MPLS	D. Frost, Ed.
Internet-Draft	S. Bryant, Ed.
Intended status: Standards Track	Cisco Systems
Expires: December 31, 2010	June 29, 2010

Packet Loss and Delay Measurement for the MPLS Transport Profile
draft-frost-mpls-tp-loss-delay-02

Abstract

An essential Operations, Administration and Maintenance requirement of the MPLS Transport Profile (MPLS-TP) is the ability to monitor performance metrics for packet loss and one-way and two-way delay for MPLS-TP pseudowires, Label Switched Paths, and Sections. This document specifies protocol mechanisms to facilitate the efficient and accurate measurement of these performance metrics.

Requirements Language

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 RFC 2119 (Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” March 1997.) [RFC2119].

Table of Contents

1. Introduction

TOC

The MPLS Transport Profile (MPLS-TP) [I-D.ietf-mpls-tp-framework] (Bocci, M., Bryant, S., Frost, D., Levrau, L., and L. Berger, “A Framework for MPLS in Transport Networks,” May 2010.) comprises the set of protocol functions that meet the requirements [RFC5654] (Niven-Jenkins, B., Brungard, D., Betts, M., Sprecher, N., and S. Ueno, “Requirements of an MPLS Transport Profile,” September 2009.) for the application of MPLS to the construction and operation of packet-switched transport networks.

RFC 5860 [RFC5860] (Vigoureux, M., Ward, D., and M. Betts, “Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks,” May 2010.) specifies Operations, Administration and Maintenance (OAM) definitions and requirements for the measurement of packet loss and one-way and two-way delay for MPLS-TP pseudowires (PWs), Label Switched Paths (LSPs), and Sections. For convenience these definitions and requirements are summarized in the following subsections.

1.1. Review of Requirements

TOC

1.1.1. Requirements for Packet Loss Measurement

TOC

The MPLS-TP OAM toolset must provide a function to enable the quantification of packet loss ratio over a PW, LSP or Section.

The loss of a packet is defined in [RFC2680] (Almes, G., Kalidindi, S., and M. Zekauskas, “A One-way Packet Loss Metric for IPPM,” September 1999.) (Section 2.4). This definition is used here.

Packet loss ratio is defined here to be the ratio of the number of user packets lost to the total number of user packets sent during a defined time interval.

This function may either be performed pro-actively or on-demand.
This function should be performed between End Points of PWs, LSPs and Sections.
It should be possible to rely on user traffic to perform this function.
The protocol solution(s) developed to perform this function must apply to point-to-point co-routed bidirectional LSPs, point-to-point associated bidirectional LSPs, point-to-point unidirectional LSPs and point-to-multipoint (unidirectional) LSPs.

1.1.2. Requirements for Delay Measurement

TOC

The MPLS-TP OAM toolset must provide a function to enable the quantification of the one-way, and if appropriate, the two-way, delay of a PW, LSP or Section.

* The one-way delay is defined in [RFC2679] (Almes, G., Kalidindi, S., and M. Zekauskas, “A One-way Delay Metric for IPPM,” September 1999.) to be the time elapsed from the start of transmission of the first bit of a packet by an End Point until the reception of the last bit of that packet by the other End Point.

* The two-way delay is defined in [RFC2681] (Almes, G., Kalidindi, S., and M. Zekauskas, “A Round-trip Delay Metric for IPPM,” September 1999.) to be the time elapsed from the start of transmission of the first bit of a packet by an End Point until the reception of the last bit of that packet by the same End Point.

Two-way delay may be quantified using data traffic loopback at the remote End Point of the PW, LSP or Section.
Accurate quantification of one-way delay may require clock synchronization, the means for which are outside the scope of this document.
This function should be performed on-demand and may be performed pro-actively.
This function should be performed between End Points of PWs, LSPs and Sections.
In addition to point-to-point co-routed bidirectional LSPs, the protocol solution(s) developed to perform this function must also apply to point-to-point associated bidirectional LSPs, point-to-point unidirectional LSPs and point-to-multipoint (unidirectional) LSPs, but only to enable the quantification of the one-way delay.

TOC

1.2. Protocol Summary

This document specifies two closely-related protocols, one for packet loss measurement (LM) and one for packet delay measurement (DM). These protocols have the following characteristics and capabilities:

* The LM and DM protocols are designed to be simple and to support efficient hardware processing.

* The LM and DM protocols support measurement of loss and delay over MPLS-TP pseudowires and sections, over associated and co-routed bidirectional point-to-point MPLS-TP LSPs, and over unidirectional point-to-point and point-to-multipoint MPLS-TP LSPs.

* The LM and DM protocols support pro-active and on-demand modes of operation.

* The LM and DM protocols use a simple query/response model over bidirectional connections that allows a single node - the querier - to measure the loss or delay of both directions of the connection.

* The LM and DM protocols use query messages to measure the loss or delay of a unidirectional connection. The measurement can either be carried out at the downstream node(s) or at the querier if an out-of-band return path is available.

* The LM and DM protocols do not require that the transmit and receive interfaces be the same at an endpoint of a bidirectional connection.

* The DM protocol is stateless.

* The LM protocol is "almost" stateless: loss is computed as a delta between successive messages, and thus the data associated with the last message received must be retained.

* The LM protocol provides perfect loss measurement if the necessary implementation support is available.

* The LM protocol supports both 32-bit and 64-bit packet counters.

* The DM protocol supports multiple timestamp formats, and provides a simple means for the two endpoints of a bidirectional connection to agree on a preferred format. This procedure reduces to a triviality for implementations supporting only a single timestamp format.

*The DM protocol supports varying the measurement message size in order to measure delays associated with different packet sizes.

