

Internet Engineering Task Force (IETF)
Request for Comments: 7759
Category: Standards Track
ISSN: 2070-1721

E. Bellagamba

G. Mirsky
Ericsson
L. Andersson
Huawei Technologies
P. Skoldstrom
Acreo AB
D. Ward
Cisco
J. Drake
Juniper
February 2016

Configuration of Proactive Operations,
Administration, and Maintenance (OAM) Functions for MPLS-Based
Transport Networks Using Label Switched Path (LSP) Ping

Abstract

This specification describes the configuration of proactive MPLS-TP Operations, Administration, and Maintenance (OAM) functions for a given Label Switched Path (LSP) using a set of TLVs that are carried by the LSP Ping protocol.

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/rfc7759>.

Copyright Notice

Copyright (c) 2016 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
1.1. Conventions Used in This Document	4
1.1.1. Terminology	4
1.1.2. Requirements Language	5
2. Theory of Operations	5
2.1. MPLS OAM Configuration Operation Overview	5
2.1.1. Configuration of BFD Sessions	5
2.1.2. Configuration of Performance Monitoring	6
2.1.3. Configuration of Fault Management Signals	6
2.2. MPLS OAM Functions TLV	7
2.2.1. BFD Configuration Sub-TLV	9
2.2.2. BFD Local Discriminator Sub-TLV	11
2.2.3. BFD Negotiation Timer Parameters Sub-TLV	11
2.2.4. BFD Authentication Sub-TLV	13
2.2.5. Traffic Class Sub-TLV	14
2.2.6. Performance Monitoring Sub-TLV	14
2.2.7. PM Loss Measurement Sub-TLV	17
2.2.8. PM Delay Measurement Sub-TLV	18
2.2.9. Fault Management Signal Sub-TLV	20
2.2.10. Source MEP-ID Sub-TLV	21
3. Summary of MPLS OAM Configuration Errors	22
4. IANA Considerations	23
4.1. TLV and Sub-TLV Allocation	23
4.2. MPLS OAM Function Flags Allocation	24
4.3. OAM Configuration Errors	25
5. Security Considerations	26
6. References	26
6.1. Normative References	26
6.2. Informative References	27
Acknowledgements	28
Authors' Addresses	29

1. Introduction

The MPLS Transport Profile (MPLS-TP) describes a profile of MPLS that enables operational models typical in transport networks while providing additional Operations, Administration, and Maintenance (OAM), survivability, and other maintenance functions not currently supported by MPLS. [RFC5860] defines the requirements for the OAM functionality of MPLS-TP.

This document describes the configuration of proactive MPLS-TP OAM functions for a given Label Switched Path (LSP) using TLVs carried in LSP Ping [RFC4379]. In particular, it specifies the mechanisms necessary to establish MPLS-TP OAM entities at the maintenance points for monitoring and performing measurements on an LSP, as well as defining information elements and procedures to configure proactive MPLS-TP OAM functions running between Label Edge Routers (LERs). Initialization and control of on-demand MPLS-TP OAM functions are expected to be carried out by directly accessing network nodes via a management interface; hence, configuration and control of on-demand OAM functions are out of scope for this document.

The Transport Profile of MPLS must, by definition [RFC5654], be capable of operating without a control plane. Therefore, there are a few options for configuring MPLS-TP OAM: without a control plane using a Network Management System (NMS), implementing LSP Ping instead or with a control plane implementing extensions to signaling protocols RSVP Traffic Engineering (RSVP-TE) [RFC3209] and/or Targeted LDP [RFC5036].

Proactive MPLS-TP OAM is performed by a set of protocols: Bidirectional Forwarding Detection (BFD) [RFC6428] for Continuity Check/Connectivity Verification, the Delay Measurement (DM) protocol [RFC6374], [RFC6375] for delay and delay variation (jitter) measurements, and the Loss Measurement (LM) protocol [RFC6374], [RFC6375] for packet loss and throughput measurements. Additionally, there are a number of Fault Management Signals that can be configured [RFC6427].

BFD is a protocol that provides low-overhead, fast detection of failures in the path between two forwarding engines, including the interfaces, data link(s), and to the extent possible, the forwarding engines themselves. BFD can be used to detect the continuity and mis-connection defects of MPLS-TP point-to-point and might also be extended to support point-to-multipoint LSPs.

The delay and loss measurements protocols [RFC6374] and [RFC6375] use a simple query/response model for performing both unidirectional and bidirectional measurements that allow the originating node to measure

packet loss and delay in forward, or forward and reverse directions. By timestamping and/or writing current packet counters to the measurement packets (four times, Transmit and Receive in both directions), current delays and packet losses can be calculated. By performing successive delay measurements, the delay and/or inter-packet delay variation (jitter) can be calculated. Current throughput can be calculated from the packet loss measurements by dividing the number of packets sent/received with the time it took to perform the measurement, given by the timestamp in the LM header. Combined with a packet generator, the throughput measurement can be used to measure the maximum capacity of a particular LSP. It should be noted that this document does not specify how to configure on-demand throughput estimates based on saturating the connection as defined in [RFC6371]; rather, it only specifies how to enable the estimation of the current throughput based on loss measurements.

1.1. Conventions Used in This Document

1.1.1. Terminology

BFD - Bidirectional Forwarding Detection

DM - Delay Measurement

FMS - Fault Management Signal

G-ACh - Generic Associated Channel

LSP - Label Switched Path

LM - Loss Measurement

MEP - Maintenance Entity Group End Point

MPLS - Multi-Protocol Label Switching

MPLS-TP - MPLS Transport Profile

NMS - Network Management System

PM - Performance Monitoring

RSVP-TE - RSVP Traffic Engineering

TC - Traffic Class

1.1.2. 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 [RFC2119].

2. Theory of Operations

2.1. MPLS OAM Configuration Operation Overview

The MPLS-TP OAM tool set is described in [RFC6669].

