

Internet Engineering Task Force (IETF)
Request for Comments: 7456
Category: Standards Track
ISSN: 2070-1721

T. Mizrahi
Marvell
T. Senevirathne
S. Salam
D. Kumar
Cisco
D. Eastlake 3rd
Huawei
March 2015

Loss and Delay Measurement in
Transparent Interconnection of Lots of Links (TRILL)

Abstract

Performance Monitoring (PM) is a key aspect of Operations, Administration, and Maintenance (OAM). It allows network operators to verify the Service Level Agreement (SLA) provided to customers and to detect network anomalies. This document specifies mechanisms for Loss Measurement and Delay Measurement in Transparent Interconnection of Lots of Links (TRILL) networks.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7456>.

Copyright Notice

Copyright (c) 2015 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 (<http://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	3
2. Conventions Used in this Document	4
2.1. Key Words	4
2.2. Definitions	4
2.3. Abbreviations	5
3. Loss and Delay Measurement in the TRILL Architecture	6
3.1. Performance Monitoring Granularity	6
3.2. One-Way vs. Two-Way Performance Monitoring	6
3.2.1. One-Way Performance Monitoring	7
3.2.2. Two-Way Performance Monitoring	7
3.3. Point-to-Point vs. Point-to-Multipoint PM	8
4. Loss Measurement	8
4.1. One-Way Loss Measurement	8
4.1.1. 1SL Message Transmission	9
4.1.2. 1SL Message Reception	10
4.2. Two-Way Loss Measurement	11
4.2.1. SLM Message Transmission	12
4.2.2. SLM Message Reception	12
4.2.3. SLR Message Reception	13
5. Delay Measurement	14
5.1. One-Way Delay Measurement	14
5.1.1. 1DM Message Transmission	15
5.1.2. 1DM Message Reception	16
5.2. Two-Way Delay Measurement	16
5.2.1. DMM Message Transmission	17
5.2.2. DMM Message Reception	17
5.2.3. DMR Message Reception	18

6. Packet Formats	19
6.1. TRILL OAM Encapsulation	19
6.2. Loss Measurement Packet Formats	21
6.2.1. Counter Format	21
6.2.2. ISL Packet Format	21
6.2.3. SLM Packet Format	22
6.2.4. SLR Packet Format	23
6.3. Delay Measurement Packet Formats	24
6.3.1. Timestamp Format	24
6.3.2. LDM Packet Format	24
6.3.3. DMM Packet Format	25
6.3.4. DMR Packet Format	26
6.4. OpCode Values	27
7. Performance Monitoring Process	28
8. Security Considerations	29
9. References	29
9.1. Normative References	29
9.2. Informative References	30
Acknowledgments	31
Authors' Addresses	32

1. Introduction

TRILL [TRILL] is a protocol for transparent least-cost routing, where Routing Bridges (RBridges) route traffic to their destination based on least cost, using a TRILL encapsulation header with a hop count.

Operations, Administration, and Maintenance [OAM] is a set of tools for detecting, isolating, and reporting connection failures and performance degradation. Performance Monitoring (PM) is a key aspect of OAM. PM allows network operators to detect and debug network anomalies and incorrect behavior. PM consists of two main building blocks: Loss Measurement and Delay Measurement. PM may also include other derived metrics such as Packet Delivery Rate, and Inter-Frame Delay Variation.

The requirements of OAM in TRILL networks are defined in [OAM-REQ], and the TRILL OAM framework is described in [OAM-FRAMEWK]. These two documents also highlight the main requirements in terms of Performance Monitoring.

This document defines protocols for Loss Measurement and for Delay Measurement in TRILL networks. These protocols are based on the Performance Monitoring functionality defined in ITU-T G.8013/Y.1731 [Y.1731-2013].

- o Loss Measurement: the Loss Measurement protocol measures packet loss between two RBridges. The measurement is performed by sending a set of synthetic packets and counting the number of packets transmitted and received during the test. The frame loss is calculated by comparing the numbers of transmitted and received packets. This provides a statistical estimate of the packet loss between the involved RBridges, with a margin of error that can be controlled by varying the number of transmitted synthetic packets. This document does not define procedures for packet loss computation based on counting user data for the reasons given in Section 5.1 of [OAM-FRAMEWK].
- o Delay Measurement: the Delay Measurement protocol measures the packet delay and packet delay variation between two RBridges. The measurement is performed using timestamped OAM messages.

2. Conventions Used in this Document

2.1. Key Words

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 [KEYWORDS].

The requirement level of PM in [OAM-REQ] is 'SHOULD'. Nevertheless, this memo uses the entire range of requirement levels, including 'MUST'; the requirements in this memo are to be read as 'A MEP (Maintenance End Point) that implements TRILL PM MUST/SHOULD/MAY/...'.

2.2. Definitions

- o One-way packet delay (based on [IPPM-1DM]) - the time elapsed from the start of transmission of the first bit of a packet by an RBridge until the reception of the last bit of the packet by the remote RBridge.
- o Two-way packet delay (based on [IPPM-2DM]) - the time elapsed from the start of transmission of the first bit of a packet from the local RBridge, receipt of the packet at the remote RBridge, the transmission of a response packet from the remote RBridge back to the local RBridge, and receipt of the last bit of that response packet by the local RBridge.
- o Packet loss (based on [IPPM-Loss]) - the number of packets sent by a source RBridge and not received by the destination RBridge. In the context of this document, packet loss is measured at a specific probe instance and a specific observation period. As in

[Y.1731-2013], this document distinguishes between near-end and far-end packet loss. Note that this semantic distinction specifies the direction of packet loss but does not affect the nature of the packet loss metric, which is defined in [IPPM-Loss].

- o Far-end packet loss - the number of packets lost on the path from the local RBridge to the remote RBridge in a specific probe instance and a specific observation period.
- o Near-end packet loss - the number of packets lost on the path from the remote RBridge to the local RBridge in a specific probe instance and a specific observation period.

2.3. Abbreviations

LDM	One-way Delay Measurement
LSL	One-way Synthetic Loss Measurement
DMM	Delay Measurement Message
DMR	Delay Measurement Reply
DoS	Denial of Service
FGL	Fine-Grained Label [FGL]
MD	Maintenance Domain
MD-L	Maintenance Domain Level
MEP	Maintenance End Point
MIP	Maintenance Intermediate Point
MP	Maintenance Point
OAM	Operations, Administration, and Maintenance [OAM]
PM	Performance Monitoring
SLM	Synthetic Loss Measurement Message
SLR	Synthetic Loss Measurement Reply
TLV	Type-Length-Value
TRILL	Transparent Interconnection of Lots of Links [TRILL]

3. Loss and Delay Measurement in the TRILL Architecture

As described in [OAM-FRAMEWK], OAM protocols in a TRILL campus operate over two types of Maintenance Points (MPs): Maintenance End Points (MEPs) and Maintenance Intermediate Points (MIPs).




