delay by passive on-path observation. Additionally, the problem can be quickly identified in the network path by moving the passive observer around.

With encrypted protocols, the equivalent transport headers are encrypted and passive packet loss and delay observations are not possible, as described in [TRANSPORT-ENCRYPT].

Measuring TCP loss and delay between similar endpoints cannot be relied upon to evaluate encrypted protocol loss and delay. Different protocols could be routed by the network differently, and the fraction of Internet traffic delivered using protocols other than TCP is increasing every year. It is imperative to measure packet loss and delay experienced by encrypted protocol users directly.

This document defines Explicit Flow Measurement Techniques. These hybrid measurement path signals (see [[IPM-Methods](#)]) are to be embedded into a transport layer protocol and are explicitly intended for exposing RTT and loss rate information to on-path measurement devices. These measurement mechanisms are applicable to any transport-layer protocol, and, as an example, the document describes QUIC and TCP bindings.

The Explicit Flow Measurement Techniques described in this document can be used alone or in combination with other Explicit Flow Measurement Techniques. Each technique uses a small number of bits and exposes a specific measurement.

Following the recommendation in [[RFC8558](#)] of making path signals explicit, this document proposes adding a small number of dedicated measurement bits to the clear portion of the protocol headers. These bits can be added to an encrypted portion of a header belonging to any protocol layer, e.g. IP (see [[IP](#)]) and IPv6 (see [[IPv6](#)]) headers or extensions, such as [[IPv6AltMark](#)], UDP surplus space (see [[UDP-OPTIONS](#)] and [[UDP-SURPLUS](#)]), reserved bits in a QUIC v1 header (see [[QUIC-TRANSPORT](#)]).

The measurements are not designed for use in automated control of the network in environments where signal bits are set by untrusted hosts. Instead, the signal is to be used for troubleshooting individual flows as well as for monitoring the network by aggregating information from multiple flows and raising operator alarms if aggregate statistics indicate a potential problem.

The spin bit, delay bit and loss bits explained in this document are inspired by [[AltMark](#)], [[SPIN-BIT](#)], [[I-D.trammell-tsvwg-spin](#)] and [[I-D.trammell-ippm-spin](#)].

Additional details about the Performance Measurements for QUIC are described in the paper [[ANRW19-PM-QUIC](#)].

## 2. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

## 3. Latency Bits

This section introduces bits that can be used for round trip latency measurements. Whenever this section of the specification refers to packets, it is referring only to packets with protocol headers that include the latency bits.

[[QUIC-TRANSPORT](#)] introduces an explicit per-flow transport-layer signal for hybrid measurement of RTT. This signal consists of a spin bit that toggles once per RTT. [[SPIN-BIT](#)] discusses an additional two-bit Valid Edge Counter (VEC) to compensate for loss and reordering of the spin bit and increase fidelity of the signal in less than ideal network conditions.

This document introduces an additional single-bit delay signal that can be used together with the spin bit by passive observers to measure the RTT of a network flow, avoiding the spin bit ambiguities that arise as soon as network conditions deteriorate.

### 3.1. Spin Bit

This section is a small recap of the spin bit working mechanism. For a comprehensive explanation of the algorithm, please see [[SPIN-BIT](#)].

The spin bit is an alternate marking [[AltMark](#)] generated signal, where the size of the alternation changes with the flight size each RTT.

The latency spin bit is a single bit signal that toggles once per RTT, enabling latency monitoring of a connection-oriented communication from intermediate observation points.

A "spin period" is a set of packets with the same spin bit value sent during one RTT time interval. A "spin period value" is the value of the spin bit shared by all packets in a spin period.

The client and server maintain an internal per-connection spin value (i.e. 0 or 1) used to set the spin bit on outgoing packets. Both endpoints initialize the spin value to 0 when a new connection starts. Then:

- \*when the client receives a packet with the packet number larger than any number seen so far, it sets the connection spin value to the opposite value contained in the received packet;

- \*when the server receives a packet with the packet number larger than any number seen so far, it sets the connection spin value to the same value contained in the received packet.

The computed spin value is used by the endpoints for setting the spin bit on outgoing packets. This mechanism allows the endpoints to generate a square wave such that, by measuring the distance in time between pairs of consecutive edges observed in the same direction, a passive on-path observer can compute the round trip delay of that network flow.

Spin bit enables round trip latency measurement by observing a single direction of the traffic flow.

Note that packet reordering can cause spurious edges that require heuristics to correct. The spin bit performance deteriorates as soon as network impairments arise as explained in [Section 3.2](#).

### **3.2. Delay Bit**

The delay bit, different from a two-bit VEC, has been designed to overcome accuracy limitations experienced by the spin bit under difficult network conditions:

- \*packet reordering leads to generation of spurious edges and errors in delay estimation;
- \*loss of edges causes wrong estimation of spin periods and therefore wrong RTT measurements;
- \*application-limited senders cause the spin bit to measure the application delays instead of network delays.


If enabled, delay bit has to be used in addition to the spin bit. Unlike the spin bit, which is set in every packet transmitted on the network, the delay bit is set only once per round trip.

When the delay bit is used, a single packet with a second marked bit (the delay bit) bounces between a client and a server during the entire connection lifetime. This single packet is called "delay sample".


An observer placed at an intermediate point, observing a single direction of traffic, tracking the delay sample and the relative timestamp in every spin period, can measure the round trip delay of the connection.

The delay sample lifetime is comprised of two phases: initialization and reflection. The initialization is the generation of the delay sample, while the reflection realizes the bounce behavior of this single packet between the two endpoints.


The next figure describes the Delay bit mechanism: the first bit is the spin bit and the second one is the delay bit.




(a) No traffic at beginning.




(b) The Client starts sending data and sets the first packet as Delay Sample.




(c) The Server starts sending data and reflects the Delay Sample.



(d) The Client inverts the spin bit and reflects the Delay Sample.



(e) The Server reflects the Delay Sample.



(f) The client reverts the spin  
bit and reflects the Delay Sample.



Figure 1: Spin bit and Delay bit

### 3.2.1. Generation Phase

Only client is actively involved in the generation phase.

When connection starts and spin bit is set to 0, the client initializes the delay bit of the first packet to 1, so it becomes the delay sample for that marking period. Only this packet is marked with the delay bit set to 1 for this round trip period; the other ones will carry the spin bit, while the delay bit will be set to 0.

The server initializes the delay bit to 0 at the beginning of the connection, and its only task during the connection is described in [Section 3.2.2](#).