LSP Ping, or alternatively RSVP-TE [RFC7487], can be used to easily enable the different OAM functions by setting the corresponding flags in the MPLS OAM Functions TLV (refer to Section 2.2). For a more detailed configuration, one may include sub-TLVs for the different OAM functions in order to specify various parameters in detail.

Typically, intermediate nodes simply forward OAM configuration TLVs to the end node without any processing or modification. At least one exception to this is if the FMS sub-TLV (refer to Section 2.2.9) is present. This sub-TLV MUST be examined even by intermediate nodes that support this extension. The sub-TLV MAY be present if a flag is set in the MPLS OAM Functions TLV.

2.1.1. Configuration of BFD Sessions

For this specification, BFD MUST run in either one of the two modes:

- o Asynchronous mode, where both sides are in active mode
- o Unidirectional mode

In the simplest scenario, LSP Ping [RFC5884], or alternatively RSVP-TE [RFC7487], is used only to bootstrap a BFD session for an LSP, without any timer negotiation.

Timer negotiation can be performed either in subsequent BFD control messages (in this case the operation is similar to bootstrapping based on LSP Ping described in [RFC5884]), or directly in the LSP Ping configuration messages.

When BFD Control packets are transported in the Associated Channel Header (ACH) encapsulation, they are not protected by any end-to-end checksum; only lower layers provide error detection/correction. A single bit error, e.g., a flipped bit in the BFD State field, could cause the receiving end to wrongly conclude that the link is down and in turn trigger protection switching. To prevent this from

happening, the BFD Configuration sub-TLV (refer to Section 2.2.1) has an Integrity flag that, when set, enables BFD Authentication using Keyed SHA1 with an empty key (all 0s) [RFC5880]. This would make every BFD Control packet carry a SHA1 hash of itself that can be used to detect errors.

If BFD Authentication using a pre-shared key/password is desired (i.e., authentication and not only error detection), the BFD Authentication sub-TLV (refer to Section 2.2.4) MUST be included in the BFD Configuration sub-TLV. The BFD Authentication sub-TLV is used to specify which authentication method that should be used and which pre-shared key/password that should be used for this particular session. How the key exchange is performed is out of scope of this document.

2.1.2. Configuration of Performance Monitoring

It is possible to configure Performance Monitoring functionalities such as Loss, Delay, Delay/Interpacket Delay variation (jitter), and throughput as described in [RFC6374].

When configuring Performance Monitoring functionalities, it is possible to choose either the default configuration, by only setting the respective flags in the MPLS OAM functions TLV, or a customized configuration. To customize the configuration, one would set the respective flags in the MPLS OAM functions TLV and include the respective Loss and/or Delay sub-TLVs.

By setting the PM Loss flag in the MPLS OAM Functions TLV and including the PM Loss sub-TLV (refer to Section 2.2.7), one can configure the measurement interval and loss threshold values for triggering protection.

Delay measurements are configured by setting the PM Delay flag in the MPLS OAM Functions TLV and by including the PM Delay sub-TLV (refer to Section 2.2.8), one can configure the measurement interval and the delay threshold values for triggering protection.

2.1.3. Configuration of Fault Management Signals

To configure Fault Management Signals (FMSs) and their refresh time, the FMS Flag in the MPLS OAM Functions TLV MUST be set and the FMS sub-TLV MUST be included. When configuring an FMS, an implementation can enable the default configuration by setting the FMS Flag in the OAM Function Flags sub-TLV. In order to modify the default configuration, the MPLS OAM FMS sub-TLV MUST be included.

If an intermediate point is meant to originate FMS messages, this means that such an intermediate point is associated with a Server MEP through a co-located MPLS-TP client/server adaptation function, and the Fault Management subscription flag in the MPLS OAM FMS sub-TLV has been set as an indication of the request to create the association at each intermediate node of the client LSP. The corresponding Server MEP needs to be configured by its own LSP Ping session or, alternatively, via a Network Management System (NMS) or RSVP-TE.

2.2. MPLS OAM Functions TLV

The MPLS OAM Functions TLV presented in Figure 1 is carried as a TLV of the MPLS Echo Request/Reply messages [RFC4379].

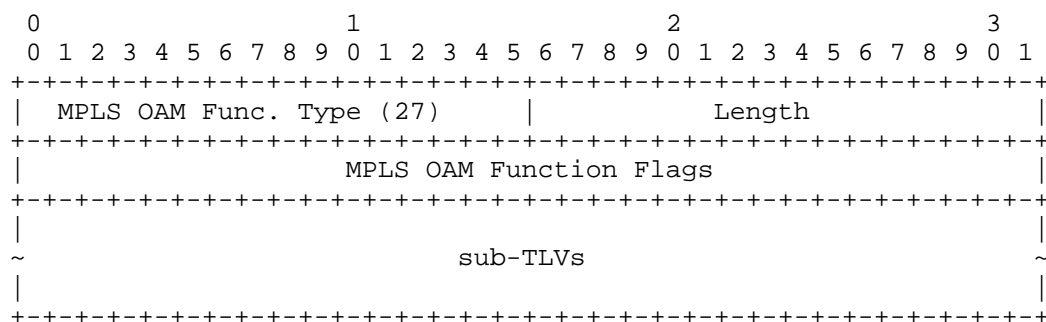


Figure 1: MPLS OAM Functions TLV Format

The MPLS OAM Functions TLV contains the MPLS OAM Function Flags field. The MPLS OAM Function Flags indicate which OAM functions should be activated as well as OAM function-specific sub-TLVs with configuration parameters for the particular function.

Type: Indicates the MPLS OAM Functions TLV (Section 4).

Length: The length of the MPLS OAM Function Flags field including the total length of the sub-TLVs in octets.

MPLS OAM Function Flags: A bitmap numbered from left to right as shown in Figure 2. These flags are managed by IANA (refer to Section 4.2). Flags defined in this document are presented in Table 2. Undefined flags MUST be set to zero and unknown flags MUST be ignored. The flags indicate what OAM is being configured and direct the presence of optional sub-TLVs as set out below.




Figure 2: MPLS OAM Function Flags Format

Sub-TLVs corresponding to the different flags are as follows. No meaning should be attached to the order of sub-TLVs.