1.3. Terminology

TOC

Term	Definition
ACH	Associated Channel Header
DM	Delay Measurement
G-ACh	Generic Associated Channel
LM	Loss Measurement
LSE	Label Stack Entry
LSP	Label Switched Path
LSR	Label Switching Router
MPLS-TP	MPLS Transport Profile
NTP	Network Time Protocol
OAM	Operations, Administration and Maintenance
PTP	Precision Time Protocol
PW	Pseudowire
TC	Traffic Class

2. Overview

TOC

The basic procedures for measuring loss and delay over a bidirectional connection are conceptually simple. The following figure shows the reference scenario.

```
          T1               T2
+-------+/     Query       \+-------+
|       | - - - - - - - - ->|       |
|   A   |===================|   B   |
|       |<- - - - - - - - - |       |
+-------+\     Response    /+-------+
          T4               T3
```

Figure 1

The figure shows a bidirectional connection between two nodes, A and B, and illustrates the temporal reference points T1-T4 associated with a measurement operation that takes place at A. The operation consists of A sending a query message to B, and B sending back a response. Each reference point indicates the point in time at which either the query or the response message is transmitted or received over the connection.

In this situation, A can arrange to measure the packet loss over the connection in the forward and reverse directions by sending Loss Measurement (LM) query messages to B each of which contains the count of packets transmitted prior to time T1 over the connection to B (A_TxP). When the message reaches B, it appends two values and reflects the message back to A: the count of packets received prior to time T2 over the connection from A (B_RxP), and the count of packets transmitted prior to time T3 over the connection to A (B_TxP). When the response reaches A, it appends a fourth value, the count of packets received prior to time T4 over the connection from B (A_RxP).

These four counter values enable A to compute the desired loss statistics. Because the transmit count at A and the receive count at B (and vice versa) may not be synchronized at the time of the first message, and to limit the effects of counter wrap, the loss is computed in the form of a delta between messages.

To measure at A the delay over the connection to B, a Delay Measurement (DM) query message is sent from A to B containing a timestamp recording the instant at which it is transmitted, i.e. T1. When the message reaches B, a timestamp is added recording the instant at which it is received (T2). The message can now be reflected from B to A, with B adding its transmit timestamp (T3) and A adding its receive timestamp (T4). These four timestamps enable A to compute the one-way delay in each direction, as well as the two-way delay for the connection. The one-way delay computations require that the clocks of A and B be synchronized; mechanisms for clock synchronization are outside the scope of this document.

In the case of a unidirectional connection rooted at A, the first half of each of the above procedures can be carried out to measure the forward one-way loss and delay associated with the connection. At this point the measurement can either take place at the terminal node(s) of the connection rather than at A, or an out-of-band channel can be used, if available, to communicate the data back to A.

In the context of MPLS-TP, LM and DM messages flow over the Generic Associated Channel (G-ACh) [RFC5586] (Bocci, M., Vigoureux, M., and S. Bryant, “MPLS Generic Associated Channel,” June 2009.) of an MPLS-TP pseudowire, LSP, or Section. The term "connection" is used in this document to mean "pseudowire, LSP, or Section". Although this document often speaks of "measuring the loss or delay associated with a connection" for simplicity, LM and DM actually occur with respect to a particular class of packets flowing over a connection. This is discussed in more detail in Section 5 (Packet Profiles and Quality of Service).

TOC

2.1. Implementation Considerations

The challenge in carrying out the above procedures lies with the implementation. For accurate loss measurement to occur, packets must not be sent between the time the transmit count for an outbound LM message is determined and the time the message is actually transmitted. Similarly, packets must not be received and processed between the time an LM message is received and the time the receive count for the message is determined. For accurate delay measurement, timestamps must be recorded in DM messages at a point in time as close as possible to when the message is actually transmitted or received over the connection.

These accuracy requirements imply that a hardware-based forwarding implementation may require hardware support for the processing of LM and DM messages. An important consideration of the LM/DM protocol and message format is therefore support for efficient hardware processing. In situations where such accuracy is not required, or the necessary level of support is not available, an implementation MAY still generate and respond to LM and DM messages but SHOULD make its accuracy limitations clear to the user. In general the DM procedures described in this document remain viable under these conditions, but the procedures for LM may be inadequate.

The LM procedures described in this document have the advantage of providing perfect packet loss accounting if the necessary implementation support is available. This is a desirable capability in an LM protocol for MPLS-TP given that loss levels for typical MPLS-TP connections are expected to be quite low, and that even small amounts of loss on such connections may be unacceptable. This capability, however, may well come at the expense of more costly hardware, and in some environments this cost may be prohibitive. Thus it is desirable to define an additional set of LM procedures for MPLS-TP that support deployments in which perfect loss accounting is not required. Such alternative procedures rely on the generation of either existing or new MPLS-TP OAM message types, which are subjected to loss accounting as a proxy for user traffic in order to infer approximate loss levels of the latter. This alternative approach to LM is for further study and will be described in a companion document.

2.2. Packet Loss Measurement

TOC

Suppose a bidirectional connection such as an MPLS-TP pseudowire, bidirectional LSP, or Section exists between the LSRs A and B. The objective is to measure at A the following two quantities associated with the connection:

> A_TxLoss (transmit loss): the number of packets transmitted by A over the connection but not received at B;
>
> A_RxLoss (receive loss): the number of packets transmitted by B over the connection but not received at A.

This is accomplished by initiating a Loss Measurement (LM) operation at A, which consists of transmission of a sequence of LM query messages (LM[1], LM[2], ...) over the connection at a specified rate, such as one every 100 milliseconds. Each message LM[n] contains the following value:

> A_TxP[n]: the total count of packets transmitted by A over the connection prior to the time this message is transmitted.

When such a message is received at B, the following value is recorded in the message:

> B_RxP[n]: the total count of packets received by B over the connection at the time this message is received (excluding the message itself).