In absence of network impairments, the delay sample should bounce between client and server continuously, for the entire duration of the connection. That is highly unlikely for two reasons:

1. the packet carrying the delay bit might be lost;
2. an endpoint could stop or delay sending packets because the application is limiting the amount of traffic transmitted;

To deal with these problems, the algorithm provides a procedure named "recovery process" to regenerate the delay sample and to inform a possible observer of the problem so the measurement can be restarted.

#### 3.2.1.1. The Recovery Process

Absent packet loss or reordering, every spin period ends with a delay sample inside. If that does not happen and a spin period terminates without a delay sample inside, the client waits a further spin period; then, it creates a new delay sample by setting the delay bit to 1 on the first outgoing packet of the subsequent period.

The spin period with all delay bits set to 0 informs observers that there was a problem and delay measurements for this flow should be reset till the next delay sample is received.

### 3.2.2. Reflection Phase

Reflection is the process that enables the bouncing of the delay sample between a client and a server. The behavior of the two endpoints is slightly different.

\*Server side reflection: when a delay sample arrives, the server marks the first packet in the opposite direction as the delay sample, if the outgoing packet has the same spin bit value as the delay sample. Otherwise, the delay sample is ignored.

\*Client side reflection: when a delay sample arrives, the client marks the first packet in the opposite direction as the delay sample, if the outgoing packet has the opposite spin bit value then the delay sample. Otherwise, the delay sample is ignored.

In both cases, if the outgoing delay sample is being transmitted with a delay greater than a predetermined threshold after the reception of the incoming delay sample (1ms by default), the delay sample is not reflected, and the outgoing delay bit is kept at 0.

Note that reflection takes place for the delay sample regardless of its position within the spin period, as long as it stays within its original spin period.

A time threshold for the retransmission of the delay sample is used to eliminate measurements that would overestimate the delay due to lack of traffic on the endpoints. Hence, the maximum estimation error would amount to twice the threshold (e.g. 2ms) per measurement.

### 3.2.3. Two Bits Delay Measurement: Spin Bit + Delay Bit


When the Delay Bit is used, a passive observer can use delay samples directly and avoid inherent ambiguities in the calculation of the RTT in spin bit analysis, such as heuristic validation of the goodness of an edge signal.

#### 3.2.3.1. RTT Measurement


The delay sample generation process ensures that only one packet marked with the delay bit set to 1 runs back and forth between two endpoints per round trip time. To determine the RTT measurement of a flow, an on-path passive observer computes the time difference between two delay samples observed in a single direction.

To ensure a valid measurement, the observer must identify spin periods in the packet flow and assign delay samples to spin periods. If a spin period is missing a delay sample, the measurement needs to be restarted from the subsequent delay sample. Hence, measurements

must take into account delay samples belonging to adjacent spin periods.



(a) client-server RTT



(b) server-client RTT


Figure 2: Round-trip time (both direction)

### 3.2.3.2. Half-RTT Measurement


An observer that is able to observe both forward and return traffic directions can use the delay samples to measure "upstream" and "downstream" RTT components, also known as the half-RTT measurements. It does this by measuring the time between a delay sample observed in one direction and the reflected delay sample observed in the opposite direction.

As with RTT measurement, the observer must identify spin periods in the packet flow and assign delay samples to spin periods. If a spin period is missing a delay sample, the measurement needs to be restarted from the subsequent delay sample. So a measurement, to be valid, must take into account delay samples belonging to adjacent periods, for the upstream component, or to the same period for the downstream component.

Note that upstream and downstream sections of paths between the endpoints and the observer, i.e. observer-to-client vs client-to-observer and observer-to-server vs server-to-observer, may have different delay characteristics due to the difference in network congestion and other factors.



(a) client-observer half-RTT




(b) observer-server half-RTT


Figure 3: Half Round-trip time (both direction)

### 3.2.3.3. Intra-Domain RTT Measurement

Intra-domain RTT is the portion of the entire RTT used by a flow to traverse the network of a provider. To measure intra-domain RTT, two observers capable of observing traffic in both directions must be employed simultaneously at ingress and egress of the network to be measured. Intra-domain RTT is difference between the two computed upstream (or downstream) RTT components.



(a) client-observer RTT components (half-RTTs)



(b) the intra-domain RTT resulting from the subtraction of the above RTT components

Figure 4: Intra-domain Round-trip time (client-observer: upstream)

### 3.2.4. Observer's Algorithm and Edge Rejection Interval

To provide a formal description of the observer behavior, we define a "matching delay sample" to be a delay sample with the spin bit value that matched the spin bit value of then-current spin period.

Upon detecting a matching delay sample, if a matching delay sample was also detected in the previous period, then the two delay samples can be used to calculate RTT measurement.

If the observer can observe both forward and return traffic flows, and it is able to determine which direction contains the client and the server (e.g. by observing the spin bit or connection handshake):

- \*If matching delay samples have been detected in both directions in the current spin period, they can be used to measure the observer-server half-RTT.

- \*If a matching delay sample has been detected in client-to-observer direction, AND a matching delay sample had been detected in observer-to-client direction in the previous spin period, they can be used to measure the observer-client half-RTT.

The described observer behavior depends on the ability to accurately identify current spin periods and to reject spurious spin edges, caused by packet reordering. Failure to do so will lead to many missed measurement opportunities and will decrease the amount of usable delay samples available to the observer.

To implement spurious edge rejection, every time a spin bit edge is detected, the observer starts a new spin period and begins a time interval during which it rejects spin edges (e.g. 5ms). This guarantees protection against spurious edges due to packets that have been reordered by less than the time interval. The mechanism only works for intervals smaller than the RTT of the observed connection; if RTT is smaller than the edge rejection interval, the observer cannot measure the RTT.

## 4. Loss Bits

This section introduces bits that can be used for loss measurements. Whenever this section of the specification refers to packets, it is referring only to packets with protocol headers that include the loss bits - the only packets whose loss can be measured.

- \*T: the "round Trip loss" bit is used in combination with the Spin bit to measure round-trip loss. See [Section 4.1](#).

- \*Q: the "sQuare signal" bit is used to measure upstream loss. See [Section 4.2](#).

\*L: the "Loss event" bit is used to measure end-to-end loss. See [Section 4.3](#).

\*R: the "Reflection square signal" bit is used in combination with Q bit to measure end-to-end loss. See [Section 4.1](#).

Loss measurements enabled by T, Q, and L bits can be implemented by those loss bits alone (T bit requires a working Spin Bit). Two-bit combinations Q+L and Q+R enable additional measurement opportunities discussed below.

Each endpoint maintains appropriate counters independently and separately for each separately identifiable flow (each sub-flow for multipath connections).