- o If a flag in the MPLS OAM Function Flags is set and the corresponding sub-TLVs listed below are absent, then this MPLS OAM function MUST be initialized according to its default settings. Default settings of MPLS OAM functions are outside the scope of this document.
- o If any sub-TLV is present without the corresponding flag being set, the sub-TLV SHOULD be ignored.
- o BFD Configuration sub-TLV, which MUST be included if either the CC, the CV, or both MPLS OAM Function flags are being set in the MPLS OAM Functions TLV.
- o Performance Monitoring sub-TLV MUST be used to carry PM Loss sub-TLV and/or PM Delay sub-TLV. If neither one of these sub-TLVs is present, then Performance Monitoring sub-TLV SHOULD NOT be included. Empty, i.e., no enclosed sub-TLVs, Performance Monitoring sub-TLV SHOULD be ignored.
- o PM Loss sub-TLV MAY be included if the PM/Loss OAM Function flag is set. If the "PM Loss sub-TLV" is not included, default configuration values are used. Such sub-TLV MAY also be included in case the Throughput function flag is set and there is the need to specify a measurement interval different from the default ones. In fact, the throughput measurement makes use of the same tool as the loss measurement; hence, the same TLV is used.
- o PM Delay sub-TLV MAY be included if the PM/Delay OAM Function flag is set. If the "PM Delay sub-TLV" is not included, default configuration values are used.
- o FMS sub-TLV, that MAY be included if the FMS OAM Function flag is set. If the "FMS sub-TLV" is not included, default configuration values are used.

If all flags in the MPLS OAM Function Flags field have the same value of zero, that MUST be interpreted as meaning that the MPLS OAM Functions TLV is not present in the MPLS Echo Request. If more than one MPLS OAM Functions TLV is present in the MPLS Echo request packet, then the first TLV SHOULD be processed and the rest ignored. Any parsing error within nested sub-TLVs that is not specified in Section 3 SHOULD be treated as described in [RFC4379].

2.2.1. BFD Configuration Sub-TLV

The BFD Configuration sub-TLV, depicted in Figure 3, is defined for BFD OAM-specific configuration parameters. The "BFD Configuration sub-TLV" is carried as a sub-TLV of the "OAM Functions TLV".

This TLV accommodates generic BFD OAM information and carries sub-TLVs.

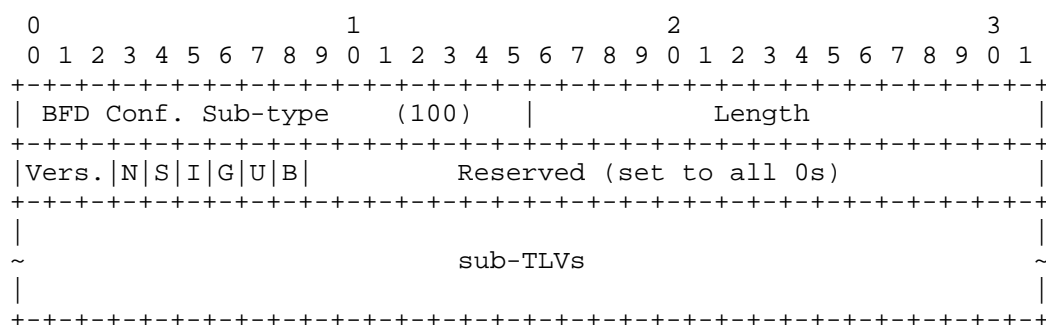


Figure 3: BFD Configuration Sub-TLV Format

Sub-type: Indicates a new sub-type, the BFD Configuration sub-TLV (value 100).

Length: Indicates the length of the Value field in octets.

Version: Identifies the BFD protocol version. If a node does not support a specific BFD version, an error must be generated: "OAM Problem/Unsupported BFD Version".

BFD Negotiation (N): If set, timer negotiation/renegotiation via BFD Control Messages is enabled. When cleared, it is disabled and timer configuration is achieved using the BFD Negotiation Timer Parameters sub-TLV as described in Section 2.2.3.

Symmetric session (S): If set, the BFD session MUST use symmetric timing values. If cleared, the BFD session MAY use any timing values either negotiated or explicitly configured.

Integrity (I): If set, BFD Authentication MUST be enabled. If the BFD Configuration sub-TLV does not include a BFD Authentication sub-TLV, the authentication MUST use Keyed SHA1 with an empty pre-shared key (all 0s). If the egress LSR does not support BFD Authentication, an error MUST be generated: "OAM Problem/BFD Authentication unsupported". If the Integrity flag is clear, then Authentication MUST NOT be used.

Encapsulation Capability (G): If set, it shows the capability of encapsulating BFD messages into the G-ACh channel. If both the G bit and U bit are set, configuration gives precedence to the G bit.

Encapsulation Capability (U): If set, it shows the capability of encapsulating BFD messages into IP/UDP packets. If both the G bit and U bit are set, configuration gives precedence to the G bit.

If the egress LSR does not support any of the ingress LSR Encapsulation Capabilities, an error MUST be generated: "OAM Problem/Unsupported BFD Encapsulation format".

Bidirectional (B): If set, it configures BFD in the Bidirectional mode. If it is not set, it configures BFD in the unidirectional mode. In the second case, the source node does not expect any Discriminator values back from the destination node.

Reserved: Reserved for future specification; set to 0 on transmission and ignored when received.

The BFD Configuration sub-TLV MUST include the following sub-TLVs in the MPLS Echo Request message:

- o BFD Local Discriminator sub-TLV, if the B flag is set in the MPLS Echo Request;
- o BFD Negotiation Timer Parameters sub-TLV, if the N flag is cleared.

The BFD Configuration sub-TLV MUST include the following sub-TLVs in the MPLS Echo Reply message:

- o BFD Local Discriminator sub-TLV;
- o BFD Negotiation Timer Parameters sub-TLV if:
 - * The N and S flags are cleared, or if:
 - * The N flag is cleared and the S flag is set, and the BFD Negotiation Timer Parameters sub-TLV received by the egress

contains unsupported values. In this case, an updated BFD Negotiation Timer Parameters sub-TLV, containing values supported by the egress node [RFC7419], is returned to the ingress.