At this point, B inserts an appropriate response code into the message and transmits it back to A, recording within it the following value:

> B_TxP[n]: the total count of packets transmitted by B over the connection prior to the time this response is transmitted.

When the message response is received back at A, the following value is recorded in the message:

> A_RxP[n]: the total count of packets received by A over the connection at the time this response is received (excluding the message itself).

The transmit loss A_TxLoss[n-1,n] and receive loss A_RxLoss[n-1,n] within the measurement interval marked by the messages LM[n-1] and LM[n] are computed by A as follows:
A_TxLoss[n-1,n] = (A_TxP[n] - A_TxP[n-1]) - (B_RxP[n] - B_RxP[n-1])
A_RxLoss[n-1,n] = (B_TxP[n] - B_TxP[n-1]) - (A_RxP[n] - A_RxP[n-1])
where the arithmetic is modulo the counter size.
The derived values

> A_TxLoss = A_TxLoss[1,2] + A_TxLoss[2,3] + ...
>
> A_RxLoss = A_RxLoss[1,2] + A_RxLoss[2,3] + ...

are updated each time a response to an LM message is received and processed, and represent the total transmit and receive loss over the connection since the LM operation was initiated.

When computing the values A_TxLoss[n-1,n] and A_RxLoss[n-1,n] the possibility of counter wrap must be taken into account. Consider for example the values of the A_TxP counter at times n-1 and n. Clearly if A_TxP[n] is allowed to wrap to 0 and then beyond to a value equal to or greater than A_TxP[n-1], the computation of an unambiguous A_TxLoss[n-1,n] value will be impossible. Therefore the LM message rate MUST be sufficiently high, given the counter size and the speed and minimum packet size of the underlying connection, that this condition cannot arise. For example, a 32-bit counter for a 100 Gbps link with a minimum packet size of 64 bytes can wrap in 2^32 / (10^11/(64*8)) = ~22 seconds, which is therefore an upper bound on the LM message interval under such conditions.

2.3. Delay Measurement

TOC

Suppose a bidirectional connection such as an MPLS-TP pseudowire, bidirectional LSP, or Section exists between the LSRs A and B. The objective is to measure at A one or more of the following quantities associated with the connection:

- The one-way delay associated with the forward (A to B) direction of the connection;
- The one-way delay associated with the reverse (B to A) direction of the connection;
- The two-way delay (A to B to A) associated with the connection.

In the case of two-way delay, there are actually two possible metrics of interest. The "strict" two-way delay is the sum of the one-way delays in each direction and reflects the two-way delay of the connection itself, irrespective of processing delays within the remote endpoint B. The "loose" two-way delay is the definition of two-way delay stated in Section 1.1.2 (Requirements for Delay Measurement) and includes in addition any delay associated with remote endpoint processing.

Measurement of the one-way delay quantities requires that the clocks of A and B be synchronized, whereas the two-way delay can be measured directly even when this is not the case (provided A and B have stable clocks).

The measurement is accomplished by sending a Delay Measurement (DM) query message over the connection to B which contains the following timestamp:

> T1: the time the DM query message is transmitted from A.

When the message arrives at B, the following timestamp is recorded in the message:

> T2: the time the DM query message is received at B.

At this point B inserts an appropriate response code into the message and transmits it back to A, recording within it the following timestamp:

> T3: the time the DM response message is transmitted from B.

When the message arrives back at A, the following timestamp is recorded in the message:

> T4: the time the DM response message is received back at A.

At this point, A can compute the strict two-way delay associated with the connection as

> strict two-way delay = (T4 - T1) - (T3 - T2)

and the loose two-way delay as

> loose two-way delay = T4 - T1.

If the clocks of A and B are known at A to be synchronized, then both one-way delay values, as well as the strict two-way delay, can be computed at A as

> forward one-way delay = T2 - T1
>
> reverse one-way delay = T4 - T3
>
> strict two-way delay = forward delay + reverse delay.

2.3.1. Timestamp Format

TOC

There are two significant timestamp formats in common use: the timestamp format of the Internet standard Network Time Protocol (NTP), described in [RFC1305] (Mills, D., “Network Time Protocol (Version 3) Specification, Implementation,” March 1992.) and [RFC2030] (Mills, D., “Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI,” October 1996.), and the timestamp format used in the IEEE 1588 Precision Time Protocol (PTP) [IEEE1588] (IEEE, “1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems,” March 2008.).

[Editor's note: There are actually two PTP timestamp formats: the 1588v1 format consists of a 32-bit seconds field and a 32-bit nanoseconds field; in 1588v2 the seconds field was extended to 48 bits.]

The NTP format has the advantages of wide use and long deployment in the Internet, and was specifically designed to make the computation of timestamp differences as simple and efficient as possible. On the other hand, there is also now a significant deployment of equipment designed to support the PTP format.

The approach taken in this document is therefore to include in DM messages fields which identify the timestamp formats used by the two devices involved in a DM operation. This implies that an LSR attempting to carry out a DM operation may be faced with the problem of computing with and possibly reconciling different timestamp formats. Support for multiple timestamp formats is OPTIONAL. An implementation SHOULD, however, make clear which timestamp formats it supports and the extent of its support for computation with and reconciliation of different formats for purposes of delay measurement.

In recognition of the wide deployment, particularly in hardware-based timing implementations, of IEEE 1588 PTP, the PTP timestamp format is the default format used in DM messages. This format MUST be supported.

2.4. Delay Variation Measurement TOC

Packet Delay Variation [RFC3393] (Demichelis, C. and P. Chimento, “IP Packet Delay Variation Metric for IP Performance Metrics (IPPM),” November 2002.) is another performance metric important in some applications. The PDV of a pair of packets within a stream of packets is defined for a selected pair of packets in the stream going from measurement point MP1 to measurement point MP2. The PDV is the difference between the one-way delay of the selected packets.