Since loss is reported independently for each flow, all bits (except for L bit) require a certain minimum number of packets to be exchanged per flow before any signal can be measured. Therefore, loss measurements work best for flows that transfer more than a minimal amount of data.

#### **4.1. T Bit - Round Trip Loss Bit**

The round Trip loss bit is used to mark a variable number of packets exchanged twice between the endpoints realizing a two round-trip reflection. A passive on-path observer, observing either direction, can count and compare the number of marked packets seen during the two reflections, estimating the loss rate experienced by the connection. The overall exchange comprises:

- \*The client selects, generates and consequently transmits a first train of packets, by setting the T bit to 1;
- \*The server, upon receiving each packet included in the first train, reflects to the client a respective second train of packets of the same size as the first train received, by setting the T bit to 1;
- \*The client, upon receiving each packet included in the second train, reflects to the server a respective third train of packets of the same size as the second train received, by setting the T bit to 1;
- \*The server, upon receiving each packet included in the third train, finally reflects to the client a respective fourth train of packets of the same size as the third train received, by setting the T bit to 1.


Packets belonging to the first round trip (first and second train) represent the Generation Phase, while those belonging to the second round trip (third and fourth train) represent the Reflection Phase.

A passive on-path observer can count and compare the number of marked packets seen during the two round trips (i.e. the first and third or the second and the fourth trains of packets, depending on which direction is observed) and estimate the loss rate experienced by the connection. This process is repeated continuously to obtain more measurements as long as the endpoints exchange traffic. These measurements can be called Round Trip losses.


Since packet rates in two directions may be different, the number of marked packets in the train is determined by the direction with the lowest packet rate. See [Section 4.1.2](#) for details on packet generation and for a mechanism to allow an observer to distinguish between trains belonging to different phases (Generation and Reflection).

#### 4.1.1.1. Round Trip Packet Loss Measurement

Since the measurements are performed on a portion of the traffic exchanged between the client and the server, the observer calculates the end-to-end Round Trip Packet Loss (RTPL) that, statistically, will correspond to the loss rate experienced by the connection along the entire network path.



(a) client-server RTPL



(b) server-client RTPL

Figure 5: Round-trip packet loss (both direction)

This methodology also allows the Half-RTPL measurement and the Intra-domain RTPL measurement in a way similar to RTT measurement.

```

=====>
= *****      -----|----->      *****
= * Client *      Obs      * Server *
= *****      <-----|-----      *****
<=====

```

(a) client-observer half-RTPL

```

=====>
*****      -----|----->      ***** =
* Client *      Obs      * Server * =
*****      <-----|-----      ***** =
<=====

```

(b) observer-server half-RTPL

Figure 6: Half Round-trip packet loss (both direction)

```

=====>
=====> =
*****      ---|-->      ---|-->      ***** = =
* Client *      Obs      Obs      * Server * = =
*****      <--|---      <--|---      ***** = =
<=====

```

(a) observer-server RTPL components (half-RTPLs)

```

=====>
*****      ---|-->      ---|-->      *****
* Client *      Obs      Obs      * Server *
*****      <--|---      <--|---      *****
<=====

```

(b) the intra-domain RTPL resulting from the subtraction of the above RTPL components

Figure 7: Intra-domain Round-trip packet loss (observer-server)

#### 4.1.2. Setting the Round Trip Loss Bit on Outgoing Packets

The round Trip loss signal requires a working Spin-bit signal to separate trains of marked packets (packets with T bit set to 1). A "pause" of at least one empty spin-bit period between each phase of the algorithm serves as such separator for the on-path observer.



The client is in charge of launching trains of marked packets and does so according to the algorithm:

1. Generation Phase. The client starts generating marked packets for two consecutive spin-bit periods; it maintains a "generation token" count that is reset to zero at the beginning of the algorithm phase and is incremented every time a packet arrives. When the client transmits a packet and a "generation token" is available, the client marks the packet and retires a "generation token". If no token is available, the outgoing packet is transmitted unmarked. At the end of the first spin-bit period spent in generation, the reflection counter is unlocked to start counting incoming marked packets that will be reflected later;
2. Pause Phase. When the generation is completed, the client pauses till it has observed one entire spin bit period with no marked packets. That spin bit period is used by the observer as a separator between generated and reflected packets. During this marking pause, all the outgoing packets are transmitted with T bit set to 0. The reflection counter is still incremented every time a marked packet arrives;
3. Reflection Phase. The client starts transmitting marked packets, decrementing the reflection counter for each transmitted marked packet until the reflection counter reached zero. The "generation token" method from the generation phase is used during this phase as well. At the end of the first spin-period spent in reflection, the reflection counter is locked to avoid incoming reflected packets incrementing it;
4. Pause Phase 2. The pause phase is repeated after the reflection phase and serves as a separator between the reflected packet train and a new packet train.

The generation token counter should be capped to limit the effects of a subsequent sudden reduction in the other endpoint's packet rate that could prevent that endpoint from reflecting collected packets. The most conservative cap value is 1.

A server maintains a "marking counter" that starts at zero and is incremented every time a marked packet arrives. When the server transmits a packet and the "marking counter" is positive, the server marks the packet and decrements the "marking counter". If the "marking counter" is zero, the outgoing packet is transmitted unmarked.

#### 4.1.3. Observer's Logic for Round Trip Loss Signal

The on-path observer counts marked packets and separates different trains by detecting spin-bit periods (at least one) with no marked packets. The Round Trip Packet Loss (RTPL) is the difference between the size of the Generation train and the Reflection train.

In the following example, packets are represented by two bits (first one is the spin bit, second one is the loss bit):




Figure 8: Round Trip Loss signal example

Note that 5 marked packets have been generated of which 4 have been reflected.

#### 4.1.4. Loss Coverage and Signal Timing

A cycle of the round Trip loss signaling algorithm contains 2 RTTs of Generation phase, 2 RTTs of Reflection phase, and two Pause phases at least 1 RTT in duration each. Hence, the loss signal is delayed by about 6 RTTs since the loss events.

The observer can only detect loss of marked packets that occurs after its initial observation of the Generation phase and before its subsequent observation of the Reflection phase. Hence, if the loss occurs on the path that sends packets at a lower rate (typically ACKs in such asymmetric scenarios),  $2/6$  ( $1/3$ ) of the packets will be sampled for loss detection.

If the loss occurs on the path that sends packets at a higher rate,  $\text{lowPacketRate}/(3*\text{highPacketRate})$  of the packets will be sampled for loss detection. For protocols that use ACKs, the portion of packets sampled for loss in the higher rate direction during unidirectional data transfer is  $1/(3*\text{packetsPerAck})$ , where the value of `packetsPerAck` can vary by protocol, by implementation, and by network conditions.