Figure 1: Maintenance Points in a TRILL Campus

Performance Monitoring (PM) allows a MEP to perform Loss and Delay Measurements on any other MEP in the campus. Performance Monitoring is performed in the context of a specific Maintenance Domain (MD).

The PM functionality defined in this document is not applicable to MIPs.

3.1. Performance Monitoring Granularity

As defined in [OAM-FRAMEWK], PM can be applied at three levels of granularity: Network, Service, and Flow.

- o Network-level PM: the PM protocol is run over a dedicated test VLAN or FGL [FGL].
- o Service-level PM: the PM protocol is used to perform measurements of actual user VLANs or FGLs.
- o Flow-level PM: the PM protocol is used to perform measurements on a per-flow basis. A flow, as defined in [OAM-REQ], is a set of packets that share the same path and per-hop behavior (such as priority). As defined in [OAM-FRAMEWK], flow-based monitoring uses a Flow Entropy field that resides at the beginning of the OAM packet header (see Section 6.1) and mimics the forwarding behavior of the monitored flow.

3.2. One-Way vs. Two-Way Performance Monitoring

Paths in a TRILL network are not necessarily symmetric, that is, a packet sent from RB1 to RB2 does not necessarily traverse the same set of RBridges or links as a packet sent from RB2 to RB1. Even within a given flow, packets from RB1 to RB2 do not necessarily traverse the same path as packets from RB2 to RB1.

3.2.1. One-Way Performance Monitoring

In one-way PM, RB1 sends PM messages to RB2, allowing RB2 to monitor the performance on the path from RB1 to RB2.

A MEP that implements TRILL PM SHOULD support one-way Performance Monitoring. A MEP that implements TRILL PM SHOULD support both the PM functionality of the sender, RB1, and the PM functionality of the receiver, RB2.

One-way PM can be applied either proactively or on-demand, although the more typical scenario is the proactive mode, where RB1 and RB2 periodically transmit PM messages to each other, allowing each of them to monitor the performance on the incoming path from the peer MEP.

3.2.2. Two-Way Performance Monitoring

In two-way PM, a sender, RB1, sends PM messages to a reflector, RB2, and RB2 responds to these messages, allowing RB1 to monitor the performance of:

- o The path from RB1 to RB2.
- o The path from RB2 to RB1.
- o The two-way path from RB1 to RB2, and back to RB1.

Note that in some cases it may be interesting for RB1 to monitor only the path from RB1 to RB2. Two-way PM allows the sender, RB1, to monitor the path from RB1 to RB2, as opposed to one-way PM (Section 3.2.1), which allows the receiver, RB2, to monitor this path.

A MEP that implements TRILL PM MUST support two-way PM. A MEP that implements TRILL PM MUST support both the sender and the reflector PM functionality.

As described in Section 3.1, flow-based PM uses the Flow Entropy field as one of the parameters that identify a flow. In two-way PM, the Flow Entropy of the path from RB1 to RB2 is typically different from the Flow Entropy of the path from RB2 to RB1. This document uses the Reflector Entropy TLV [TRILL-FM], which allows the sender to specify the Flow Entropy value to be used in the response message.

Two-way PM can be applied either proactively or on-demand.

3.3. Point-to-Point vs. Point-to-Multipoint PM

PM can be applied either as a point-to-point measurement protocol, or as a point-to-multi-point measurement protocol.

The point-to-point approach measures the performance between two RBridges using unicast PM messages.

In the point-to-multipoint approach, an RBridge RB1 sends PM messages to multiple RBridges using multicast messages. The reflectors (in two-way PM) respond to RB1 using unicast messages. To protect against reply storms, the reflectors MUST send the response messages after a random delay in the range of 0 to 2 seconds. This ensures that the responses are staggered in time and that the initiating RBridge is not overwhelmed with responses. Moreover, an RBridge Scope TLV [TRILL-FM] can be used to limit the set of RBridges from which a response is expected, thus reducing the impact of potential response bursts.

4. Loss Measurement

The Loss Measurement protocol has two modes of operation: one-way Loss Measurement and two-way Loss Measurement.

Note: The terms 'one-way' and 'two-way' Loss Measurement should not be confused with the terms 'single-ended' and 'dual-ended' Loss Measurement used in [Y.1731-2013]. As defined in Section 3.2, the terms 'one-way' and 'two-way' specify whether the protocol monitors performance on one direction or on both directions. The terms 'single-ended' and 'dual-ended', on the other hand, describe whether the protocol is asymmetric or symmetric, respectively.

4.1. One-Way Loss Measurement

One-way Loss Measurement measures the one-way packet loss from one MEP to another. The loss ratio is measured using a set of One-way Synthetic Loss Measurement (1SL) messages. The packet format of the 1SL message is specified in Section 6.2.2. Figure 2 illustrates a one-way Loss Measurement message exchange.




Figure 2: One-Way Loss Measurement

The one-way Loss Measurement procedure uses a set of 1SL messages to measure the packet loss. The figure shows two non-consecutive messages from the set.

The sender maintains a counter of transmitted 1SL messages, and includes the value of this counter, TX, in each 1SL message it transmits. The receiver maintains a counter of received 1SL messages, RX, and can calculate the loss by comparing its counter values to the counter values received in the 1SL messages.

In Figure 2, the subscript 'c' is an abbreviation for current, and 'p' is an abbreviation for previous.

4.1.1.1. 1SL Message Transmission

One-way Loss Measurement can be applied either proactively or on-demand, although as mentioned in Section 3.2.1, it is more likely to be applied proactively.

The term 'on-demand' in the context of one-way Loss Measurement implies that the sender transmits a fixed set of 1SL messages, allowing the receiver to perform the measurement based on this set.

A MEP that supports one-way Loss Measurement MUST support unicast transmission of 1SL messages.

A MEP that supports one-way Loss Measurement MAY support multicast transmission of 1SL messages.

The sender MUST maintain a packet counter for each peer MEP and probe instance (test ID). Every time the sender transmits a 1SL packet, it increments the corresponding counter and then integrates the value of the counter into the Counter TX field of the 1SL packet.

The 1SL message MAY be sent with a variable-size Data TLV, allowing Loss Measurement for various packet sizes.

4.1.2. 1SL Message Reception

The receiver MUST maintain a reception counter for each peer MEP and probe instance (test ID). Upon receiving a 1SL packet, the receiver MUST verify that:

- o The 1SL packet is destined to the current MEP.
- o The packet's MD level matches the MEP's MD level.