A PDV measurement can therefore be derived from successive delay measurements obtained through the procedures in Section 2.3 (Delay Measurement). An important point regarding PDV measurement, however, is that it can be carried out based on one-way delay measurements even when the clocks of the two systems involved in those measurements are not synchronized.

2.5. Unidirectional Connections TOC

In the case that the connection from A to (B1, ..., Bk) is unidirectional, i.e. is a unidirectional LSP, LM and DM measurements can be carried out at B1, ..., Bk instead of at A.

For LM this is accomplished by initiating an LM operation at A and carrying out the same procedures as for bidirectional connections, except that no responses from B1, ..., Bk to A are generated. Instead, each terminal node B uses the A_TxP and B_RxP values in the LM messages it receives to compute the receive loss associated with the connection in essentially the same way as described previously, i.e.
B_RxLoss[n-1,n] = (A_TxP[n] - A_TxP[n-1]) - (B_RxP[n] - B_RxP[n-1])
For DM, of course, only the forward one-way delay can be measured and the clock synchronization requirement applies.
Alternatively, if an out-of-band connection from a terminal node B back to A is available, the LM and DM message responses can be communicated to A via this connection so that the measurements can be carried out at A.

2.6. Distributed Systems

TOC

The overview of the bidirectional measurement process presented in Section 2 (Overview) is also applicable when the transmit and receive interfaces at A or B differ from one another, as may occur when the connection is an MPLS-TP LSP that is not co-routed. Some additional considerations, however, do apply in this case:

* If the transmit and receive interfaces reside on different line cards, the clocks of those line cards must be synchronized in order to compute the two-way delay.

* The DM protocol specified in this document requires that the timestamp formats used by the interfaces that receive a DM query and transmit a DM response agree.

* The LM protocol specified in this document supports both 32-bit and 64-bit counter sizes, but the use of 32-bit counters at any of the up to four interfaces involved in an LM operation will result in 32-bit LM calculations for both directions of the connection.

[Editor's note: The last two restrictions could be relaxed if desired, at the expense of some additional protocol complexity.]

---

3. Packet Format

TOC

Loss Measurement and Delay Measurement messages flow over the Generic Associated Channel (G-ACh) [RFC5586] (Bocci, M., Vigoureux, M., and S.

Bryant, “MPLS Generic Associated Channel,” June 2009.) of an MPLS-TP connection (pseudowire, LSP or Section).
[Editor's note: The question of ACH TLV usage and the manner of supporting metadata such as authentication keys and node identifiers is deliberately omitted. These issues will be addressed in a future version of the document.]

3.1. Loss Measurement Message Format

TOC

The format of a Loss Measurement message, beginning with the Associated Channel Header (ACH), is as follows:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 1|Version|   Reserved    |     0xHH (MPLS-TP Loss)       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Version| Flags |  Control Code |     Session Identifier        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Sequence Number                         |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Counter 1                            |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
.                                                               .
.                                                               .
.                                                               .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Counter 4                            |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: Loss Measurement Message Format

The meanings of the fields following the ACH are summarized in the following table.

Field	Meaning
Version	Protocol version
Flags	Message control flags

Control Code	Code identifying the query or response type
Session Identifier	Set arbitrarily by the querier
Sequence Number	64-bit sequence number, incremented for each message
Counter 1-4	Packet counter values in network byte order

The possible values for these fields are as follows.

Version: Currently set to 0.

Flags: Each bit represents a message control flag. The flags, listed in left-to-right (most- to least-significant-bit) order, are:

> Q/R: Set to 0 for a Query and 1 for a Response.
>
> X: Extended data format. Indicates support for extended (64-bit) counter values. Initialized to 1 upon creation (and prior to transmission) of an LM Query and copied from an LM Query to an LM response. Set to 0 when the LM message is transmitted or received over an interface that writes 32-bit counter values.
>
> Remaining bits: Reserved for future specification and set to 0.

Control Code: Set as follows according to whether the message is a Query or a Response as identified by the Q/R flag.

> For a Query:
>
> > 0x0: Query (in-band response requested). Indicates that this query has been sent over a bidirectional connection and the response is expected over the same connection.
> >
> > 0x1: Query (out-of-band response requested). Indicates that the response should be sent via an out-of-band channel.
> >
> > 0x2: Query (no response requested). Indicates that no response to the query should be sent.
>
> For a Response:
>
> > 0x1: Success. Indicates that the operation was successful.
> >
> > 0x8: Notification - Data Format Invalid. Indicates that the query was processed but the format of the data fields in this response may be inconsistent. Consequently these data fields MUST NOT be used for measurement.
> >
> > 0x10: Error - Unspecified Error. Indicates that the operation failed for an unspecified reason.

> 0x11: Error - Unsupported Version. Indicates that the operation failed because the protocol version supplied in the query message is not supported.
>
> 0x12: Error - Unsupported Control Code. Indicates that the operation failed because the Control Code requested an operation that is not available for this connection.
>
> 0x13: Error - Authentication Failure. Indicates that the operation failed because the authentication data supplied in the query was missing or incorrect.
>
> 0x14: Error - Invalid Source Node Identifier. Indicates that the operation failed because the Source Node Identifier supplied in the query is not expected.
>
> 0x15: Error - Invalid Destination Node Identifier. Indicates that the operation failed because the Destination Node Identifier supplied in the query is not the identifier of this node.
>
> 0x16: Error - Connection Mismatch. Indicates that the operation failed because the connection identifier supplied in the query did not match the connection over which the query was received.
>
> 0x17: Error - Query Rate Exceeded. Indicates that the operation failed because the rate of query messages exceeded the configured threshold.
>
> 0x18: Error - Administrative Block. Indicates that the operation failed because it has been administratively disallowed.
>
> 0x19: Error - Temporary Resource Exhaustion. Indicates that the operation failed because node resources were not available.