#### 4.2. Q Bit - Square Bit

The sQuare bit (Q bit) takes its name from the square wave generated by its signal. Every outgoing packet contains the Q bit value, which is initialized to the 0 and inverted after sending N packets (a

square Block or simply Q Block). Hence, Q Period is  $2 \cdot N$ . The Q bit represents "packet color" as defined by [\[AltMark\]](#).

Observation points can estimate upstream losses by watching a single direction of the traffic flow and counting the number of packets in each observed Q Block, as described in [Section 4.2.2](#).

#### 4.2.1. Q Block Length Selection

The length of the block must be known to the on-path network probes. There are two alternatives to selecting the Q Block length. The first one requires that the length is known a priori and therefore set within the protocol specifications that implements the marking mechanism. The second requires the sender to select it.

In this latter scenario, the sender is expected to choose N (Q Block length) based on the expected amount of loss and reordering on the path. The choice of N strikes a compromise - the observation could become too unreliable in case of packet reordering and/or severe loss if N is too small, while short flows may not yield a useful upstream loss measurement if N is too large (see [Section 4.2.2](#)).

The value of N should be at least 64 and be a power of 2. This requirement allows an Observer to infer the Q Block length by observing one period of the square signal. It also allows the Observer to identify flows that set the loss bits to arbitrary values (see [Section 7](#)).

If the sender does not have sufficient information to make an informed decision about Q Block length, the sender should use  $N=64$ , since this value has been extensively tried in large-scale field tests and yielded good results. Alternatively, the sender may also choose a random power-of-2 N for each flow, increasing the chances of using a Q Block length that gives the best signal for some flows.

The sender must keep the value of N constant for a given flow.

#### 4.2.2. Upstream Loss

Blocks of N (Q Block length) consecutive packets are sent with the same value of the Q bit, followed by another block of N packets with an inverted value of the Q bit. Hence, knowing the value of N, an on-path observer can estimate the amount of upstream loss after observing at least N packets. The upstream loss rate (uloss) is one minus the average number of packets in a block of packets with the same Q value (p) divided by N ( $uloss = 1 - avg(p)/N$ ).

The observer needs to be able to tolerate packet reordering that can blur the edges of the square signal, as explained in [Section 4.2.3](#).

```

=====>
*****      -----Obs----->      *****
* Client *                               * Server *
*****      <-----          *****

```

(a) in client-server channel (uloss\_up)

```

*****      ----->      *****
* Client *                               * Server *
*****      <----Obs-----      *****
<=====

```

(b) in server-client channel (uloss\_down)

Figure 9: Upstream loss

#### 4.2.3. Identifying Q Block Boundaries

Packet reordering can produce spurious edges in the square signal. To address this, the observer should look for packets with the current Q bit value up to X packets past the first packet with a reverse Q bit value. The value of X, a "Marking Block Threshold", should be less than  $N/2$ .

The choice of X represents a trade-off between resiliency to reordering and resiliency to loss. A very large Marking Block Threshold will be able to reconstruct Q Blocks despite a significant amount of reordring, but it may erroneously coalesce packets from multiple Q Blocks into fewer Q Blocks, if loss exceeds 50% for some Q Blocks.

#### 4.3. L Bit - Loss Event Bit

The Loss Event bit uses an Unreported Loss counter maintained by the protocol that implements the marking mechanism. To use the Loss Event bit, the protocol must allow the sender to identify lost packets. This is true of protocols such as QUIC, partially true for TCP and SCTP (losses of pure ACKs are not detected) and is not true of protocols such as UDP and IP/IPv6.

The Unreported Loss counter is initialized to 0, and L bit of every outgoing packet indicates whether the Unreported Loss counter is positive (L=1 if the counter is positive, and L=0 otherwise).

The value of the Unreported Loss counter is decremented every time a packet with L=1 is sent.

The value of the Unreported Loss counter is incremented for every packet that the protocol declares lost, using whatever loss detection machinery the protocol employs. If the protocol is able to

rescind the loss determination later, a positive Unreported Loss counter may be decremented due to the rescission, but it should NOT become negative due to the rescission.

This loss signaling is similar to loss signaling in [\[ConEx\]](#), except the Loss Event bit is reporting the exact number of lost packets, whereas Echo Loss bit in [\[ConEx\]](#) is reporting an approximate number of lost bytes.

For protocols, such as TCP ([\[TCP\]](#)), that allow network devices to change data segmentation, it is possible that only a part of the packet is lost. In these cases, the sender must increment Unreported Loss counter by the fraction of the packet data lost (so Unreported Loss counter may become negative when a packet with L=1 is sent after a partial packet has been lost).

Observation points can estimate the end-to-end loss, as determined by the upstream endpoint, by counting packets in this direction with the L bit equal to 1, as described in [Section 4.3.1](#).

#### **4.3.1. End-To-End Loss**

The Loss Event bit allows an observer to estimate the end-to-end loss rate by counting packets with L bit value of 0 and 1 for a given flow. The end-to-end loss rate is the fraction of packets with L=1.

The assumption here is that upstream loss affects packets with L=0 and L=1 equally. If some loss is caused by tail-drop in a network device, this may be a simplification. If the sender's congestion controller reduces the packet send rate after loss, there may be a sufficient delay before sending packets with L=1 that they have a greater chance of arriving at the observer.

#### **4.3.2. Loss Profile Characterization**

In addition to measuring the end-to-end loss rate, the Loss Event bit allows an observer to characterize loss profile, since the distribution of observed packets with L bit set to 1 roughly corresponds to the distribution of packets lost between 1 RTT and 1 RTO before (see [Section 4.4.1](#)). Hence, observing random single instances of L bit set to 1 indicates random single packet loss, while observing blocks of packets with L bit set to 1 indicates loss affecting entire blocks of packets.

#### 4.4. L+Q Bits - Upstream, Downstream, and End-to-End Loss Measurements

Combining L and Q bits allows a passive observer watching a single direction of traffic to accurately measure:

- \*upstream loss: sender-to-observer loss (see [Section 4.2.2](#))
- \*downstream loss: observer-to-receiver loss (see [Section 4.4.1.1](#))
- \*end-to-end loss: sender-to-receiver loss on the observed path (see [Section 4.3.1](#)) with loss profile characterization (see [Section 4.3.2](#))

##### 4.4.1. Correlating End-to-End and Upstream Loss

Upstream loss is calculated by observing packets that did not suffer the upstream loss ([Section 4.2.2](#)). End-to-end loss, however, is calculated by observing subsequent packets after the sender's protocol detected the loss. Hence, end-to-end loss is generally observed with a delay of between 1 RTT (loss declared due to multiple duplicate acknowledgments) and 1 RTO (loss declared due to a timeout) relative to the upstream loss.

