CCAMP Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 15, 2015

E. Bellagamba
A. Takacs
G. Mirsky
Ericsson
L. Andersson
Huawei Technologies
P. Skoldstrom
Acreo AB
D. Ward
Cisco

January 11, 2015

Configuration of Pro-Active Operations, Administration, and Maintenance (OAM) Functions for MPLS-based Transport Networks using RSVP-TE

draft-ietf-ccamp-rsvp-te-mpls-tp-oam-ext-15

Abstract

This specification describes the configuration of pro-active MPLS-TP (MPLS-Transport Profile) Operations, Administration, and Maintenance (OAM) Functions for a given LSP using a set of TLVs that are carried by the GMPLS RSVP-TE protocol based on the OAM Configuration Framework for GMPLS RSVP-TE.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 15, 2015.

Copyright Notice

Table of Contents

1. Introduction

This document describes the configuration of pro-active MPLS-TP (MPLS-Transport Profile) Operations, Administration, and Maintenance (OAM) Functions for a given LSP using TLVs using GMPLS RSVP-TE [RFC3473]. The use of GMPLS RSVP-TE for the configuration of OAM functions is defined in a technology agnostic way in [RFC7260]. This document specifies the additional 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 pro-active MPLS-TP OAM functions running between 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.

MPLS-TP, the Transport Profile of MPLS, must, by definition [RFC5654], be capable of operating without a control plane. Therefore, there are several options for configuring MPLS-TP OAM, without a control plane by either using an NMS or LSP Ping, or with a control plane using signaling protocols such as RSVP-TE.

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

Pro-active MPLS-TP OAM is performed by three different protocols, Bi-directional Forwarding Detection (BFD) [RFC6428] for Continuity Check/Connectivity Verification, the delay measurement protocol (DM) [RFC6374] for delay and delay variation (jitter) measurements, and the loss measurement protocol (LM) [RFC6374] for packet loss and throughput measurements. Additionally there is a number of Fault Management Signals that can be configured.

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 track the liveliness and detect data plane failures of MPLS-TP point-to-point and might also be extended to support point-to-multipoint connections.

The delay and loss measurements protocols [RFC6374] use a simple query/response model for performing bidirectional measurements that allows the originating node to measure packet loss and delay in both directions. By timestamping and/or writing current packet counters

to the measurement packets at four times (Tx and Rx in both directions) current delays and packet losses can be calculated. By performing successive delay measurements the 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 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 here we are not configuring on-demand throughput estimates based on saturating the connection as defined in [RFC6371]. Rather, we only 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

CV - Connectivity Verification

CC - Continuity Check

DM - Delay Measurement

FMS - Fault Management Signal

G-ACh - Generic Associated Channel

GMPLS - Generalized Multi-Protocol Label Switching

LER - Label switching Edge Router

LM - Loss Measurement

LSP - Label Switched Path

LSR - Label Switching Router

MEP - Maintenance Entity Group End Point

MPLS - Multi-Protocol Label Switching

MPLS-TP - MPLS Transport Profile

NMS - Network Management System

PM - Performance Measurement

RSVP-TE - Reservation Protocol 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. Overview of MPLS OAM for Transport Applications

[RFC6371] describes how MPLS-TP OAM mechanisms are operated to meet transport requirements outlined in [RFC5860].

[RFC6428] specifies two BFD operation modes: 1) "CC mode", which uses periodic BFD message exchanges with symmetric timer settings, supporting Continuity Check, 2) "CV/CC mode" which sends unique maintenance entity identifiers in the periodic BFD messages supporting Connectivity Verification (CV) as well as Continuity Check (CC).

[RFC6374] specifies mechanisms for performance monitoring of LSPs, in particular it specifies loss and delay measurement OAM functions.

[RFC6427] specifies fault management signals with which a server LSP can notify client LSPs about various fault conditions to suppress alarms or to be used as triggers for actions in the client LSPs. The following signals are defined: Alarm Indication Signal (AIS), Link Down Indication (LDI) and Lock Report (LKR).