If both conditions are satisfied, the receiver increments the corresponding reception counter and records the new value of the counter, RX1.

A MEP that supports one-way Loss Measurement MUST support reception of both unicast and multicast 1SL messages.

The receiver computes the one-way packet loss with respect to a probe instance measurement interval. A probe instance measurement interval includes a sequence of 1SL messages with the same test ID. The one-way packet loss is computed by comparing the counter values TXp and RXp at the beginning of the measurement interval and the counter values TXc and RXc at the end of the measurement interval (see Figure 2):

$$\text{one-way packet loss} = (\text{TXc} - \text{TXp}) - (\text{RXc} - \text{RXp}) \quad (1)$$

The calculation in Equation (1) is based on counter value differences, implying that the sender's counter, TX, and the receiver's counter, RX, are not required to be synchronized with respect to a common initial value.

It is noted that if the sender or receiver resets one of the counters, TX or RX, the calculation in Equation (1) produces a false measurement result. Hence, the sender and receiver SHOULD NOT clear the TX and RX counters during a measurement interval.

When the receiver calculates the packet loss per Equation (1), it MUST perform a wraparound check. If the receiver detects that one of the counters has wrapped around, the receiver adjusts the result of Equation (1) accordingly.

A 1SL receiver MUST support reception of 1SL messages with a Data TLV.

Since synthetic one-way Loss Measurement is performed using 1SL messages, obviously, some 1SL messages may be dropped during a measurement interval. Thus, when the receiver does not receive a 1SL, the receiver cannot perform the calculations in Equation (1) for that specific 1SL message.

4.2. Two-Way Loss Measurement

Two-way Loss Measurement allows a MEP to measure the packet loss on the paths to and from a peer MEP. Two-way Loss Measurement uses a set of Synthetic Loss Measurement Messages (SLMs) to compute the packet loss. Each SLM is answered with a Synthetic Loss Measurement Reply (SLR). The packet formats of the SLM and SLR packets are specified in Sections 6.2.3 and 6.2.4, respectively. Figure 3 illustrates a two-way Loss Measurement message exchange.




Figure 3: Two-Way Loss Measurement

The two-way Loss Measurement procedure uses a set of SLM-SLR handshakes. The figure shows two non-consecutive handshakes from the set.

The sender maintains a counter of transmitted SLM messages and includes the value of this counter, TX, in each transmitted SLM message. The reflector maintains a counter of received SLM messages, TRX. The reflector generates an SLR and incorporates TRX into the SLR packet. The sender maintains a counter of received SLR messages, RX. Upon receiving an SLR message, the sender can calculate the loss by comparing the local counter values to the counter values received in the SLR messages.

The subscript 'c' is an abbreviation for current, and 'p' is an abbreviation for previous.

4.2.1. SLM Message Transmission

Two-way Loss Measurement can be applied either proactively or on-demand.

A MEP that supports two-way Loss Measurement MUST support unicast transmission of SLM messages.

A MEP that supports two-way Loss Measurement MAY support multicast transmission of SLM messages.

The sender MUST maintain a counter of transmitted SLM packets for each peer MEP and probe instance (test ID). Every time the sender transmits an SLM packet, it increments the corresponding counter and then integrates the value of the counter into the Counter TX field of the SLM packet.

A sender MAY include a Reflector Entropy TLV in an SLM message. The Reflector Entropy TLV format is specified in [TRILL-FM].

An SLM message MAY be sent with a Data TLV, allowing Loss Measurement for various packet sizes.

4.2.2. SLM Message Reception

The reflector MUST maintain a reception counter, TRX, for each peer MEP and probe instance (test ID).

Upon receiving an SLM packet, the reflector MUST verify that:

- o The SLM packet is destined to the current MEP.
- o The packet's MD level matches the MEP's MD level.

If both conditions are satisfied, the reflector increments the corresponding packet counter and records the value of the new counter, TRX. The reflector then generates an SLR message that is identical to the received SLM, except for the following modifications:

- o The reflector incorporates TRX into the Counter TRX field of the SLR.
- o The OpCode field in the OAM header is set to the SLR OpCode.
- o The reflector assigns its MEP ID in the Reflector MEP ID field.

- o If the received SLM includes a Reflector Entropy TLV [TRILL-FM], the reflector copies the value of the Flow Entropy from the TLV into the Flow Entropy field of the SLR message. The outgoing SLR message does not include a Reflector Entropy TLV.
- o The TRILL Header and transport header are modified to reflect the source and destination of the SLR packet. The SLR is always a unicast message.

A MEP that supports two-way Loss Measurement MUST support reception of both unicast and multicast SLM messages.

A reflector MUST support reception of SLM packets with a Data TLV. When receiving an SLM with a Data TLV, the reflector includes the unmodified TLV in the SLR.

4.2.3. SLR Message Reception

The sender MUST maintain a reception counter, RX, for each peer MEP and probe instance (test ID).

Upon receiving an SLR message, the sender MUST verify that:

- o The SLR packet is destined to the current MEP.
- o The Sender MEP ID field in the SLR packet matches the current MEP.
- o The packet's MD level matches the MEP's MD level.

If the conditions above are met, the sender increments the corresponding reception counter, and records the new value, RX.

The sender computes the packet loss with respect to a probe instance measurement interval. A probe instance measurement interval includes a sequence of SLM messages and their corresponding SLR messages, all with the same test ID. The packet loss is computed by comparing the counters at the beginning of the measurement interval, denoted with a subscript 'p', and the counters at the end of the measurement interval, denoted with a subscript 'c' (as illustrated in Figure 3).

$$\text{far-end packet loss} = (\text{TXc} - \text{TXp}) - (\text{TRXc} - \text{TRXp}) \quad (2)$$

$$\text{near-end packet loss} = (\text{TRXc} - \text{TRXp}) - (\text{RXc} - \text{RXp}) \quad (3)$$

Note: The total two-way packet loss is the sum of the far-end and near-end packet losses, that is $(\text{TXc} - \text{TXp}) - (\text{RXc} - \text{RXp})$.

The calculations in the two equations above are based on counter value differences, implying that the sender's counters, TX and RX, and the reflector's counter, TRX, are not required to be synchronized with respect to a common initial value.

It is noted that if the sender or reflector resets one of the counters, TX, TRX, or RX, the calculation in Equations (2) and (3) produces a false measurement result. Hence, the sender and reflector SHOULD NOT clear the TX, TRX, and RX counters during a measurement interval.