The flow RTT can sometimes be estimated by timing protocol handshake messages. This RTT estimate can be greatly improved by observing a dedicated protocol mechanism for conveying RTT information, such as the Spin bit (see [Section 3.1](#)) or Delay bit (see [Section 3.2](#)).

Whenever the observer needs to perform a computation that uses both upstream and end-to-end loss rate measurements, it should use upstream loss rate leading the end-to-end loss rate by approximately 1 RTT. If the observer is unable to estimate RTT of the flow, it should accumulate loss measurements over time periods of at least 4 times the typical RTT for the observed flows.

If the calculated upstream loss rate exceeds the end-to-end loss rate calculated in [Section 4.3.1](#), then either the Q Period is too short for the amount of packet reordering or there is observer loss, described in [Section 4.4.1.2](#). If this happens, the observer should adjust the calculated upstream loss rate to match end-to-end loss rate, unless the following applies.

In case of a protocol like TCP and SCTP that does not track losses of pure ACK packets, observing a direction of traffic dominated by pure ACK packets could result in measured upstream loss that is higher than measured end-to-end loss, if said pure ACK packets are lost upstream. Hence, if the measurement is applied to such protocols, and the observer can confirm that pure ACK packets dominate the observed traffic direction, the observer should adjust the calculated end-to-end loss rate to match upstream loss rate.

#### 4.4.1.1. Downstream Loss

Because downstream loss affects only those packets that did not suffer upstream loss, the end-to-end loss rate (eloss) relates to the upstream loss rate (uloss) and downstream loss rate (dloss) as  $(1-\text{uloss})(1-\text{dloss})=1-\text{eloss}$ . Hence,  $\text{dloss}=(\text{eloss}-\text{uloss})/(1-\text{uloss})$ .

#### 4.4.1.2. Observer Loss

A typical deployment of a passive observation system includes a network tap device that mirrors network packets of interest to a device that performs analysis and measurement on the mirrored packets. The observer loss is the loss that occurs on the mirror path.

Observer loss affects upstream loss rate measurement, since it causes the observer to account for fewer packets in a block of identical Q bit values (see [Section 4.2.2](#)). The end-to-end loss rate measurement, however, is unaffected by the observer loss, since it is a measurement of the fraction of packets with the L bit value of 1, and the observer loss would affect all packets equally (see [Section 4.3.1](#)).

The need to adjust the upstream loss rate down to match end-to-end loss rate as described in [Section 4.4.1](#) is an indication of the observer loss, whose magnitude is between the amount of such adjustment and the entirety of the upstream loss measured in [Section 4.2.2](#). Alternatively, a high apparent upstream loss rate could be an indication of significant packet reordering, possibly due to packets belonging to a single flow being multiplexed over several upstream paths with different latency characteristics.

#### 4.5. R Bit - Reflection Square Bit

R bit requires a deployment alongside Q bit. Unlike the square signal for which packets are transmitted into blocks of fixed size, the Reflection square signal (being an alternate marking signal too) produces blocks of packets whose size varies according to these rules:

- \*when the transmission of a new block starts, its size is set equal to the size of the last Q Block whose reception has been completed;

- \*if, before transmission of the block is terminated, the reception of at least one further Q Block is completed, the size of the block is updated to the average size of the further received Q Blocks. Implementation details follow.

The Reflection square value is initialized to 0 and is applied to the R-bit of every outgoing packet. The Reflection square value is toggled for the first time when the completion of a Q Block is detected in the incoming square signal (produced by the opposite node using the Q-bit). When this happens, the number of packets ( $p$ ), detected within this first Q Block, is used to generate a reflection square signal which toggles every  $M=p$  packets (at first). This new signal produces blocks of  $M$  packets (marked using the R-bit) and each of them is called "Reflection Block" (R Block).

The  $M$  value is then updated every time a completed Q Block in the incoming square signal is received, following this formula:  
 $M = \text{round}(\text{avg}(p))$ .

The parameter  $\text{avg}(p)$  is the average number of packets in a marking period computed considering all the Q Blocks received since the beginning of the current R Block.

To ensure a proper computation of the  $M$  value, endpoints implementing the R bit must identify the boundaries of incoming Q Blocks. The same approach described in {#endmarkingblock} should be used.

Looking at the R-bit, unidirectional observation points have an indication of losses experienced by the entire unobserved channel plus those occurred in the path from the sender up to them.

Since the Q Block is sent in one direction, and the corresponding reflected R Block is sent in the opposite direction, the reflected R signal is transmitted with the packet rate of the slowest direction. Namely, if the observed direction is the slowest, there can be multiple Q Blocks transmitted in the unobserved direction before a complete R Block is transmitted in the observed direction. If the unobserved direction is the slowest, the observed direction can be sending R Blocks of the same size repeatedly before it can update the signal to account for a newly-completed Q Block.

#### **4.5.1. R+Q Bits - Using R and Q Bits for Passive Loss Measurement**

Since both sQuare and Reflection square bits are toggled at most every  $N$  packets (except for the first transition of the R-bit as explained before), an on-path observer can count the number of packets of each marking block and, knowing the value of  $N$ , can estimate the amount of loss experienced by the connection. An observer can calculate different measurements depending on whether it is able to observe a single direction of the traffic or both directions.



Single directional observer:

- \*upstream loss in the observed direction: the loss between the sender and the observation point (see [Section 4.2.2](#))
- \*"three-quarters" connection loss: the loss between the receiver and the sender in the unobserved direction plus the loss between the sender and the observation point in the observed direction
- \*end-to-end loss in the unobserved direction: the loss between the receiver and the sender in the opposite direction

Two directions observer (same metrics seen previously applied to both direction, plus):

- \*client-observer half round-trip loss: the loss between the client and the observation point in both directions
- \*observer-server half round-trip loss: the loss between the observation point and the server in both directions
- \*downstream loss: the loss between the observation point and the receiver (applicable to both directions)

#### **4.5.1.1. Three-Quarters Connection Loss**

Except for the very first block in which there is nothing to reflect (a complete Q Block has not been yet received), packets are continuously R-bit marked into alternate blocks of size lower or equal than N. Knowing the value of N, an on-path observer can estimate the amount of loss occurred in the whole opposite channel plus the loss from the sender up to it in the observation channel. As for the previous metric, the three-quarters connection loss rate (tqloss) is one minus the average number of packets in a block of packets with the same R value (t) divided by N ( $tqloss = 1 - avg(t)/N$ ).

```

=====>
= *****      -----Obs----->      *****
= * Client *                                * Server *
= *****      <-----              *****
<=====

```

(a) in client-server channel (tqloss\_up)

```

=====>
*****      ----->      ***** =
* Client *                                * Server * =
*****      <-----Obs-----      ***** =
<=====

```

(b) in server-client channel (tqloss\_down)

Figure 10: Three-quarters connection loss

The following metrics derive from this last metric and the upstream loss produced by the Q Bit.