Session Identifier: Set arbitrarily in a query and copied in the response, if any.

Counter 1-4: Referring to Section 2.2 (Packet Loss Measurement), when a query is sent from A, Counter 1 is set to A_TxP and the other counter fields are set to 0. When the query is received at B, Counter 2 is set to B_RxP. At this point, B copies Counter 1 to Counter 3 and Counter 2 to Counter 4, and re-initializes Counter 1 and Counter 2 to 0. When B transmits the response, Counter 1 is set to B_TxP. When the response is received at A, Counter 2 is set to A_RxP. All counter values MUST be in network byte order.

When a 32-bit counter value is written to one of the counter fields, that value SHALL be written to the low-order 32 bits of the field; the high-order 32 bits of the field MUST, in this case, be set to 0.

3.2. Delay Measurement Message Format

TOC

The format of a Delay Measurement message, beginning with the Associated Channel Header (ACH), is as follows:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 1|Version|   Reserved    |     0xHH (MPLS-TP Delay)      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Version| Flags |  Control Code |     Session Identifier        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Message Length        |  QTF  |  RTF  | RPTF  | Resv  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Timestamp 1                          |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
.                                                               .
.                                                               .
.                                                               .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Timestamp 4                          |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
~                            Padding                            ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: Delay Measurement Message Format

The meanings of the fields following the ACH are summarized in the following table.

Field	**Meaning**
Version	Protocol version
Flags	Message control flags
Control Code	Code identifying the query or response type
Session Identifier	Set arbitrarily by the querier
Message Length	Total length of this message in bytes
QTF	Querier timestamp format
RTF	Responder timestamp format
RPTF	Responder's preferred timestamp format

Resv (Reserved)	Reserved for future specification
Timestamp 1-4	64-bit timestamp values
Padding	Optional padding

The possible values for these fields are as follows.

Version: Currently set to 0.

Flags: As specified in Section 3.1 (Loss Measurement Message Format).

Control Code: As specified in Section 3.1 (Loss Measurement Message Format).

Session Identifier: Set arbitrarily in a query and copied in the response, if any.

Message Length: Set to the total length of this message, excluding the ACH, in bytes.

Querier Timestamp Format: The format of the timestamp values written by the querier, as specified in Section 3.3 (Timestamp Field Formats).

Responder Timestamp Format: The format of the timestamp values written by the responder, as specified in Section 3.3 (Timestamp Field Formats).

Responder's Preferred Timestamp Format: The timestamp format preferred by the responder, as specified in Section 3.3 (Timestamp Field Formats).

Resv (Reserved): Currently set to 0.

Timestamp 1-4: Referring to Section 2.3 (Delay Measurement), when a query is sent from A, Timestamp 1 is set to T1 and the other timestamp fields are set to 0. When the query is received at B, Timestamp 2 is set to T2. At this point, B copies Timestamp 1 to Timestamp 3 and Timestamp 2 to Timestamp 4, and re-initializes Timestamp 1 and Timestamp 2 to 0. When B transmits the response, Timestamp 1 is set to T3. When the response is received at A, Timestamp 2 is set to T4. The actual formats of the timestamp fields written by A and B are indicated by the Querier Timestamp Format and Responder Timestamp Format fields respectively.

Padding: One or more octets of padding may optionally follow the Timestamp 4 field in a query, in order to allow for delay measurement based on packets of a particular size. The value of the first octet of padding provides information about the padding. If in a Query the first bit of the first pad octet is 1, the padding SHALL be copied to the response, assuming one was requested. If this bit is 0, the response MUST NOT include padding. The remaining bits in the first pad octet are reserved and SHALL be set to 0. The values of the remaining pad octets, if present, are arbitrary.

3.3. Timestamp Field Formats

TOC

The following timestamp format field values are specified in this document:

> 0x0: Network Time Protocol version 4 timestamp format [RFC2030] (Mills, D., “Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI,” October 1996.). This format consists of a 32-bit seconds field followed by a 32-bit fractional seconds field, so that it can be regarded as a fixed-point 64-bit quantity.
>
> 0x2: IEEE 1588-2002 (1588v1) Precision Time Protocol timestamp format [IEEE1588] (IEEE, “1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems,” March 2008.). This format consists of a 32-bit seconds field followed by a 32-bit nanoseconds field.

In recognition of the wide deployment, particularly in hardware-based timing implementations, of IEEE 1588 PTP, the PTP timestamp format is the default format used in Delay Measurement messages. This format MUST be supported. Support for other timestamp formats is OPTIONAL. Timestamp formats of n < 64 bits in size SHALL be encoded in the 64-bit timestamp fields specified in this document using the n high-order bits of the field. The remaining 64 - n low-order bits in the field SHOULD be set to 0 and MUST be ignored when reading the field.

---

4. Operation

TOC

4.1. Loss Measurement Procedures

TOC

4.1.1. Initiating a Loss Measurement Operation

TOC

An LM operation for a particular MPLS-TP connection consists of sending a sequence (LM[1], LM[2], ...) of LM query messages over the connection at a specific rate and processing the responses received, if any. As described in Section 2.2 (Packet Loss Measurement), the packet loss associated with the connection during the operation is computed as a delta between successive messages; these deltas can be accumulated to obtain a running total of the packet loss for the connection.
The query message transmission rate MUST be sufficiently high, given the LM message counter size (which can be either 32 or 64 bits) and the speed and minimum packet size of the underlying connection, that the ambiguity condition noted in Section 2.2 (Packet Loss Measurement) cannot arise. The implementation SHOULD assume, in evaluating this

rate, that the counter size is 32 bits unless explicitly configured otherwise, or unless (in the case of a bidirectional connection) all local and remote interfaces involved in the LM operation are known to be 64-bit-capable, which can be inferred from the value of the X flag in an LM response.