When the sender calculates the packet loss per Equations (2) and (3), it MUST perform a wraparound check. If the reflector detects that one of the counters has wrapped around, the reflector adjusts the result of Equations (2) and (3) accordingly.

Since synthetic two-way Loss Measurement is performed using SLM and SLR messages, obviously, some SLM and SLR messages may be dropped during a measurement interval. When an SLM or an SLR is dropped, the corresponding two-way handshake (Figure 3) is not completed successfully; thus, the reflector does not perform the calculations in Equations (2) and (3) for that specific message exchange.

A sender MAY choose to monitor only the far-end packet loss, that is, perform the computation in Equation (2), and ignore the computation in Equation (3). Note that, in this case, the sender can run flow-based PM of the path to the peer MEP without using the Reflector Entropy TLV.

5. Delay Measurement

The Delay Measurement protocol has two modes of operation: one-way Delay Measurement and two-way Delay Measurement.

5.1. One-Way Delay Measurement

One-way Delay Measurement is used for computing the one-way packet delay from one MEP to another. The packet format used in one-way Delay Measurement is referred to as 1DM and is specified in Section 6.3.2. The one-way Delay Measurement message exchange is illustrated in Figure 4.




Figure 4: One-Way Delay Measurement

The sender transmits a 1DM message incorporating its time of transmission, T1. The receiver then receives the message at time T2, and calculates the one-way delay as:

$$\text{one-way delay} = T2 - T1 \quad (4)$$

Equation (4) implies that T2 and T1 are measured with respect to a common reference time. Hence, two MEPs running a one-way Delay Measurement protocol MUST be time-synchronized. The method used for synchronizing the clocks associated with the two MEPs is outside the scope of this document.

5.1.1.1. 1DM Message Transmission

1DM packets can be transmitted proactively or on-demand, although, as mentioned in Section 3.2.1, they are typically transmitted proactively.

A MEP that supports one-way Delay Measurement MUST support unicast transmission of 1DM messages.

A MEP that supports one-way Delay Measurement MAY support multicast transmission of 1DM messages.

A 1DM message MAY be sent with a variable size Data TLV, allowing packet Delay Measurement for various packet sizes.

The sender incorporates the 1DM packet's time of transmission into the Timestamp T1 field.

5.1.2. 1DM Message Reception

Upon receiving a 1DM packet, the receiver records its time of reception, T2. The receiver MUST verify two conditions:

- o The 1DM packet is destined to the current MEP.
- o The packet's MD level matches the MEP's MD level.

If both conditions are satisfied, the receiver terminates the packet and calculates the one-way delay as specified in Equation (4).

A MEP that supports one-way Delay Measurement MUST support reception of both unicast and multicast 1DM messages.

A 1DM receiver MUST support reception of 1DM messages with a Data TLV.

When one-way Delay Measurement packets are received periodically, the receiver MAY compute the packet delay variation based on multiple measurements. Note that packet delay variation can be computed even when the two peer MEPs are not time-synchronized.

5.2. Two-Way Delay Measurement

Two-way Delay Measurement uses a two-way handshake for computing the two-way packet delay between two MEPs. The handshake includes two packets: a Delay Measurement Message (DMM) and a Delay Measurement Reply (DMR). The DMM and DMR packet formats are specified in Sections 6.3.3 and 6.3.4, respectively.

The two-way Delay Measurement message exchange is illustrated in Figure 5.




Figure 5: Two-Way Delay Measurement

The sender generates a DMM message incorporating its time of transmission, T1. The reflector receives the DMM message and records its time of reception, T2. The reflector then generates a DMR

message, incorporating T1, T2, and the DMR's transmission time, T3. The sender receives the DMR message at T4, and using the four timestamps, it calculates the two-way packet delay.

5.2.1. DMM Message Transmission

DMM packets can be transmitted periodically or on-demand.

A MEP that supports two-way Delay Measurement **MUST** support unicast transmission of DMM messages.

A MEP that supports two-way Delay Measurement **MAY** support multicast transmission of DMM messages.

A sender **MAY** include a Reflector Entropy TLV in a DMM message. The Reflector Entropy TLV format is specified in [TRILL-FM].

A DMM **MAY** be sent with a variable size Data TLV, allowing packet Delay Measurement for various packet sizes.

The sender incorporates the DMM packet's time of transmission into the Timestamp T1 field.

5.2.2. DMM Message Reception

Upon receiving a DMM packet, the reflector records its time of reception, T2. The reflector **MUST** verify two conditions:

- o The DMM packet is destined to the current MEP.
- o The packet's MD level matches the MEP's MD level.

If both conditions are satisfied, the reflector terminates the packet and generates a DMR packet. The DMR is identical to the received DMM, except for the following modifications:

- o The reflector incorporates T2 into the Timestamp T2 field of the DMR.
- o The reflector incorporates the DMR's transmission time, T3, into the Timestamp T3 field of the DMR.
- o The OpCode field in the OAM header is set to the DMR OpCode.
- o If the received DMM includes a Reflector Entropy TLV [TRILL-FM], the reflector copies the value of the Flow Entropy from the TLV into the Flow Entropy field of the DMR message. The outgoing DMR message does not include a Reflector Entropy TLV.

- o The TRILL Header and transport header are modified to reflect the source and destination of the DMR packet. The DMR is always a unicast message.

A MEP that supports two-way Delay Measurement MUST support reception of both unicast and multicast DMM messages.

A reflector MUST support reception of DMM packets with a Data TLV. When receiving a DMM with a Data TLV, the reflector includes the unmodified TLV in the DMR.

5.2.3. DMR Message Reception

Upon receiving the DMR message, the sender records its time of reception, T4. The sender MUST verify:

- o The DMR packet is destined to the current MEP.
- o The packet's MD level matches the MEP's MD level.

If both conditions above are met, the sender uses the four timestamps to compute the two-way delay:

$$\text{two-way delay} = (T4 - T1) - (T3 - T2) \quad (5)$$

Note that two-way delay can be computed even when the two peer MEPs are not time-synchronized. One-way Delay Measurement, on the other hand, requires the two MEPs to be synchronized.

Two MEPs running a two-way Delay Measurement protocol MAY be time-synchronized. If two-way Delay Measurement is run between two time-synchronized MEPs, the sender MAY compute the one-way delays as follows:

$$\text{one-way delay \{sender->reflector\}} = T2 - T1 \quad (6)$$

$$\text{one-way delay \{reflector->sender\}} = T4 - T3 \quad (7)$$

When two-way Delay Measurement is run periodically, the sender MAY also compute the delay variation based on multiple measurements.

A sender MAY choose to monitor only the sender->reflector delay, that is, perform the computation in Equation (6) and ignore the computations in Equations (5) and (7). Note that in this case, the sender can run flow-based PM of the path to the peer MEP without using the Reflector Entropy TLV.