#### 4.5.1.2. End-To-End Loss in the Opposite Direction

End-to-end loss in the unobserved direction (eloss\_unobserved) relates to the "three-quarters" connection loss (tqloss) and upstream loss in the observed direction (uloss) as  $(1 - \text{eloss\_unobserved})(1 - \text{uloss}) = 1 - \text{tqloss}$ . Hence,  $\text{eloss\_unobserved} = (\text{tqloss} - \text{uloss}) / (1 - \text{uloss})$ .

```

*****      -----Obs----->      *****
* Client *                                * Server *
*****      <-----              *****
<=====

```

(a) in client-server channel (eloss\_down)

```

=====>
*****      ----->      *****
* Client *                                * Server *
*****      <-----Obs-----      *****

```


(b) in server-client channel (eloss\_up)

Figure 11: End-To-End loss in the opposite direction


#### 4.5.1.3. Half Round-Trip Loss

If the observer is able to observe both directions of traffic, it is able to calculate two "half round-trip" loss measurements - loss

from the observer to the receiver (in a given direction) and then back to the observer in the opposite direction. For both directions, "half round-trip" loss (hrtloss) relates to "three-quarters" connection loss (tqloss\_opposite) measured in the opposite direction and the upstream loss (uloss) measured in the given direction as  $(1 - \text{uloss})(1 - \text{hrtloss}) = 1 - \text{tqloss\_opposite}$ . Hence,  $\text{hrtloss} = (\text{tqloss\_opposite} - \text{uloss}) / (1 - \text{uloss})$ .



(a) client-observer half round-trip loss (hrtloss\_co)




(b) observer-server half round-trip loss (hrtloss\_os)

Figure 12: Half Round-trip loss (both direction)


#### 4.5.1.4. Downstream Loss

If the observer is able to observe both directions of traffic, it is able to calculate two downstream loss measurements using either end-to-end loss and upstream loss, similar to the calculation in [Section 4.4.1.1](#) or using "half round-trip" loss and upstream loss in the opposite direction.

For the latter,  $\text{dloss} = (\text{hrtloss} - \text{uloss\_opposite}) / (1 - \text{uloss\_opposite})$ .



(a) in client-server channel (dloss<sub>up</sub>)



(b) in server-client channel (dloss<sub>down</sub>)

Figure 13: Downstream loss

#### 4.5.2. Enhancement of R Block Length Computation

The use of the rounding function used in the M computation introduces errors that can be minimized by storing the rounding applied each time M is computed, and using it during the computation of the M value in the following R Block.

This can be achieved introducing the new  $r_{avg}$  parameter in the computation of M. The new formula is  $M_r = \text{avg}(p) + r_{avg}$ ;  $M = \text{round}(M_r)$ ;  $r_{avg} = M_r - M$  where the initial value of  $r_{avg}$  is equal to 0.

#### 4.5.3. Improved Resilience to Packet Reordering

When a protocol implementing the marking mechanism is able to detect when packets are received out of order, it can improve resilience to packet reordering beyond what is possible using methods described in [Section 4.2.3](#).

This can be achieved by updating the size of the current R Block while this is being transmitted. The reflection block size is then updated every time an incoming reordered packet of the previous Q Block is detected. This can be done if and only if the transmission of the current reflection block is in progress and no packets of the following Q Block have been received.

### 5. Summary of Delay and Loss Marking Methods

This section summarizes the marking methods described in this draft.

For the Delay measurement, it is possible to use the spin bit alone or combined with the delay bit. A unidirectional or bidirectional observer can be used.

Method	# of bits	Available Delay Metrics		Impairments Resiliency
		UNIDIR	BIDIR	
		Observer	Observer	
S: Spin Bit	1	RTT	x2	lower
			Half RTT	
SD: Spin Bit + Delay Bit	2	RTT	x2	higher
			Half RTT	

x2 Same metric for both directions.

Figure 14: Delay Comparison

For the Loss measurement, each row in the table of [Figure 15](#) represents a loss marking method. For each method the table specifies the number of bits required in the header, the available metrics using an unidirectional or bidirectional observer, applicable protocols, measurement fidelity and delay.

Method	B	Available		P	Measurement Aspects	
	i	Loss Metrics		r		
	t	UNIDIR	BIDIR	t	Fidelity	Delay
	s	Observer	Observer	o		
T: Round Trip	\$	RT	x2		Rate by	~6 RTT
Loss Bit	1		Half RT	*	sampling	
					1/3 to 1/(3*ppa) of	
					pkts over 2 RTT	
Q: Square Bit	1	Upstream	x2	*	Rate over	N pkts
					N pkts	(e.g. 64)
					(e.g. 64)	
L: Loss Event	1	E2E	x2	#	Loss shape	Min: RTT
Bit					(and rate)	Max: RTT
QL: Square +	2	Upstream	x2		-> see Q	Up: see Q
Loss Ev.		Downstream	x2	#	-> see Q L	Others:
Bits		E2E	x2		-> see L	see L
QR: Square +	2	Upstream	x2		Rate over	Up: see Q
Ref. Sq.		3/4 RT	x2		N*ppa pkts	Others:
Bits		!E2E	E2E	*	(see Q bit	N*ppa pk
			Downstream		for N)	(see Q
			Half RT			for N)

\* All protocols

# Protocols employing loss detection (w/ or w/o pure ACK loss detection)

\$ Require a working spin bit

! Metric relative to the opposite channel

x2 Same metric for both directions

ppa Packets-Per-Ack

Q|L See Q if Upstream loss is significant; L otherwise

Figure 15: Loss Comparison

## 6. ECN-Echo Event Bit

While the primary focus of the draft is on exposing packet loss and delay, modern networks can report congestion before they are forced to drop packets, as described in [ECN]. When transport protocols keep ECN-Echo feedback under encryption, this signal cannot be observed by the network operators. When tasked with diagnosing network performance problems, knowledge of a congestion downstream of an observation point can be instrumental.

If downstream congestion information is desired, this information can be signaled with an additional bit.