4.1.2. Transmitting a Loss Measurement Query

TOC

When transmitting an LM Query over an MPLS-TP connection, the Version and Reserved fields MUST be set to 0. The Q/R flag MUST be set to 0. The X flag MUST be set to 1 if the transmitting interface writes 64-bit LM counters, and otherwise MUST be set to 0 to indicate that 32-bit counters are written. The remaining flag bits MUST be set to 0.
The Control Code field MUST be set to one of the values for Query messages listed in Section 3.1 (Loss Measurement Message Format); if the connection is unidirectional, this field MUST NOT be set to 0x0 (Query: in-band response requested).
The Session Identifier field can be set arbitrarily.
The Sequence Number field MUST be set to 0 for the first message sent after device initialization or explicit reset, and incremented by 1 for each subsequent message sent.
The Counter 1 field SHOULD be set to the total count of packets transmitted over the connection prior to this LM Query. The remaining Counter fields MUST be set to 0.

4.1.3. Receiving a Loss Measurement Query

TOC

Upon receipt of an LM Query message, the Counter 2 field SHOULD be set to the total count of packets received over the connection prior to this LM Query. If the receiving interface writes 32-bit LM counters, the X flag MUST be set to 0.
At this point the LM Query message must be inspected. If the Control Code field is set to 0x2 (no response requested), an LM Response message MUST NOT be transmitted. If the Control Code field is set to 0x0 (in-band response requested) or 0x1 (out-of-band response requested), then an in-band or out-of-band response, respectively, SHOULD be transmitted unless this has been prevented by an administrative, security or congestion control mechanism.

TOC

4.1.4. Transmitting a Loss Measurement Response

When constructing a Response to an LM Query, the Version and Reserved fields MUST be set to 0. The Q/R flag MUST be set to 1. The the X flag MUST be set to 0 if the transmitting interface writes 32-bit LM counters; otherwise its value MUST be copied from the LM Query. The remaining flag bits MUST be set to 0.

The Session Identifier and Sequence Number fields MUST be copied from the LM Query. The Counter 1 and Counter 2 fields from the LM Query MUST be copied to the Counter 3 and Counter 4 fields, respectively, of the LM Response.

The Control Code field MUST be set to one of the values for Response messages listed in Section 3.1 (Loss Measurement Message Format). The value 0x10 (Unspecified Error) SHOULD NOT be used if one of the other more specific error codes is applicable.

If the response is transmitted in-band, the Counter 1 field SHOULD be set to the total count of packets transmitted over the connection prior to this LM Response. If the response is transmitted out-of-band, the Counter 1 field MUST be set to 0. In either case, the Counter 2 field MUST be set to 0.

4.1.5. Receiving a Loss Measurement Response

TOC

Upon in-band receipt of an LM Response message, the Counter 2 field SHOULD be set to the total count of packets received over the connection prior to this LM Response. If the receiving interface writes 32-bit LM counters, the X flag MUST be set to 0.

Upon out-of-band receipt of an LM Response message, the Counter 1 and Counter 2 fields MUST NOT be used for purposes of loss measurement.

If the Control Code in an LM Response is anything other than 0x1 (Success), the counter values in the response MUST NOT be used for purposes of loss measurement. When the Control Code indicates an error condition, the LM operation SHOULD be suspended and an appropriate notification to the user generated. If a temporary error condition is indicated, the LM operation MAY be restarted automatically.

4.1.6. Loss Calculation

TOC

Calculation of packet loss is carried out according to the procedures in Section 2.2 (Packet Loss Measurement). The X flag in an LM message informs the device performing the calculation whether to perform 32-bit or 64-bit arithmetic. If the flag value is equal to 1, all interfaces involved in the LM operation have written 64-bit counter values, and 64-bit arithmetic can be used. If the flag value is equal to 0, at least one interface involved in the operation has written a 32-bit counter value, and 32-bit arithmetic is carried out using the low-order 32 bits of each counter value.

4.1.7. Message Loss and Packet Misorder Conditions

TOC

Because an LM operation consists of a message sequence with state maintained from one message to the next, LM is subject to the effects of lost messages and misordered packets in a way that DM is not. Because this state exists only on the querier, the handling of these conditions is, strictly speaking, a local matter. This section, however, presents RECOMMENDED procedures for handling such conditions.

The first kind of anomaly that may occur is that one or more LM messages may be lost in transit. The effect of such loss is that when an LM Response is next received at the querier, an unambiguous interpretation of the counter values it contains may be impossible, for the reasons described at the end of Section 2.2 (Packet Loss Measurement). Whether this is so depends on the number of messages lost and the other variables mentioned in that section, such as the LM message rate and the connection parameters.

Another possibility is that LM messages are misordered in transit, so that for instance the response to LM[n] is received prior to the response to LM[n-1]. A typical implementation will discard the late response to LM[n-1], so that the effect is the same as the case of a lost message.

Finally, LM is subject to the possibility that data packets are misordered relative to LM messages. This condition can result, for example, in a transmit count of 100 and a corresponding receive count of 101. The effect here is that the A_TxLoss[n-1,n] value (for example) for a given measurement interval will appear to be extremely (if not impossibly) large. The other case, where an LM message arrives earlier than some of the packets, simply results in those packets being counted as lost, which is usually what is desired.

[Editor's note: Text to be added here about handling the above conditions with sequence numbers and thresholds.]