6. Packet Formats

6.1. TRILL OAM Encapsulation

The TRILL OAM packet format is generally discussed in [OAM-FRAMEWK] and specified in detail in [TRILL-FM]. It is quoted in this document for convenience.

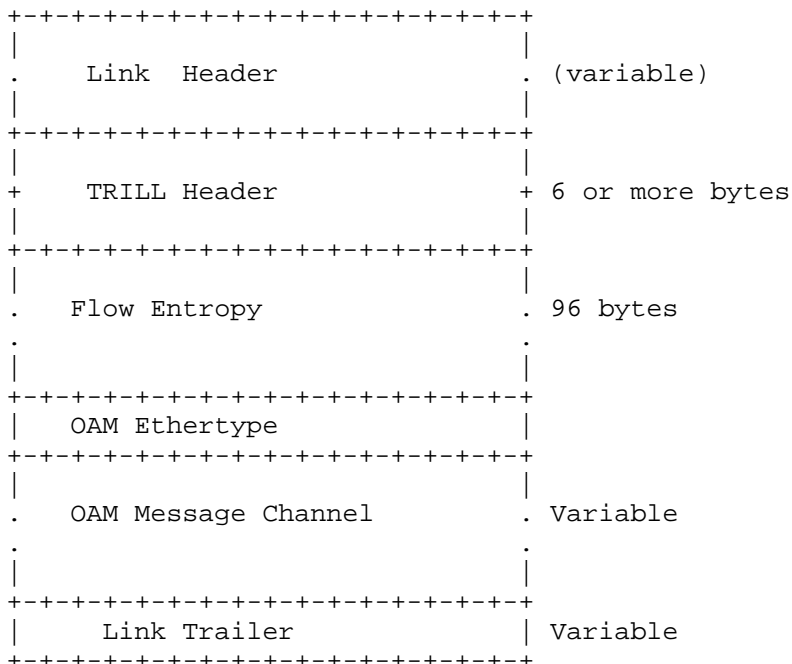


Figure 6: TRILL OAM Encapsulation

The OAM Message Channel used in this document is defined in [TRILL-FM] and has the following structure:




Figure 7: OAM Packet Format

The first four octets of the OAM Message Channel are common to all OpCodes, whereas the rest is OpCode-specific. Below is a brief summary of the fields in the first 4 octets:

- o MD-L: Maintenance Domain Level.
- o Version: indicates the version of this protocol. Always zero in the context of this document.
- o OpCode: Operation Code (8 bits). Specifies the operation performed by the message. Specific packet formats are presented in Sections 6.2 and 6.3 of this document. A list of the PM message OpCodes is provided in Section 6.4.
- o Flags: The definition of flags is OpCode-specific. The value of this field is zero unless otherwise stated.
- o FirstTLVOffset: defines the location of the first TLV, in octets, starting from the end of the FirstTLVOffset field.
- o TLVs: one or more TLV fields. The last TLV field is always an End TLV.

For further details about the OAM packet format, including the format of TLVs, see [TRILL-FM].

6.2. Loss Measurement Packet Formats

6.2.1. Counter Format

Loss Measurement packets use a 32-bit packet counter field. When a counter is incremented beyond its maximal value, 0xFFFFFFFF, it wraps around back to 0.

6.2.2. 1SL Packet Format

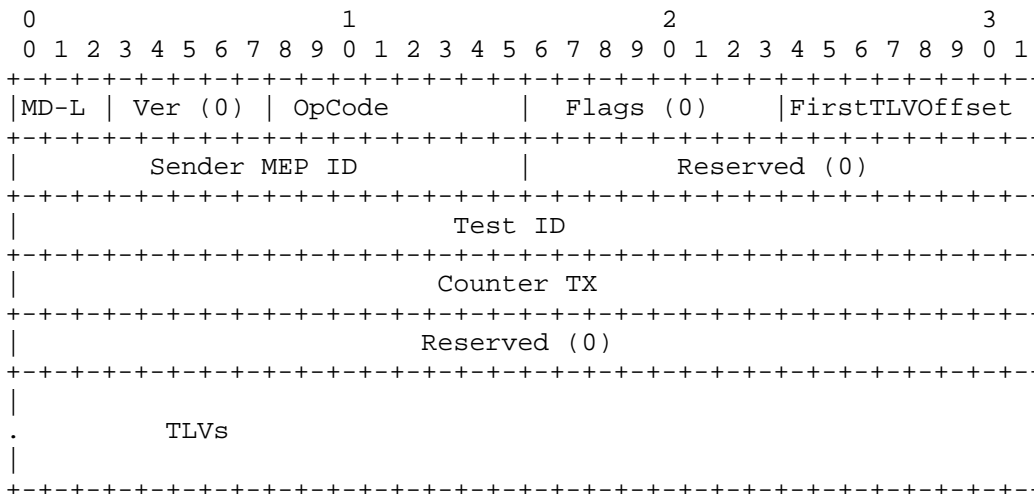


Figure 8: 1SL Packet Format

For fields not listed below, see Section 6.1.

- o OpCode: see Section 6.4.
- o FirstTLVOffset: defines the location of the first TLV, in octets, starting from the end of the FirstTLVOffset field. The value of this field MUST be 16 in 1SL packets.
- o Sender MEP ID: the MEP ID of the MEP that initiated the 1SL.
- o Reserved (0): set to 0 by the sender and ignored by the receiver.
- o Test ID: a 32-bit unique test identifier.
- o Counter TX: the value of the sender’s transmission counter, including this packet, at the time of transmission.

6.2.3. SLM Packet Format

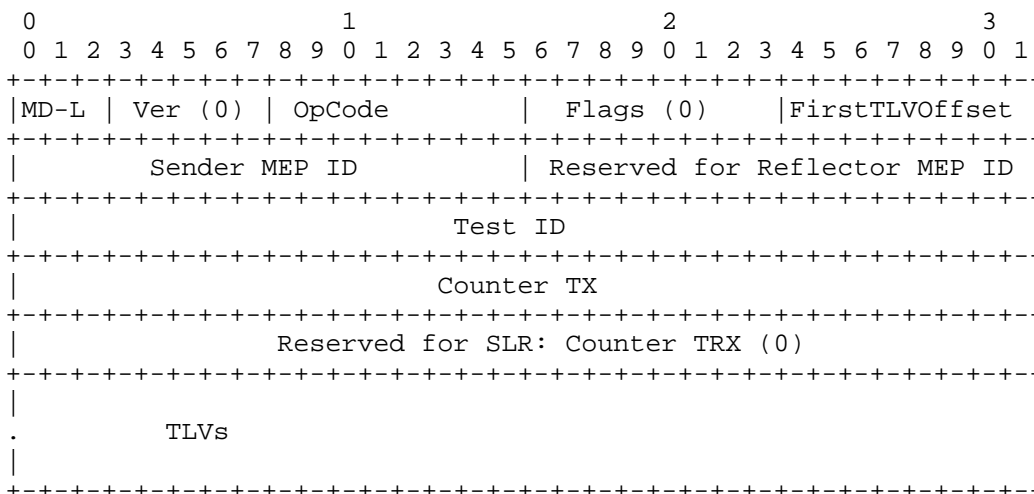


Figure 9: SLM Packet Format

For fields not listed below, see Section 6.1.