2.2.2. BFD Local Discriminator Sub-TLV

The BFD Local Discriminator sub-TLV is carried as a sub-TLV of the "BFD Configuration sub-TLV" and is depicted in Figure 4.




Figure 4: BFD Local Discriminator Sub-TLV Format

Sub-type: Indicates a new sub-type, the "BFD Local Discriminator sub-TLV" (value 101).

Length: Indicates the length of the Value field in octets (4).

Local Discriminator: A nonzero discriminator value that is unique in the context of the transmitting system that generates it. It is used to demultiplex multiple BFD sessions between the same pair of systems.

2.2.3. BFD Negotiation Timer Parameters Sub-TLV

The BFD Negotiation Timer Parameters sub-TLV is carried as a sub-TLV of the BFD Configuration sub-TLV and is depicted in Figure 5.

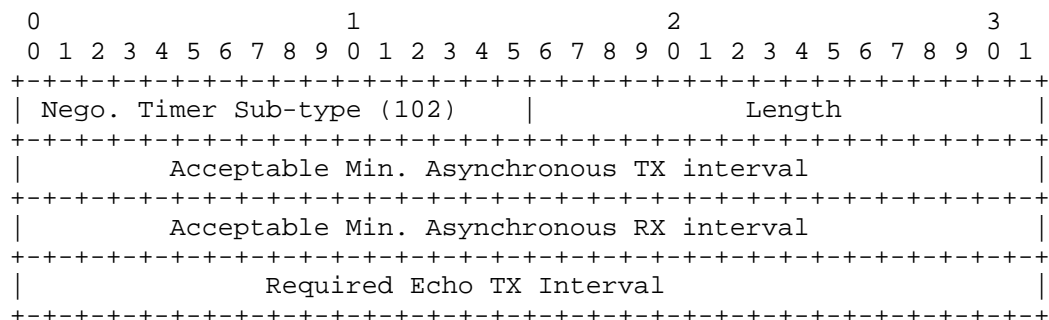


Figure 5: BFD Negotiation Timer Parameters Sub-TLV Format

Sub-type: Indicates a new sub-type, the BFD Negotiation Timer Parameters sub-TLV (value 102).

Length: Indicates the length of the Value field in octets (12).

Acceptable Min. Asynchronous TX interval: If the S (symmetric) flag is set in the BFD Configuration sub-TLV, defined in Section 2.2.1, it expresses the desired time interval (in microseconds) at which the ingress LER intends to both transmit and receive BFD periodic control packets. If the receiving edge LSR cannot support such a value, it SHOULD reply with an interval greater than the one proposed.

If the S (symmetric) flag is cleared in the BFD Configuration sub-TLV, this field expresses the desired time interval (in microseconds) at which an edge LSR intends to transmit BFD periodic control packets in its transmitting direction.

Acceptable Min. Asynchronous RX interval: If the S (symmetric) flag is set in the BFD Configuration sub-TLV, Figure 3, this field MUST be equal to Acceptable Min. Asynchronous TX interval and has no additional meaning respect to the one described for "Acceptable Min. Asynchronous TX interval".

If the S (symmetric) flag is cleared in the BFD Configuration sub-TLV, it expresses the minimum time interval (in microseconds) at which edge LSRs can receive BFD periodic control packets. If this value is greater than the value of Acceptable Min. Asynchronous TX interval received from the other edge LSR, such an edge LSR MUST adopt the interval expressed in this Acceptable Min. Asynchronous RX interval.

Required Echo TX Interval: The minimum interval (in microseconds) between received BFD Echo packets that this system is capable of supporting, less any jitter applied by the sender as described in Section 6.8.9 of [RFC5880]. This value is also an indication for the receiving system of the minimum interval between transmitted BFD Echo packets. If this value is zero, the transmitting system does not support the receipt of BFD Echo packets. If the receiving system cannot support this value, the "Unsupported BFD TX Echo rate interval" error MUST be generated. By default, the value is set to 0.

2.2.4. BFD Authentication Sub-TLV

The "BFD Authentication sub-TLV" is carried as a sub-TLV of the "BFD Configuration sub-TLV" and is depicted in Figure 6.




Figure 6: BFD Authentication Sub-TLV Format

Sub-type: Indicates a new sub-type, the BFD Authentication sub-TLV (value 103).

Length: Indicates the length of the Value field in octets (4).

Auth Type: Indicates which type of authentication to use. The same values as are defined in Section 4.1 of [RFC5880] are used. Simple Password SHOULD NOT be used if other authentication types are available.

Auth Key ID: Indicates which authentication key or password (depending on Auth Type) should be used. How the key exchange is performed is out of scope of this document. If the egress LSR does not support this Auth Key ID, an "OAM Problem/Mismatch of BFD Authentication Key ID" error MUST be generated.

Reserved: Reserved for future specification; set to 0 on transmission and ignored when received.

An implementation MAY change the mode of authentication if an operator re-evaluates the security situation in and around the administrative domain. If the BFD Authentication sub-TLV is used for a BFD session in Up state, then the Sender of the MPLS LSP Echo Request SHOULD ensure that old and new modes of authentication, i.e., a combination of Auth.Type and Auth. Key ID, are used to send and receive BFD control packets, until the Sender can confirm that its peer has switched to the new authentication.

2.2.5. Traffic Class Sub-TLV

The Traffic Class sub-TLV is carried as a sub-TLV of the "BFD Configuration sub-TLV" and "Fault Management Signal Sub-TLV" (Section 2.2.9) and is depicted in Figure 7.




Figure 7: Traffic Class Sub-TLV Format

Sub-type: Indicates a new sub-type, the "Traffic Class sub-TLV" (value 104).

Length: Indicates the length of the Value field in octets (4).

TC: Identifies the Traffic Class (TC) [RFC5462] for periodic continuity monitoring messages or packets with fault management information.