\*E: The "ECN-Echo Event" bit is set to 0 or 1 according to the Unreported ECN Echo counter, as explained below in [Section 6.1](#).

### **6.1. Setting the ECN-Echo Event Bit on Outgoing Packets**

The Unreported ECN-Echo counter operates identically to Unreported Loss counter ([Section 4.3](#)), except it counts packets delivered by the network with CE markings, according to the ECN-Echo feedback from the receiver.

This ECN-Echo signaling is similar to ECN signaling in [\[ConEx\]](#). ECN-Echo mechanism in QUIC provides the number of packets received with CE marks. For protocols like TCP, the method described in [\[ConEx-TCP\]](#) can be employed. As stated in [\[ConEx-TCP\]](#), such feedback can be further improved using a method described in [\[ACCURATE\]](#).

### **6.2. Using E Bit for Passive ECN-Reported Congestion Measurement**

A network observer can count packets with CE codepoint and determine the upstream CE-marking rate directly.

Observation points can also estimate ECN-reported end-to-end congestion by counting packets in this direction with a E bit equal to 1.

The upstream CE-marking rate and end-to-end ECN-reported congestion can provide information about downstream CE-marking rate. Presence of E bits along with L bits, however, can somewhat confound precise estimates of upstream and downstream CE-markings in case the flow contains packets that are not ECN-capable.

## **7. Protocol Ossification Considerations**

Accurate loss and delay information is not critical to the operation of any protocol, though its presence for a sufficient number of flows is important for the operation of networks.

The delay and loss bits are amenable to "greasing" described in [\[RFC8701\]](#), if the protocol designers are not ready to dedicate (and ossify) bits used for loss reporting to this function. The greasing could be accomplished similarly to the Latency Spin bit greasing in [\[QUIC-TRANSPORT\]](#). Namely, implementations could decide that a fraction of flows should not encode loss and delay information and, instead, the bits would be set to arbitrary values. The observers would need to be ready to ignore flows with delay and loss information more resembling noise than the expected signal.

## 8. Examples of Application

### 8.1. QUIC

The binding of the delay bit signal to QUIC is partially described in [\[QUIC-TRANSPORT\]](#), which adds the spin bit to the first byte of the short packet header, leaving two reserved bits for future experiments.

To implement the additional signals discussed in this document, the first byte of the short packet header can be modified as follows:

\*the delay bit (D) can be placed in the first reserved bit (i.e. the fourth most significant bit  $0x10$ ) while the loss bit (L) in the second reserved bit (i.e. the fifth most significant bit  $0x08$ ); the proposed scheme is:

```
 0 1 2 3 4 5 6 7
+---+---+---+---+
|0|1|S|D|L|K|P|P|
+---+---+---+---+
```

Figure 16: Scheme 1

\*alternatively, a two bits loss signal (QL or QR) can be placed in both reserved bits; the proposed schemes, in this case, are:

```
 0 1 2 3 4 5 6 7
+---+---+---+---+
|0|1|S|Q|L|K|P|P|
+---+---+---+---+
```

Figure 17: Scheme 2A

```
 0 1 2 3 4 5 6 7
+---+---+---+---+
|0|1|S|Q|R|K|P|P|
+---+---+---+---+
```

Figure 18: Scheme 2B

### 8.2. TCP

The signals can be added to TCP by defining bit 4 of byte 13 of the TCP header to carry the spin bit, and eventually bits 5 and 6 to carry additional information, like the delay bit and the round-trip loss bit, or a two bits loss signal (QL or QR).



## 9. Security Considerations

Passive loss and delay observations have been a part of the network operations for a long time, so exposing loss and delay information to the network does not add new security concerns for protocols that are currently observable.

In the absence of packet loss, Q and R bits signals do not provide any information that cannot be observed by simply counting packets transiting a network path. In the presence of packet loss, Q and R bits will disclose the loss, but this is information about the environment and not the endpoint state. The L bit signal discloses internal state of the protocol's loss detection machinery, but this state can often be gleamed by timing packets and observing congestion controller response.

Hence, loss bits do not provide a viable new mechanism to attack data integrity and secrecy.

### 9.1. Optimistic ACK Attack

A defense against an Optimistic ACK Attack, described in [[QUIC-TRANSPORT](#)], involves a sender randomly skipping packet numbers to detect a receiver acknowledging packet numbers that have never been received. The Q bit signal may inform the attacker which packet numbers were skipped on purpose and which had been actually lost (and are, therefore, safe for the attacker to acknowledge). To use the Q bit for this purpose, the attacker must first receive at least an entire Q Block of packets, which renders the attack ineffective against a delay-sensitive congestion controller.

A protocol that is more susceptible to an Optimistic ACK Attack with the loss signal provided by Q bit and uses a loss-based congestion controller, should shorten the current Q Block by the number of skipped packets numbers. For example, skipping a single packet number will invert the square signal one outgoing packet sooner.

Similar considerations apply to the R Bit, although a shortened R Block along with a matching skip in packet numbers does not necessarily imply a lost packet, since it could be due to a lost packet on the reverse path along with a deliberately skipped packet by the sender.

## 10. Privacy Considerations

To minimize unintentional exposure of information, loss bits provide an explicit loss signal - a preferred way to share information per [[RFC8558](#)].

New protocols commonly have specific privacy goals, and loss reporting must ensure that loss information does not compromise those privacy goals. For example, [QUIC-TRANSPORT] allows changing Connection IDs in the middle of a connection to reduce the likelihood of a passive observer linking old and new sub-flows to the same device. A QUIC implementation would need to reset all counters when it changes the destination (IP address or UDP port) or the Connection ID used for outgoing packets. It would also need to avoid incrementing Unreported Loss counter for loss of packets sent to a different destination or with a different Connection ID.

## 11. IANA Considerations

This document makes no request of IANA.

## 12. Change Log

TBD

## 13. Contributors

The following people provided valuable contributions to this document:

\*Marcus Ihlar, Ericsson, marcus.ihlar@ericsson.com

\*Jari Arkko, Ericsson, jari.arkko@ericsson.com

\*Emile Stephan, Orange, emile.stephan@orange.com

## 14. Acknowledgements

TBD

## 15. References

### 15.1. Normative References

- [ConEx] Mathis, M. and B. Briscoe, "Congestion Exposure (ConEx) Concepts, Abstract Mechanism, and Requirements", RFC 7713, DOI 10.17487/RFC7713, December 2015, <<https://www.rfc-editor.org/info/rfc7713>>.
- [ConEx-TCP] Kuehlewind, M., Ed. and R. Scheffenegger, "TCP Modifications for Congestion Exposure (ConEx)", RFC 7786, DOI 10.17487/RFC7786, May 2016, <<https://www.rfc-editor.org/info/rfc7786>>.
- [ECN] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC

3168, DOI 10.17487/RFC3168, September 2001, <<https://www.rfc-editor.org/info/rfc3168>>.