- o OpCode: see Section 6.4.
- o FirstTLVOffset: defines the location of the first TLV, in octets, starting from the end of the FirstTLVOffset field. The value of this field MUST be 16 in SLM packets.
- o Sender MEP ID: the MEP ID of the MEP that initiated this packet.
- o Reserved for Reflector MEP ID: this field is reserved for the reflector’s MEP ID, to be added in the SLR.
- o Test ID: a 32-bit unique test identifier.
- o Counter TX: the value of the sender’s transmission counter, including this packet, at the time of transmission.
- o Reserved for SLR: this field is reserved for the SLR corresponding to this packet. The reflector uses this field in the SLR for carrying TRX, the value of its reception counter.

6.2.4. SLR Packet Format

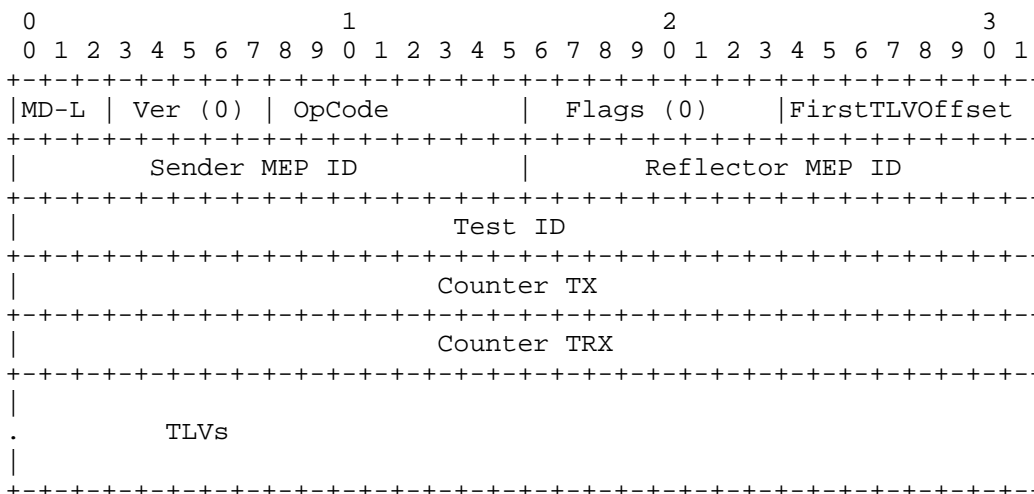


Figure 10: SLR Packet Format

For fields not listed below, see Section 6.1.

- o OpCode: see Section 6.4.
- o FirstTLVOffset: defines the location of the first TLV, in octets, starting from the end of the FirstTLVOffset field. The value of this field MUST be 16 in SLR packets.
- o Sender MEP ID: the MEP ID of the MEP that initiated the SLM that this SLR replies to.
- o Reflector MEP ID: the MEP ID of the MEP that transmits this SLR message.
- o Test ID: a 32-bit unique test identifier, copied from the corresponding SLM message.
- o Counter TX: the value of the sender’s transmission counter at the time of the SLM transmission.
- o Counter TRX: the value of the reflector’s reception counter, including this packet, at the time of reception of the corresponding SLM packet.

6.3. Delay Measurement Packet Formats

6.3.1. Timestamp Format

The timestamps used in Delay Measurement packets are 64 bits long. These timestamps use the 64 least significant bits of the IEEE 1588-2008 (1588v2) Precision Time Protocol timestamp format [IEEE1588v2].

This truncated format consists of a 32-bit seconds field followed by a 32-bit nanoseconds field. This truncated format is also used in IEEE 1588v1 [IEEE1588v1], in [Y.1731-2013], and in [MPLS-LM-DM].

6.3.2. lDM Packet Format

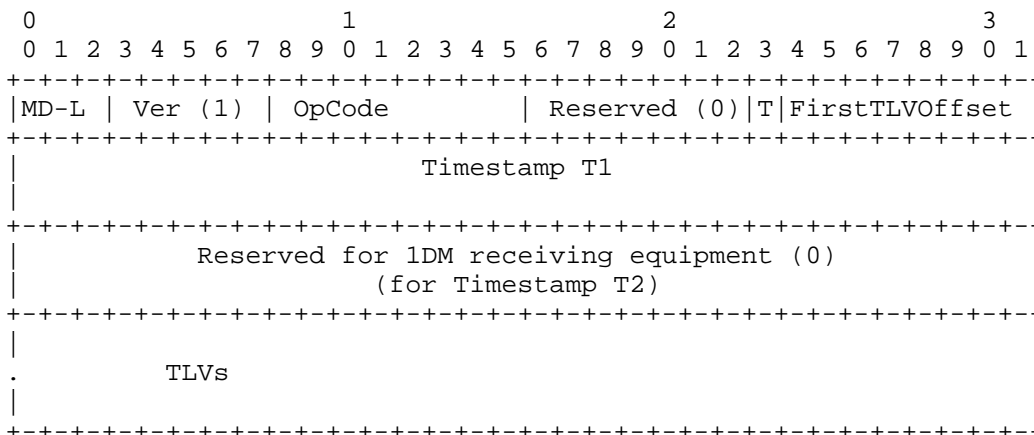


Figure 11: lDM Packet Format

For fields not listed below, see Section 6.1.

- o OpCode: see Section 6.4.
- o Reserved (0): Upper part of Flags field. Set to 0 by the sender and ignored by the receiver.
- o T: Type flag. When this flag is set, it indicates proactive operation; when cleared, it indicates on-demand mode.
- o FirstTLVOffset: defines the location of the first TLV, in octets, starting from the end of the FirstTLVOffset field. The value of this field MUST be 16 in lDM packets.
- o Timestamp T1: specifies the time of transmission of this packet.

- o Reserved for 1DM: this field is reserved for internal usage of the 1DM receiver. The receiver can use this field for carrying T2, the time of reception of this packet.