If the TC sub-TLV is present, then the sender of any periodic continuity monitoring messages or packets with fault management information on the LSP, with a Forwarding Equivalence Class (FEC) that corresponds to the FEC for which fault detection is being performed, MUST use the value contained in the TC field of the sub-TLV as the value of the TC field in the top label stack entry of the MPLS label stack. If the TC sub-TLV is absent from either "BFD Configuration sub-TLV" or "Fault Management Signal sub-TLV", then selection of the TC value is a local decision.

2.2.6. Performance Monitoring Sub-TLV

If the MPLS OAM Functions TLV has any of the L (Loss), D (Delay), and T (Throughput) flags set, the Performance Monitoring sub-TLV MUST be present. Failure to include the correct sub-TLVs MUST result in an "OAM Problem/PM Configuration Error" being generated.

The Performance Monitoring sub-TLV provides the configuration information mentioned in Section 7 of [RFC6374]. It includes support for the configuration of quality thresholds and, as described in [RFC6374]:

...the crossing of which will trigger warnings or alarms, and result in reporting and exception notification will be integrated into the system-wide network management and reporting framework.

In case the values need to be different than the default ones, the Performance Monitoring sub-TLV MAY include the following sub-TLVs:

- o PM Loss sub-TLV, if the L flag is set in the MPLS OAM Functions TLV;
- o PM Delay sub-TLV, if the D flag is set in the MPLS OAM Functions TLV.

The Performance Monitoring sub-TLV depicted in Figure 8 is carried as a sub-TLV of the MPLS OAM Functions TLV.




Figure 8: Performance Monitoring Sub-TLV Format

Sub-type: Indicates a new sub-type, the Performance Monitoring sub-TLV (value 200).

Length: Indicates the length of the Value field in octets, including PM Configuration Flags and optional sub-TLVs.




Figure 9: PM Configuration Flags Format

The PM Configuration Flags format is presented in Figure 9. For the specific function description, please refer to [RFC6374]:

- D: Delay inferred/direct (0=INFERRED, 1=DIRECT). If the egress LSR does not support the specified mode, an "OAM Problem/Unsupported Delay Mode" error MUST be generated.
- L: Loss inferred/direct (0=INFERRED, 1=DIRECT). If the egress LSR does not support the specified mode, an "OAM Problem/Unsupported Loss Mode" error MUST be generated.
- J: Delay variation/jitter (1=ACTIVE, 0=NOT ACTIVE). If the egress LSR does not support Delay variation measurements and the J flag is set, an "OAM Problem/Delay variation unsupported" error MUST be generated.
- Y: Dyadic (1=ACTIVE, 0=NOT ACTIVE). If the egress LSR does not support Dyadic mode and the Y flag is set, an "OAM Problem/Dyadic mode unsupported" error MUST be generated.
- K: Loopback (1=ACTIVE, 0=NOT ACTIVE). If the egress LSR does not support Loopback mode and the K flag is set, an "OAM Problem/Loopback mode unsupported" error MUST be generated.
- C: Combined (1=ACTIVE, 0=NOT ACTIVE). If the egress LSR does not support Combined mode and the C flag is set, an "OAM Problem/Combined mode unsupported" error MUST be generated.
- Reserved: Reserved for future specification; set to 0 on transmission and ignored when received.

2.2.7. PM Loss Measurement Sub-TLV

The PM Loss Measurement sub-TLV depicted in Figure 10 is carried as a sub-TLV of the Performance Monitoring sub-TLV.

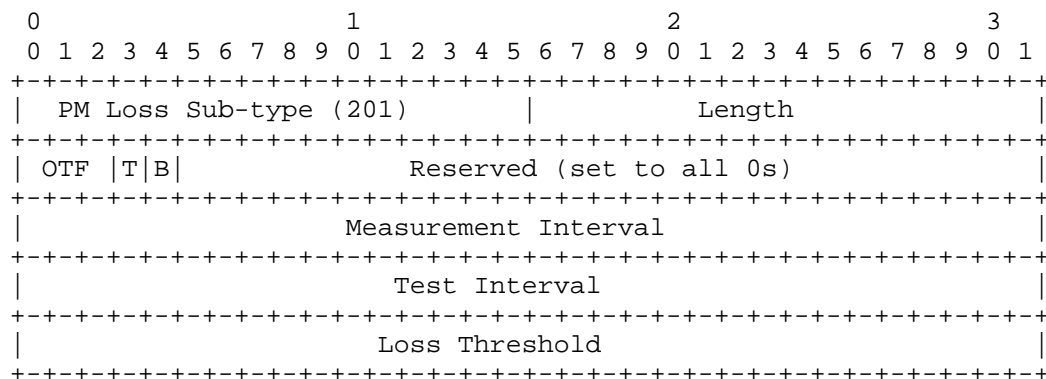


Figure 10: PM Loss Measurement Sub-TLV Format

Sub-type: Indicates a new sub-type, the PM Loss Measurement sub-TLV (value 201).

Length: Indicates the length of the Value field in octets (16).

OTF: Origin Timestamp Format of the Origin Timestamp field described in [RFC6374]. By default, it is set to IEEE 1588 version 1. If the egress LSR cannot support this value, an "OAM Problem/Unsupported Timestamp Format" error MUST be generated.

Configuration Flags, please refer to [RFC6374] for further details:

T: Traffic-class-specific measurement indicator. Set to 1 when the measurement operation is scoped to packets of a particular traffic class (Differentiated Services Code Point value), and 0 otherwise. When set to 1, the Differentiated Services (DS) field of the message indicates the measured traffic class. By default, it is set to 1.

B: Octet (byte) count. When set to 1, indicates that the Counter 1-4 fields represent octet counts. When set to 0, indicates that the Counter 1-4 fields represent packet counts. By default, it is set to 0.

Reserved: Reserved for future specification; set to 0 on transmission and ignored when received.