4.2. Delay Measurement Procedures

TOC

TOC

4.2.1. Transmitting a Delay Measurement Query

When transmitting a DM Query over an MPLS-TP connection, the Version and Reserved fields MUST be set to 0. The Q/R flag MUST be set to 0 and the remaining flag bits MUST be set to 0.
The Control Code field MUST be set to one of the values for Query messages listed in Section 3.1 (Loss Measurement Message Format); if the connection is unidirectional, this field MUST NOT be set to 0x0 (Query: in-band response requested).
The Session Identifier field can be set arbitrarily.
The Querier Timestamp Format field MUST be set to the timestamp format used by the querier when writing timestamp fields in this message; the possible values for this field are listed in Section 3.3 (Timestamp Field Formats). The Responder Timestamp Format and Responder's Preferred Timestamp Format fields MUST be set to 0.
The Timestamp 1 field SHOULD be set to the time at which this DM Query is transmitted, in the format indicated by the Querier Timestamp Format field. The other timestamp fields MUST be set to 0.
One or more pad octets MAY follow the Timestamp 4 field, as described in Section 3.2 (Delay Measurement Message Format).

4.2.2. Receiving a Delay Measurement Query

TOC

Upon receipt of a DM Query message, the Timestamp 2 field SHOULD be set to the time at which this DM Query is received.
At this point the DM Query message must be inspected. If the Control Code field is set to 0x2 (no response requested), a DM Response message MUST NOT be transmitted. If the Control Code field is set to 0x0 (in-band response requested) or 0x1 (out-of-band response requested), then an in-band or out-of-band response, respectively, SHOULD be transmitted unless this has been prevented by an administrative, security or congestion control mechanism.

4.2.3. Transmitting a Delay Measurement Response

TOC

When constructing a Response to a DM Query, the Version and Reserved fields MUST be set to 0. The Q/R flag MUST be set to 1 and the remaining flag bits MUST be set to 0.
The Session Identifier and Querier Timestamp Format (QTF) fields MUST be copied from the DM Query. The Timestamp 1 and Timestamp 2 fields from the DM Query MUST be copied to the Timestamp 3 and Timestamp 4 fields, respectively, of the DM Response.
The Responder Timestamp Format (RTF) field MUST be set to the timestamp format used by the responder when writing timestamp fields in this

message, i.e. Timestamp 4 and (if applicable) Timestamp 1; the possible values for this field are listed in Section 3.3 (Timestamp Field Formats). Furthermore, the RTF field MUST be set equal either to the QTF or the RPTF field. See Section 4.2.5 (Timestamp Format Negotiation) for guidelines on selection of the value for this field.

The Responder's Preferred Timestamp Format (RPTF) field MUST be set to one of the values listed in Section 3.3 (Timestamp Field Formats) and SHOULD be set to indicate the timestamp format with which the responder can provide the best accuracy for purposes of delay measurement.

The Control Code field MUST be set to one of the values for Response messages listed in Section 3.1 (Loss Measurement Message Format). The value 0x10 (Unspecified Error) SHOULD NOT be used if one of the other more specific error codes is applicable.

If the response is transmitted in-band, the Timestamp 1 field SHOULD be set to the time at which this DM Response is transmitted. If the response is transmitted out-of-band, the Timestamp 1 field MUST be set to 0. In either case, the Timestamp 2 field MUST be set to 0.

If the response is transmitted in-band and the Control Code in the message is 0x1 (Success), then the Timestamp 1 and Timestamp 4 fields MUST have the same format, which will be the format indicated in the Responder Timestamp Format field.

Padding SHALL be included in the response if, and only if, padding was present in the DM Query and the first bit of the first octet of that padding was set to 1, in which case the response padding MUST be identical to the query padding.

4.2.4. Receiving a Delay Measurement Response

TOC

Upon in-band receipt of a DM Response message, the Timestamp 2 field SHOULD be set to the time at which this DM Response is received.

Upon out-of-band receipt of a DM Response message, the Timestamp 1 and Timestamp 2 fields MUST NOT be used for purposes of delay measurement.

If the Control Code in a DM Response is anything other than 0x1 (Success), the timestamp values in the response MUST NOT be used for purposes of delay measurement. When the Control Code indicates an error condition, an appropriate notification to the user SHOULD be generated.

4.2.5. Timestamp Format Negotiation

TOC

In case either the querier or the responder in a DM transaction is capable of supporting multiple timestamp formats, it is desirable to determine the optimal format for purposes of delay measurement on a particular connection. The procedures for making this determination SHALL be as follows.

Upon sending an initial DM Query over a connection, the querier sets the Querier Timestamp Format (QTF) field to its preferred timestamp format.
Upon receiving any DM Query message, the responder determines whether it is capable of writing timestamps in the format specified by the QTF field. If so, the Responder Timestamp Format (RTF) field is set equal to the QTF field. If not, the RTF field is set equal to the Responder's Preferred Timestamp Format (RPTF) field.
The process of changing from one timestamp format to another at the responder may result in the Timestamp 1 and Timestamp 4 fields in an in-band DM Response having different formats. If this is the case, the Control Code in the response MUST NOT be set to 0x1 (Success). Unless an error condition has occurred, the Control Code MUST be set to 0x2 (Notification - Data Format Invalid).
Upon receiving a DM Response, the querier knows from the RTF field in the message whether the responder is capable of supporting its preferred timestamp format: if it is, the RTF will be equal to the QTF. The querier also knows the responder's preferred timestamp format from the RPTF field. The querier can then decide whether to retain its current QTF or to change it and repeat the negotiation procedures.

4.2.5.1. Single-Format Procedures

TOC

When an implementation supports only one timestamp format, the procedures above reduce to the following simple behavior:

*All DM Queries are transmitted with the same QTF;

*All DM Responses are transmitted with the same RTF, and the RPTF is always set equal to the RTF;

*All DM Responses received with RTF not equal to QTF are discarded;

*On a unidirectional connection, all DM Queries received with QTF not equal to the supported format are discarded.

---

5. Packet Profiles and Quality of Service

TOC

Although this document has referred, for simplicity, to measuring the packet loss or delay associated with a connection, it is more precise to say that these measurement operations occur with respect to a specific class of packets transiting the connection. Such a class is referred to as a "packet profile".