6.3.3. DMM Packet Format

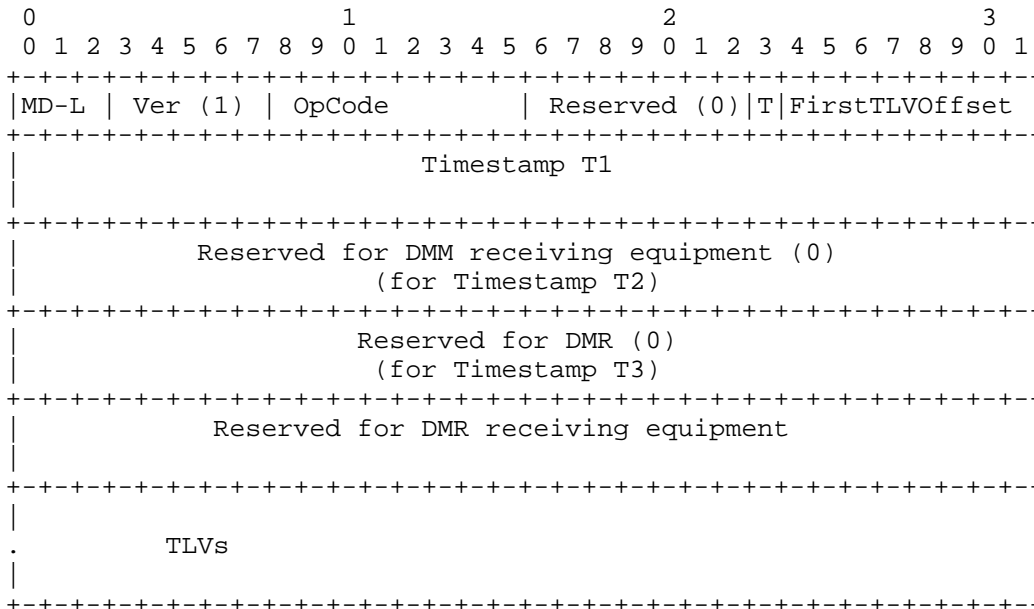


Figure 12: DMM Packet Format

For fields not listed below, see Section 6.1.

- o OpCode: see Section 6.4.
- o Reserved (0): Upper part of Flags field. Set to 0 by the sender and ignored by the receiver.
- o T: Type flag. When this flag is set, it indicates proactive operation; when cleared, it indicates on-demand mode.
- o FirstTLVOffset: defines the location of the first TLV, in octets, starting from the end of the FirstTLVOffset field. The value of this field MUST be 32 in DMM packets.
- o Timestamp T1: specifies the time of transmission of this packet.

- o Reserved for DMM: this field is reserved for internal usage of the MEP that receives the DMM (the reflector). The reflector can use this field for carrying T2, the time of reception of this packet.
- o Reserved for DMR: two timestamp fields are reserved for the DMR message. One timestamp field is reserved for T3, the DMR transmission time, and the other field is reserved for internal usage of the MEP that receives the DMR.

6.3.4. DMR Packet Format

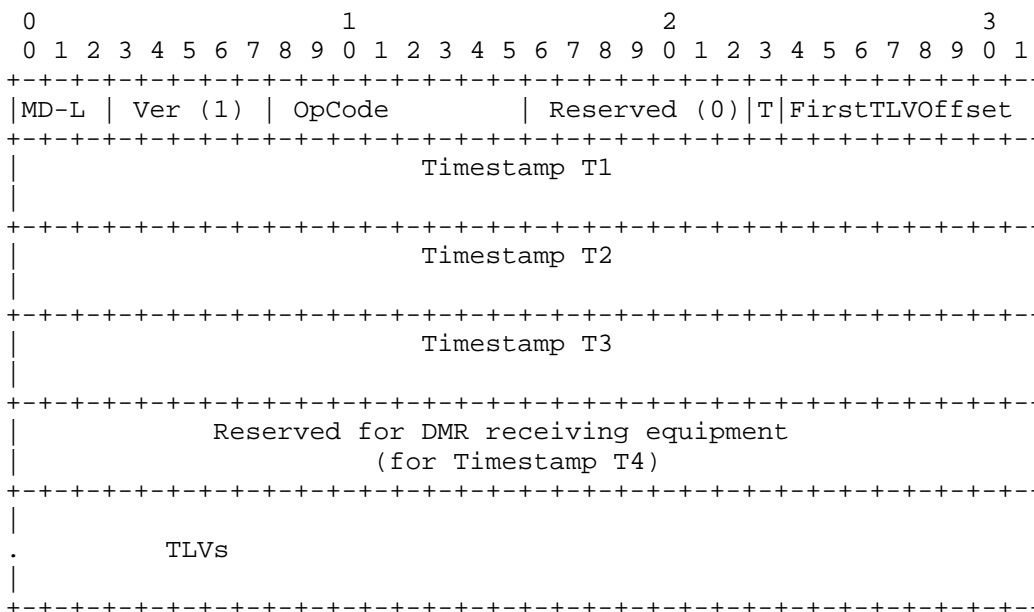


Figure 13: DMR Packet Format

For fields not listed below, see Section 6.1.

- o OpCode: see Section 6.4.
- o Reserved (0): Upper part of Flags field. Set to 0 by the sender and ignored by the receiver.
- o T: Type flag. When this flag is set, it indicates proactive operation; when cleared, it indicates on-demand mode.
- o FirstTLVOffset: defines the location of the first TLV, in octets, starting from the end of the FirstTLVOffset field. The value of this field MUST be 32 in DMR packets.

- o Timestamp T1: specifies the time of transmission of the DMM packet that this DMR replies to.
- o Timestamp T2: specifies the time of reception of the DMM packet that this DMR replies to.
- o Timestamp T3: specifies the time of transmission of this DMR packet.
- o Reserved for DMR: this field is reserved for internal usage of the MEP that receives the DMR (the sender). The sender can use this field for carrying T4, the time of reception of this packet.

6.4. OpCode Values

As the OAM packets specified herein conform to [Y.1731-2013], the same OpCodes are used:

OpCode value	OAM packet type
-----	-----
45	1DM
46	DMR
47	DMM
53	1SL
54	SLR
55	SLM

These OpCodes are from the range of values that has been allocated by IEEE 802.1 [802.1Q] for control by ITU-T.

7. Performance Monitoring Process

The Performance Monitoring process is made up of a number of Performance Monitoring instances, known as PM Sessions. A PM session can be initiated between two MEPs on a specific flow and be defined as either a Loss Measurement session or Delay Measurement session.