Measurement Interval: The time interval (in milliseconds) at which Loss Measurement query messages MUST be sent on both directions. If the edge LSR receiving the Path message cannot support such a value, it SHOULD reply with a higher interval. By default, it is set to (100) as per [RFC6375].

Test Interval: Test messages interval in milliseconds as described in [RFC6374]. By default, it is set to (10) as per [RFC6375].

Loss Threshold: The threshold value of measured lost packets per measurement over which action(s) SHOULD be triggered.

2.2.8. PM Delay Measurement Sub-TLV

The "PM Delay Measurement sub-TLV" depicted in Figure 11 is carried as a sub-TLV of the Performance Monitoring sub-TLV.




Figure 11: PM Delay Measurement Sub-TLV Format

Sub-type: Indicates a new sub-type, the "PM Delay Measurement sub-TLV" (value 202).

Length: Indicates the length of the Value field in octets (16).

OTF: Origin Timestamp Format of the Origin Timestamp field described in [RFC6374]. By default, it is set to IEEE 1588 version 1. If the egress LSR cannot support this value, an "OAM Problem/Unsupported Timestamp Format" error MUST be generated.

Configuration Flags, please refer to [RFC6374] for further details:

- T: Traffic-class-specific measurement indicator. Set to 1 when the measurement operation is scoped to packets of a particular traffic class (Differentiated Services Code Point value), and 0 otherwise. When set to 1, the DS field of the message indicates the measured traffic class. By default, it is set to 1.
- B: Octet (byte) count. When set to 1, indicates that the Counter 1-4 fields represent octet counts. When set to 0, indicates that the Counter 1-4 fields represent packet counts. By default, it is set to 0.

Reserved: Reserved for future specification; set to 0 on transmission and ignored when received.

Measurement Interval: The time interval (in milliseconds) at which Delay Measurement query messages MUST be sent on both directions. If the edge LSR receiving the Path message cannot support such a value, it can reply with a higher interval. By default, it is set to (1000) as per [RFC6375].

Test Interval: Test messages interval (in milliseconds) as described in [RFC6374]. By default, it is set to (10) as per [RFC6375].

Delay Threshold: The threshold value of measured two-way delay (in milliseconds) over which action(s) SHOULD be triggered.

2.2.9. Fault Management Signal Sub-TLV

The FMS sub-TLV depicted in Figure 12 is carried as a sub-TLV of the MPLS OAM Configuration sub-TLV. When both working and protection paths are configured, both LSPs SHOULD be configured with identical settings of the E flag, T flag, and the refresh timer. An implementation MAY configure the working and protection LSPs with different settings of these fields in case of 1:N protection.

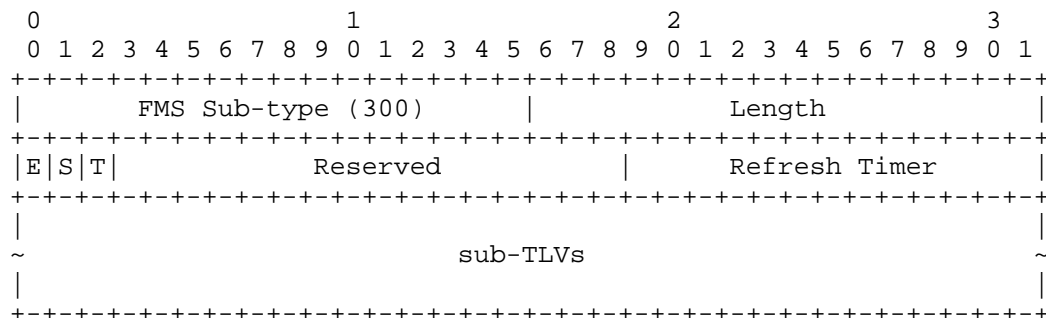


Figure 12: Fault Management Signal Sub-TLV Format

Sub-type: Indicates a new sub-type, the FMS sub-TLV (value 300).

Length: Indicates the length of the Value field in octets.

FMS Flags are used to enable the FMS Flags at end point MEPs and the Server MEPs of the links over which the LSP is forwarded. In this document, only the S flag pertains to Server MEPs.

The following flags are defined:

- E: Enable Alarm Indication Signal (AIS) and Lock Report (LKR) signaling as described in [RFC6427]. Default value is 1 (enabled). If the egress MEP does not support FMS Flag generation, an "OAM Problem/Fault management signaling unsupported" error MUST be generated.
- S: Indicate to a Server MEP that it should transmit AIS and LKR signals on the client LSP. Default value is 0 (disabled). If a Server MEP that is capable of generating FMS messages is, for some reason, unable to do so for the LSP being signaled, an "OAM Problem/Unable to create fault management association" error MUST be generated.
- T: Set timer value, enabled the configuration of a specific timer value. Default value is 0 (disabled).

Reserved: Bits 4-16 that follow the FMS Flags are reserved for future allocation. These bits MUST be set to 0 on transmit and ignored on receipt if not allocated.

Refresh Timer: Indicates the refresh timer of fault indication messages, in seconds. The value MUST be between 1 to 20 seconds as specified for the Refresh Timer field in [RFC6427]. If the edge LSR receiving the Path message cannot support the value, it SHOULD reply with a higher timer value.

FMS sub-TLV MAY include Traffic Class sub-TLV (Section 2.2.5). If the TC sub-TLV is present, the value of the TC field MUST be used as the value of the TC field of an MPLS label stack entry for FMS messages. If the TC sub-TLV is absent, then selection of the TC value is a local decision.

2.2.10. Source MEP-ID Sub-TLV

The Source MEP-ID sub-TLV depicted in Figure 13 is carried as a sub-TLV of the MPLS OAM Functions TLV.

Note that support of ITU IDs is out of scope.