[RFC6371] describes the mapping of fault conditions to consequent actions. Some of these mappings may be configured by the operator, depending on the application of the LSP. The following defects are identified: Loss Of Continuity (LOC), Misconnectivity, MEP Misconfiguration and Period Misconfiguration. Out of these defect conditions, the following consequent actions may be configurable: 1) whether or not the LOC defect should result in blocking the outgoing data traffic; 2) whether or not the "Period Misconfiguration defect" should result in a signal fail condition.

3. Theory of Operations

3.1. MPLS-TP OAM Configuration Operation Overview

GMPLS RSVP-TE, or alternatively LSP Ping [LSP-PING-CONF], can be used to simply enable the different OAM functions, by setting the corresponding flags in the "OAM Function Flags sub-TLV" [RFC7260].

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 should not process or modify any of the OAM configuration TLVs but simply forward them to the end-node. There is one exception to this and that is if the "MPLS OAM FMS sub-TLV" is present. This sub-TLV has to be examined even by intermediate nodes, but only acted upon by nodes capable of transmitting FMS signals into the LSP being established. The sub-TLV MAY be present when the FMS flag is set in the "OAM Function Flags sub-TLV". If this sub-TLV is present, then the "OAM MIP entities desired" and "OAM MEP entities desired" flags (described in [RFC7260]) in the "LSP Attributes Flags TLV" MUST be set and the entire OAM Configuration TLV placed either in the LSP_REQUIRED_ATTRIBUTES object or in the LSP_ATTRIBUTES object in order to ensure that capable intermediate nodes process the configuration. If placed in LSP_ATTRIBUTES object, nodes that are not able to process the OAM Configuration TLV will forward the message without generating an error. If the ?MPLS OAM FMS sub-TLV? been placed in the LSP_REQUIRED_ATTRIBUTES object a node that supports the RFC 7260 but does not support the "MPLS OAM FMS sub-TLV" MUST generate PathErr message with "OAM Problem/Configuration Error" [RFC7260]. Otherwise, if the node doesn't support the RFC 7260, it will not raise any errors as described in the Section 4.1 [RFC7260].

Finally, if the "MPLS OAM FMS sub-TLV" is not included only the "OAM MEP entities desired" flag is set and the OAM Configuration TLV may be placed in either LSP_ATTRIBUTES or LSP_REQUIRED_ATTRIBUTES.

3.1.1. Configuration of BFD sessions

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

- Asynchronous mode, where both sides should be in active mode

- Unidirectional mode

In the simplest scenario, RSVP-TE, or alternatively LSP Ping [LSP-PING-CONF], 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 LSP Ping based bootstrapping described in [RFC5884]) or directly in the RSVP-TE signaling messages.

When BFD Control packets are transported in the G-ACh they are not protected by any end-to-end checksum, only lower-layers are providing 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" has an Integrity flag that when set enables BFD Authentication using Keyed SHA1 with an empty key (all 0s) [RFC5880]. This would ensure that every BFD Control packet carries an 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" MUST be included in the "BFD Configuration sub-TLV". The "BFD Authentication sub-TLV" is used to specify which authentication method should be used and which pre-shared key / password should be used for this particular session. How the key exchange is performed is out of scope of this document.

3.1.2. Configuration of Performance Monitoring

It is possible to configure Performance Monitoring functionalities such as Loss, Delay, 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 "OAM Function Flags sub-TLV", or a customized configuration. To customize the configuration, one would set the respective flags and including the respective Loss and/or Delay sub-TLVs).

By setting the PM/Loss flag in the "OAM Function Flags sub-TLV" and by including the "MPLS OAM PM Loss sub-TLV", one can configure the measurement interval and loss threshold values for triggering protection.

Delay measurements are configured by setting PM/Delay flag in the "OAM Function Flags sub-TLV", and by including the "MPLS OAM PM Loss sub-TLV", one can configure the measurement interval and the delay threshold values for triggering protection.