The Loss Measurement session can be used to determine the performance metrics Frame Loss Ratio, availability, and resiliency. The Delay Measurement session can be used to determine the performance metrics Frame Delay, Inter-Frame Delay Variation, Frame Delay Range, and Mean Frame Delay.

The PM session is defined by the specific PM function (PM tool) being run and also by the Start Time, Stop Time, Message Period, Measurement Interval, and Repetition Time. These terms are defined as follows:

- o Start Time - the time that the PM session begins.
- o Stop Time - the time that the measurement ends.
- o Message Period - the message transmission frequency (the time between message transmissions).
- o Measurement Interval - the time period over which measurements are gathered and then summarized. The Measurement Interval can align with the PM Session duration, but it doesn't need to. PM messages are only transmitted during a PM Session.
- o Repetition Time - the time between start times of the Measurement Intervals.

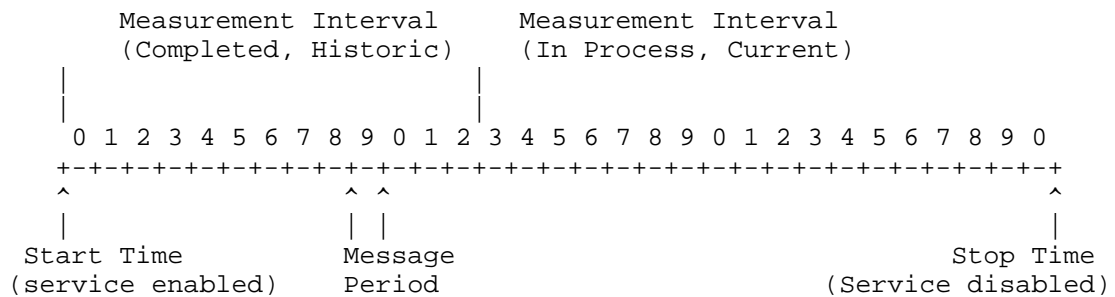


Figure 14: Relationship between Different Timing Parameters

8. Security Considerations

The security considerations of TRILL OAM are discussed in [OAM-REQ], [OAM-FRAMEWK], and [TRILL-FM]. General TRILL security considerations are discussed in [TRILL].

As discussed in [OAM-Over], an attack on a PM protocol can falsely indicate nonexistent performance issues or prevent the detection of actual ones, consequently resulting in DoS (Denial of Service). Furthermore, synthetic PM messages can be used maliciously as a means to implement DoS attacks on RBridges. Another security aspect is network reconnaissance; by passively eavesdropping on PM messages, an attacker can gather information that can be used maliciously to attack the network.

As in [TRILL-FM], TRILL PM OAM messages MAY include the OAM Authentication TLV. It should be noted that an Authentication TLV requires a cryptographic algorithm, which may have performance implications on the RBridges that take part in the protocol; thus, they may, in some cases, affect the measurement results. Based on a system-specific threat assessment, the benefits of the security TLV must be weighed against the potential measurement inaccuracy it may inflict, and based on this trade-off, operators should make a decision on whether or not to use authentication.

9. References

9.1. Normative References

- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [TRILL] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, July 2011, <<http://www.rfc-editor.org/info/rfc6325>>.
- [FGL] Eastlake 3rd, D., Zhang, M., Agarwal, P., Perlman, R., and D. Dutt, "Transparent Interconnection of Lots of Links (TRILL): Fine-Grained Labeling", RFC 7172, May 2014, <<http://www.rfc-editor.org/info/rfc7172>>.
- [TRILL-FM] Senevirathne, T., Finn, N., Salam, S., Kumar, D., Eastlake 3rd, D., Aldrin, S., and Y. Li, "Transparent Interconnection of Lots of Links (TRILL): Fault Management", RFC 7455, March 2015, <<http://www.rfc-editor.org/info/rfc7455>>.

9.2. Informative References

- [OAM-REQ] Senevirathne, T., Bond, D., Aldrin, S., Li, Y., and R. Watve, "Requirements for Operations, Administration, and Maintenance (OAM) in Transparent Interconnection of Lots of Links (TRILL)", RFC 6905, March 2013, <<http://www.rfc-editor.org/info/rfc6905>>.
- [OAM-FRAMEWK] Salam, S., Senevirathne, T., Aldrin, S., and D. Eastlake 3rd, "Transparent Interconnection of Lots of Links (TRILL) Operations, Administration, and Maintenance (OAM) Framework", RFC 7174, May 2014, <<http://www.rfc-editor.org/info/rfc7174>>.
- [Y.1731-2013] ITU-T, "OAM functions and mechanisms for Ethernet based Networks", ITU-T Recommendation G.8013/Y.1731, November 2013.
- [802.1Q] IEEE, "IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks", IEEE Std 802.1Q, December 2014.
- [IEEE1588v1] IEEE, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems Version 1", IEEE Standard 1588, 2002.
- [IEEE1588v2] IEEE, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems Version 2", IEEE Standard 1588, 2008.
- [MPLS-LM-DM] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, September 2011, <<http://www.rfc-editor.org/info/rfc6374>>.
- [OAM] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, June 2011, <<http://www.rfc-editor.org/info/rfc6291>>.
- [IPPM-1DM] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, September 1999, <<http://www.rfc-editor.org/info/rfc2679>>.
- [IPPM-2DM] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", RFC 2681, September 1999, <<http://www.rfc-editor.org/info/rfc2681>>.

- [IPPM-Loss] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC 2680, September 1999, <<http://www.rfc-editor.org/info/rfc2680>>.
- [OAM-Over] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", RFC 7276, June 2014, <<http://www.rfc-editor.org/info/rfc7276>>.

Acknowledgments

The authors gratefully acknowledge Adrian Farrel, Alexey Melnikov, Jan Novak, Carlos Pignataro, Gagan Mohan Goel, Pete Resnick, and Prabhu Raj for their helpful comments.

Authors' Addresses

Tal Mizrahi
Marvell
6 Hamada St.
Yokneam, 20692
Israel

EEmail: talmi@marvell.com

Tissa Senevirathne
Cisco
375 East Tasman Drive
San Jose, CA 95134
United States

EEmail: tsenevir@cisco.com

Samer Salam
Cisco
595 Burrard Street, Suite 2123
Vancouver, BC V7X 1J1
Canada

EEmail: ssalam@cisco.com

Deepak Kumar
Cisco
510 McCarthy Blvd,
Milpitas, CA 95035
United States

Phone : +1 408-853-9760
EEmail: dekumar@cisco.com

Donald Eastlake 3rd
Huawei Technologies
155 Beaver Street
Milford, MA 01757
United States

Phone: +1-508-333-2270
EEmail: d3e3e3@gmail.com