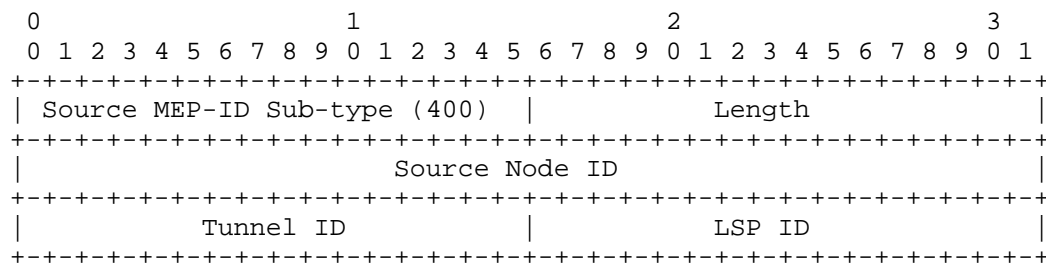


Figure 13: Source MEP-ID Sub-TLV Format

Sub-type: Indicates a new sub-type, the Source MEP-ID sub-TLV (value 400).

Length: Indicates the length of the Value field in octets (8).

Source Node ID: 32-bit node identifier as defined in [RFC6370].

Tunnel ID: A 16-bit unsigned integer unique to the node as defined in [RFC6370].

LSP ID: A 16-bit unsigned integer unique within the Tunnel_ID as defined in [RFC6370].

3. Summary of MPLS OAM Configuration Errors

This is the summary of Return Codes [RFC4379] defined in this document:

- o If an egress LSR does not support the specified BFD version, an error MUST be generated: "OAM Problem/Unsupported BFD Version".
- o If an egress LSR does not support the specified BFD Encapsulation format, an error MUST be generated: "OAM Problem/Unsupported BFD Encapsulation format".
- o If an egress LSR does not support BFD Authentication, and it is requested, an error MUST be generated: "OAM Problem/BFD Authentication unsupported".
- o If an egress LSR does not support the specified BFD Authentication Type, an error MUST be generated: "OAM Problem/Unsupported BFD Authentication Type".
- o If an egress LSR is not able to use the specified Authentication Key ID, an error MUST be generated: "OAM Problem/Mismatch of BFD Authentication Key ID".
- o If PM flags in MPLS OAM Functions TLV don't have corresponding PM sub-TLVs present, an error MUST be generated: "OAM Problem/PM Configuration Error".
- o If an egress LSR does not support the specified Timestamp Format, an error MUST be generated: "OAM Problem/Unsupported Timestamp Format".
- o If an egress LSR does not support specified Delay mode, an "OAM Problem/Unsupported Delay Mode" error MUST be generated.
- o If an egress LSR does not support specified Loss mode, an "OAM Problem/Unsupported Loss Mode" error MUST be generated.
- o If an egress LSR does not support Delay variation measurements, and it is requested, an "OAM Problem/Delay variation unsupported" error MUST be generated.
- o If an egress LSR does not support Dyadic mode, and it is requested, an "OAM Problem/Dyadic mode unsupported" error MUST be generated.

- o If an egress LSR does not support Loopback mode, and it is requested, an "OAM Problem/Loopback mode unsupported" error MUST be generated.
- o If an egress LSR does not support Combined mode, and it is requested, an "OAM Problem/Combined mode unsupported" error MUST be generated.
- o If an egress LSR does not support Fault Monitoring Signals, and it is requested, an "OAM Problem/Fault management signaling unsupported" error MUST be generated.
- o If an intermediate Server MEP supports Fault Monitoring Signals, but is unable to create an association, when requested to do so, an "OAM Problem/Unable to create fault management association" error MUST be generated.

Ingress LSR MAY combine multiple MPLS OAM configuration TLVs and sub-TLVs into single MPLS echo request. In case an egress LSR doesn't support any of the requested modes, it MUST set the return code to report the first unsupported mode in the list of TLVs and sub-TLVs. And if any of the requested OAM configuration is not supported, the egress LSR SHOULD NOT process OAM Configuration TLVs and sub-TLVs listed in the MPLS echo request.

4. IANA Considerations

4.1. TLV and Sub-TLV Allocation

IANA maintains the "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry and, within that registry, a subregistry for TLVs and sub-TLVs.

IANA has allocated a new MPLS OAM Functions TLV from the Standards Action [RFC5226] range (0-16383) and sub-TLVs as follows from subregistry presented in Table 1, called "Sub-TLVs for TLV Type 27".

Registration procedures for Sub-TLVs from ranges 0-16383 and 32768-49161 are by Standards Action. Ranges 16384-31743 and 49162-64511 are through Specification Required (Experimental RFC Needed).

Type	Sub-type	Value Field	Reference
27		MPLS OAM Functions	This document
	100	BFD Configuration	This document
	101	BFD Local Discriminator	This document
	102	BFD Negotiation Timer Parameters	This document
	103	BFD Authentication	This document
	104	Traffic Class	This document
	200	Performance Monitoring	This document
	201	PM Loss Measurement	This document
	202	PM Delay Measurement	This document
	300	Fault Management Signal	This document
	400	Source MEP-ID	This document

Table 1: IANA TLV Type Allocation

4.2. MPLS OAM Function Flags Allocation

IANA has created a new registry called the "MPLS OAM Function Flags" registry. Assignments of bit positions 0 through 31 are via Standards Action. The new registry is to be populated as follows.

Bit Position	MPLS OAM Function Flag	Description
0	C	Continuity Check (CC)
1	V	Connectivity Verification (CV)
2	F	Fault Management Signal (FMS)
3	L	Performance Monitoring/Loss (PM/Loss)
4	D	Performance Monitoring/Delay (PM/Delay)
5	T	Throughput Measurement
6-30		Unassigned (Must be zero)
31		Reserved

Table 2: MPLS OAM Function Flags

4.3. OAM Configuration Errors

IANA maintains a registry "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters", and within that registry a subregistry "Return Codes".