3.1.3. Configuration of Fault Management Signals

To configure Fault Management Signals and their refresh time, the FMS flag in the "OAM Function Flags sub-TLV" MUST be set and the "MPLS OAM FMS sub-TLV" included. When configuring Fault Management Signals, 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 intended to originate fault management signal 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? been set as indication of the request to create the association at each intermediate node of the client LSP. Corresponding server MEP needs to be configured by its own RSVP-TE session (or, alternatively, via an NMS or LSP-ping).

3.2. MPLS OAM Configuration sub-TLV

The "OAM Configuration TLV", defined in [RFC7260], specifies the OAM functions that are used for the LSP. This document extends the "OAM Configuration TLV" by defining a new OAM Type: "MPLS OAM" (TBA1). The "MPLS OAM" type is set to request the establishment of OAM functions for MPLS-TP LSPs. The specific OAM functions are specified in the "OAM Function Flags sub-TLV" as depicted in [RFC7260].

When an egress LSR receives an "OAM Configuration TLV" indicating the MPLS OAM type, the LSR will first process any present "OAM Function Flags sub-TLV" and then it MUST process technology specific configuration TLVs. This document defines a sub-TLV, the "MPLS OAM Configuration sub-TLV" which is carried in the "OAM Configuration TLV".

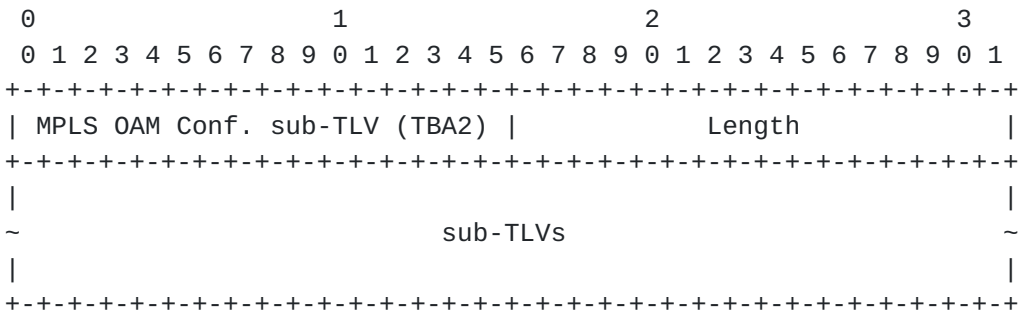

Type: TBA2, the "MPLS OAM Configuration sub-TLV".

Length: indicates the total length in octets, including sub-TLVs as well as the Type and Length fields.

The following MPLS OAM specific sub-TLVs MAY be included in the "MPLS OAM Configuration sub-TLV":

- "BFD Configuration sub-TLV", which MUST be included if the CC and/or the CV OAM Function flag is set. This sub-TLV carries additional sub-TLVs, failure to include the correct sub-TLVs MUST result in an error being generated: "OAM Problem/Configuration Error". The sub-TLVs are:

  - "BFD Identifiers sub-TLV", MUST always be included.

  - "Timer Negotiation Parameters sub-TLV", MUST be included if the N flag is not set.

  - "BFD Authentication sub-TLV", MAY be included if the I flag is set.

- "Performance Monitoring sub-TLV", which MUST be included if any of the PM/Delay, PM/Loss or PM/Throughput flags are set in the "OAM Function Flag sub-TLV". This sub-TLV MAY carry additional sub-TLVs:

  - "MPLS OAM PM Loss sub-TLV" MAY be included if the PM/Loss OAM Function flag is set. If the "MPLS OAM PM Loss sub-TLV" is not included, default configuration values are used. The same sub-TLV MAY also be included in case the PM/Throughput OAM Function flag is set and there is the need to specify measurement interval different from the default ones. Since throughput measurements use the same tool as loss measurements the same TLV is used.

  - "MPLS OAM PM Delay sub-TLV" MAY be included if the PM/Delay OAM Function flag is set. If the "MPLS OAM PM Delay sub-TLV" is not included, default configuration values are used.

- "MPLS OAM FMS sub-TLV" MAY be included if the FMS OAM Function flag is set. If the "MPLS OAM FMS sub-TLV" is not included, default configuration values are used.