Care must be taken to ensure that the endpoints of an LM or DM operation agree on the packet profile. For DM this reduces to ensuring that query and response messages are assigned to the same traffic class, while for LM it requires that the LM counters at each endpoint count the same kinds of packets.

This document considers two aspects of packet profile support pertinent to loss and delay measurement:

*Quality of Service

*Loss Measurement of OAM Messages

5.1. Quality of Service

TOC

For connections that support multiple traffic classes, such as those that employ the Traffic Class (TC) field [RFC5462] (Andersson, L. and R. Asati, “Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field,” February 2009.) in the MPLS Label Stack Entry (LSE) for Differentiated Services [RFC3270] (Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, “Multi-Protocol Label Switching (MPLS) Support of Differentiated Services,” May 2002.), the implementation MUST provide the capability to perform delay measurement on a per-traffic-class basis, by assigning the DM messages themselves to the corresponding class.

For connections that support multiple traffic classes, the implementation SHOULD provide the capability to perform loss measurement on a per-traffic-class basis, and MAY provide the more general capability to perform loss measurement on a subset of the traffic classes supported by the connection, by restricting the LM packet profile (i.e. the class of packets counted by the LM counters) accordingly. LM messages themselves SHOULD be assigned to a traffic class equal to or better than the best traffic class within the LM packet profile.

5.2. Loss Measurement of OAM Messages

TOC

By default the LM packet profile MUST include packets transmitted and received over the Generic Associated Channel (G-ACh) associated with a connection. An implementation MAY provide the means to alter the LM packet profile to exclude some or all G-ACh messages.

6. Congestion Considerations

TOC

An MPLS-TP network may be traffic-engineered in such a way that the bandwidth required both for client traffic and for control, management and OAM traffic is always available. The following congestion considerations therefore apply only when this is not the case.

The proactive generation of Loss Measurement and Delay Measurement messages for purposes of monitoring the performance of an MPLS-TP connection naturally results in a degree of additional load placed on both the network and the terminal nodes of the connection. When configuring such monitoring, operators should be mindful of the overhead involved and should choose transmit rates that do not stress network resources unduly; such choices must be informed by the deployment context. In case of slower links or lower-speed devices, for example, lower Loss Measurement message rates can be chosen, up to the limits noted at the end of Section 2.2 (Packet Loss Measurement).

In general, lower measurement message rates place less load on the network at the expense of reduced granularity. For delay measurement this reduced granularity translates to a greater possibility that the delay associated with a connection temporarily exceeds the expected threshold without detection. For loss measurement, it translates to a larger gap in loss information in case of exceptional circumstances such as lost LM messages or misordered packets.

When carrying out a sustained measurement operation such as an LM operation or continuous pro-active DM operation, the querier SHOULD take note of the number of lost measurement messages (queries for which a response is never received) and set a corresponding Measurement Message Loss Threshold. If this threshold is exceeded, the measurement operation SHOULD be suspended so as not to exacerbate the possible congestion condition. This suspension SHOULD be accompanied by an appropriate notification to the user so that the condition can be investigated and corrected.

From the receiver perspective, the main consideration is the possibility of receiving an excessive quantity of measurement messages. An implementation SHOULD employ a mechanism such as rate-limiting to guard against the effects of this case. Authentication procedures can also be used to ensure that only queries from authorized devices are processed.

7. Security Considerations

TOC

There are two main types of security considerations associated with the exchange of performance monitoring messages such as those described in

this document: the possibility of a malicious or misconfigured device generating an excessive quantity of messages, causing service impairment; and the possibility of an unauthorized device learning the data contained in or implied by such messages.
The first consideration is discussed in Section 6 (Congestion Considerations). If reception of performance-related data by unauthorized devices is an operational concern, message authentication procedures such as those described in [xref] should be used to ensure that only queries from authorized devices are processed.

8. IANA Considerations

TOC

A future version of this document will detail IANA considerations for:

- *ACH Channel Types for LM and DM messages
- *Timestamp format registry
- *LM and DM Control Codes

9. References

TOC

9.1. Normative References

TOC

[RFC2119]	Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).
[RFC5586]	Bocci, M., Vigoureux, M., and S. Bryant, “MPLS Generic Associated Channel,” RFC 5586, June 2009 (TXT).
[RFC5654]	Niven-Jenkins, B., Brungard, D., Betts, M., Sprecher, N., and S. Ueno, “Requirements of an MPLS Transport Profile,” RFC 5654, September 2009 (TXT).
[RFC5860]	Vigoureux, M., Ward, D., and M. Betts, “Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks,” RFC 5860, May 2010 (TXT).

9.2. Informative References

TOC

[I-D.ietf-mpls-tp-framework]	Bocci, M., Bryant, S., Frost, D., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks," draft-ietf-mpls-tp-framework-12 (work in progress), May 2010 (TXT).
[IEEE1588]	IEEE, "1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems," March 2008.
[RFC1305]	Mills, D., "Network Time Protocol (Version 3) Specification, Implementation," RFC 1305, March 1992 (TXT, PDF).
[RFC2030]	Mills, D., "Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI," RFC 2030, October 1996 (TXT).
[RFC2679]	Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," RFC 2679, September 1999 (TXT).
[RFC2680]	Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," RFC 2680, September 1999 (TXT).
[RFC2681]	Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM," RFC 2681, September 1999 (TXT).
[RFC3270]	Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services," RFC 3270, May 2002 (TXT).
[RFC3393]	Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)," RFC 3393, November 2002 (TXT).
[RFC5462]	Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field," RFC 5462, February 2009 (TXT).

Authors' Addresses

TOC

	Dan Frost (editor)
	Cisco Systems
Email:	danfrost@cisco.com
	Stewart Bryant (editor)
	Cisco Systems

Email:	stbryant@cisco.com