**[IP]** Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.

**[IPM-Methods]** Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.

**[IPv6]** Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

**[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>>.

**[RFC8558]** Hardie, T., Ed., "Transport Protocol Path Signals", RFC 8558, DOI 10.17487/RFC8558, April 2019, <<https://www.rfc-editor.org/info/rfc8558>>.

**[TCP]** Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/info/rfc793>>.

## 15.2. Informative References

**[ACCURATE]** Briscoe, B., Kuehlewind, M., and R. Scheffenegger, "More Accurate ECN Feedback in TCP", Work in Progress, Internet-Draft, draft-ietf-tcpm-accurate-ecn-12, 28 October 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-tcpm-accurate-ecn-12.txt>>.

**[AltMark]** Fioccola, G., Ed., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, "Alternate-Marking Method for Passive and Hybrid Performance Monitoring", RFC 8321, DOI 10.17487/RFC8321, January 2018, <<https://www.rfc-editor.org/info/rfc8321>>.

**[ANRW19-PM-QUIC]** Bulgarella, F., Cociglio, M., Fioccola, G., Marchetto, G., and R. Sisto, "Performance measurements of QUIC communications", DOI 10.1145/3340301.3341127, Proceedings of the Applied Networking Research Workshop, July 2019, <<https://doi.org/10.1145/3340301.3341127>>.

**[I-D.trammell-ippm-spin]**

Trammell, B., "An Explicit Transport-Layer Signal for Hybrid RTT Measurement", Work in Progress, Internet-Draft, draft-trammell-ippm-spin-00, 9 January 2019, <<http://www.ietf.org/internet-drafts/draft-trammell-ippm-spin-00.txt>>.

**[I-D.trammell-tsvwg-spin]** Trammell, B., "A Transport-Independent Explicit Signal for Hybrid RTT Measurement", Work in Progress, Internet-Draft, draft-trammell-tsvwg-spin-00, 2 July 2018, <<http://www.ietf.org/internet-drafts/draft-trammell-tsvwg-spin-00.txt>>.

**[IPv6AltMark]** Fioccola, G., Zhou, T., Cociglio, M., Qin, F., and R. Pang, "IPv6 Application of the Alternate Marking Method", Work in Progress, Internet-Draft, draft-ietf-6man-ipv6-alt-mark-02, 13 October 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-6man-ipv6-alt-mark-02.txt>>.

**[QUIC-TRANSPORT]** Iyengar, J. and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport", Work in Progress, Internet-Draft, draft-ietf-quic-transport-32, 20 October 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-quic-transport-32.txt>>.

**[RFC8517]** Dolson, D., Ed., Snellman, J., Boucadair, M., Ed., and C. Jacquenet, "An Inventory of Transport-Centric Functions Provided by Middleboxes: An Operator Perspective", RFC 8517, DOI 10.17487/RFC8517, February 2019, <<https://www.rfc-editor.org/info/rfc8517>>.

**[RFC8701]** Benjamin, D., "Applying Generate Random Extensions And Sustain Extensibility (GREASE) to TLS Extensibility", RFC 8701, DOI 10.17487/RFC8701, January 2020, <<https://www.rfc-editor.org/info/rfc8701>>.

**[SPIN-BIT]** Trammell, B., Vaere, P., Even, R., Fioccola, G., Fossati, T., Ihlar, M., Morton, A., and S. Emile, "Adding Explicit Passive Measurability of Two-Way Latency to the QUIC Transport Protocol", Work in Progress, Internet-Draft, draft-trammell-quic-spin-03, 14 May 2018, <<http://www.ietf.org/internet-drafts/draft-trammell-quic-spin-03.txt>>.

**[TRANSPORT-ENCRYPT]**

Fairhurst, G. and C. Perkins, "Considerations around Transport Header Confidentiality, Network Operations, and the Evolution of Internet Transport Protocols", Work in Progress, Internet-Draft, draft-ietf-tsvwg-transport-encrypt-17, 8 September 2020, <<http://www.ietf.org/>

[internet-drafts/draft-ietf-tsvwg-transport-encrypt-17.txt](http://www.ietf.org/internet-drafts/draft-ietf-tsvwg-transport-encrypt-17.txt)>.

**[UDP-OPTIONS]** Touch, J., "Transport Options for UDP", Work in Progress, Internet-Draft, draft-ietf-tsvwg-udp-options-08, 12 September 2019, <<http://www.ietf.org/internet-drafts/draft-ietf-tsvwg-udp-options-08.txt>>.

**[UDP-SURPLUS]** Herbert, T., "UDP Surplus Header", Work in Progress, Internet-Draft, draft-herbert-udp-space-hdr-01, 8 July 2019, <<http://www.ietf.org/internet-drafts/draft-herbert-udp-space-hdr-01.txt>>.

#### Authors' Addresses

Mauro Cociglio  
Telecom Italia  
Via Reiss Romoli, 274  
10148 Torino  
Italy

Email: [mauro.cociglio@telecomitalia.it](mailto:mauro.cociglio@telecomitalia.it)

Alexandre Ferrieux  
Orange Labs

Email: [alexandre.ferrieux@orange.com](mailto:alexandre.ferrieux@orange.com)

Giuseppe Fioccola  
Huawei Technologies  
Riesstrasse, 25  
80992 Munich  
Germany

Email: [giuseppe.fioccola@huawei.com](mailto:giuseppe.fioccola@huawei.com)

Igor Lubashev  
Akamai Technologies

Email: [ilubashe@akamai.com](mailto:ilubashe@akamai.com)

Fabio Bulgarella  
Telecom Italia  
Via Reiss Romoli, 274  
10148 Torino  
Italy

Email: [fabio.bulgarella@guest.telecomitalia.it](mailto:fabio.bulgarella@guest.telecomitalia.it)

Isabelle Hamchaoui

Orange Labs

Email: [isabelle.hamchaoui@orange.com](mailto:isabelle.hamchaoui@orange.com)

Massimo Nilo  
Telecom Italia

Email: [massimo.nilo@telecomitalia.it](mailto:massimo.nilo@telecomitalia.it)

Riccardo Sisto  
Politecnico di Torino

Email: [riccardo.sisto@polito.it](mailto:riccardo.sisto@polito.it)

Dmitri Tikhonov  
LiteSpeed Technologies

Email: [dtikhonov@litespeedtech.com](mailto:dtikhonov@litespeedtech.com)