IANA has assigned new Return Codes from the Standards Action range (0-191) as follows:

Error Value Sub-codes	Description	Reference
21	OAM Problem/Unsupported BFD Version	This document
22	OAM Problem/Unsupported BFD Encapsulation format	This document
23	OAM Problem/Unsupported BFD Authentication Type	This document
24	OAM Problem/Mismatch of BFD Authentication Key ID	This document
25	OAM Problem/Unsupported Timestamp Format	This document
26	OAM Problem/Unsupported Delay Mode	This document
27	OAM Problem/Unsupported Loss Mode	This document
28	OAM Problem/Delay variation unsupported	This document
29	OAM Problem/Dyadic mode unsupported	This document
30	OAM Problem/Loopback mode unsupported	This document
31	OAM Problem/Combined mode unsupported	This document
32	OAM Problem/Fault management signaling unsupported	This document
33	OAM Problem/Unable to create fault management association	This document
34	OAM Problem/PM Configuration Error	This document

Table 3: IANA Return Codes Allocation

5. Security Considerations

The signaling of OAM-related parameters and the automatic establishment of OAM entities introduces additional security considerations to those discussed in [RFC4379]. In particular, a network element could be overloaded if an attacker were to request high-frequency liveliness monitoring of a large number of LSPs, targeting a single network element. Implementations must be made cognizant of available OAM resources and MAY refuse new OAM configurations that would overload a node. Additionally, policies to manage OAM resources may be used to provide some fairness in OAM resource distribution among monitored LSPs.

Security of OAM protocols configured with extensions to LSP Ping described in this document are discussed in [RFC5880], [RFC5884], [RFC6374], [RFC6427], and [RFC6428].

In order that the configuration of OAM functionality can be achieved securely through the techniques described in this document, security mechanisms must already be in place and operational for LSP Ping. Thus, the exchange of security parameters (such as keys) for use in securing OAM is outside the scope of this document and is assumed to use an off-line mechanism or an established secure key exchange protocol.

Additional discussion of security for MPLS protocols can be found in [RFC5920].

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, DOI 10.17487/RFC4379, February 2006, <<http://www.rfc-editor.org/info/rfc4379>>.
- [RFC5654] Niven-Jenkins, B., Ed., Brungard, D., Ed., Betts, M., Ed., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, DOI 10.17487/RFC5654, September 2009, <<http://www.rfc-editor.org/info/rfc5654>>.

- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<http://www.rfc-editor.org/info/rfc5880>>.
- [RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, DOI 10.17487/RFC5884, June 2010, <<http://www.rfc-editor.org/info/rfc5884>>.
- [RFC6370] Bocci, M., Swallow, G., and E. Gray, "MPLS Transport Profile (MPLS-TP) Identifiers", RFC 6370, DOI 10.17487/RFC6370, September 2011, <<http://www.rfc-editor.org/info/rfc6370>>.
- [RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, DOI 10.17487/RFC6374, September 2011, <<http://www.rfc-editor.org/info/rfc6374>>.
- [RFC6427] Swallow, G., Ed., Fulignoli, A., Ed., Vigoureux, M., Ed., Boutros, S., and D. Ward, "MPLS Fault Management Operations, Administration, and Maintenance (OAM)", RFC 6427, DOI 10.17487/RFC6427, November 2011, <<http://www.rfc-editor.org/info/rfc6427>>.
- [RFC6428] Allan, D., Ed., Swallow Ed., G., and J. Drake Ed., "Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile", RFC 6428, DOI 10.17487/RFC6428, November 2011, <<http://www.rfc-editor.org/info/rfc6428>>.

6.2. Informative References

- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<http://www.rfc-editor.org/info/rfc3209>>.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, DOI 10.17487/RFC5036, October 2007, <<http://www.rfc-editor.org/info/rfc5036>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.

- [RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, DOI 10.17487/RFC5462, February 2009, <<http://www.rfc-editor.org/info/rfc5462>>.
- [RFC5860] Vigoureux, M., Ed., Ward, D., Ed., and M. Betts, Ed., "Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks", RFC 5860, DOI 10.17487/RFC5860, May 2010, <<http://www.rfc-editor.org/info/rfc5860>>.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<http://www.rfc-editor.org/info/rfc5920>>.
- [RFC6371] Busi, I., Ed. and D. Allan, Ed., "Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks", RFC 6371, DOI 10.17487/RFC6371, September 2011, <<http://www.rfc-editor.org/info/rfc6371>>.
- [RFC6375] Frost, D., Ed. and S. Bryant, Ed., "A Packet Loss and Delay Measurement Profile for MPLS-Based Transport Networks", RFC 6375, DOI 10.17487/RFC6375, September 2011, <<http://www.rfc-editor.org/info/rfc6375>>.
- [RFC6669] Sprecher, N. and L. Fang, "An Overview of the Operations, Administration, and Maintenance (OAM) Toolset for MPLS-Based Transport Networks", RFC 6669, DOI 10.17487/RFC6669, July 2012, <<http://www.rfc-editor.org/info/rfc6669>>.
- [RFC7419] Akiya, N., Binderberger, M., and G. Mirsky, "Common Interval Support in Bidirectional Forwarding Detection", RFC 7419, DOI 10.17487/RFC7419, December 2014, <<http://www.rfc-editor.org/info/rfc7419>>.
- [RFC7487] Bellagamba, E., Takacs, A., Mirsky, G., Andersson, L., Skoldstrom, P., and D. Ward, "Configuration of Proactive Operations, Administration, and Maintenance (OAM) Functions for MPLS-Based Transport Networks Using RSVP-TE", RFC 7487, DOI 10.17487/RFC7487, March 2015, <<http://www.rfc-editor.org/info/rfc7487>>.

Acknowledgements

The authors would like to thank Nobo Akiya, David Allan, and Adrian Farrel for their thorough reviews and insightful comments.

Authors' Addresses

Elisa Bellagamba

Email: elisa.bellagamba@gmail.com

Gregory Mirsky
Ericsson

Email: Gregory.Mirsky@ericsson.com

Loa Andersson
Huawei Technologies

Email: loa@mail01.huawei.com

Pontus Skoldstrom
Acreo AB
Electrum 236
Kista 164 40
Sweden

Phone: +46 8 6327731
Email: pontus.skoldstrom@acreo.se

Dave Ward
Cisco

Email: dward@cisco.com

John Drake
Juniper

Email: jdrake@juniper.net