Following are some additional rules of processing MPLS OAM Configuration sub-TLV:

- MPLS OAM Configuration sub-TLV MAY be empty, i.e. have no Value. Then its Length MUST be 8. Then all OAM functions that have their corresponding flags set in the "OAM Function Flags sub-TLV" MUST be assigned their default values or left disabled.

- sub-TLV that doesn't have corresponding flag set MUST be silently ignored;

- if multiple copies of a sub-TLV are present, then only the first sub-TLV MUST be used and the remaining sub-TLVs MUST be silently ignored.

However, not all the values can be derived from the standard RSVP-TE objects, in particular the locally assigned Tunnel ID at the egress cannot be derived by the ingress node. Therefore, the full LSP MEP-ID used by the ingress has to be carried in the "BFD Identifiers sub-TLV" in the Path message and the egress LSP MEP-ID in the same way in the Resv message.

3.2.1. CV Flag Rules of Use

If the CV flag is set, then the CC flag MUST be set as well because performing Connectivity Verification implies performing Continuity Check as well. The format of an MPLS-TP CV/CC message is shown in [RFC6428]. In order to perform Connectivity Verification the CV/CC message MUST contain the ?LSP MEP-ID? in addition to the BFD Control packet information. The "LSP MEP-ID" contains four identifiers:

MPLS-TP Global_ID

MPLS-TP Node Identifier

Tunnel_Num

LSP_Num

These values need to be correctly set by both ingress and egress when transmitting a CV packet and both ingress and egress needs to know what to expect when receiving a CV packet. Most of these values can be derived from the Path and Resv messages [RFC3473], which uses a 5-tuple to uniquely identify an LSP within an operator's network. This tuple is composed of a Tunnel Sender Address, Tunnel Endpoint Address, Tunnel_ID, Extended Tunnel ID, and (GMPLS) LSP_ID.

3.3. BFD Configuration sub-TLV

The "BFD Configuration sub-TLV" (depicted below) is defined for BFD OAM specific configuration parameters. The "BFD Configuration sub-TLV" is carried as a sub-TLV of the "MPLS OAM Configuration sub-TLV".

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

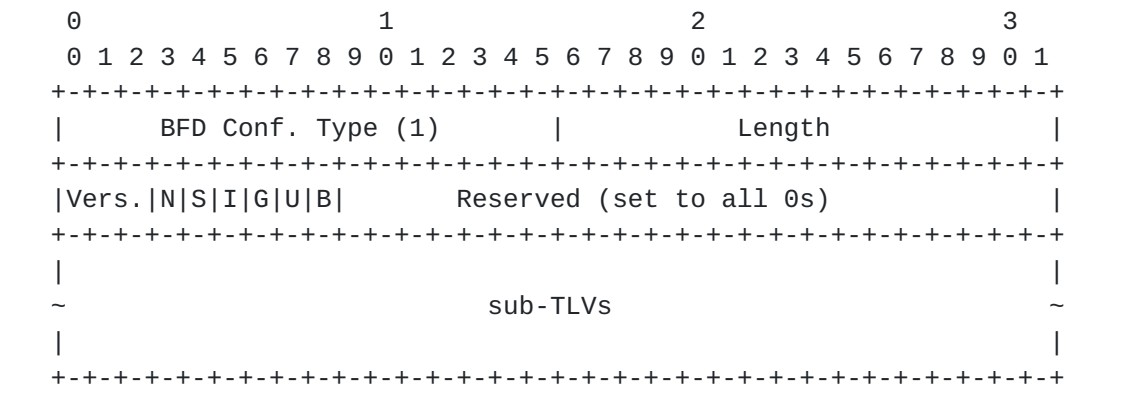

Type: 1, the "BFD Configuration sub-TLV".

Length: indicates the total length in octets, including sub-TLVs as well as the Type and Length fields.

Version: identifies the BFD protocol version. If the egress LSR does not support the version an error MUST be generated: "OAM Problem/ Unsupported BFD Version".

BFD Negotiation (N): If set timer negotiation/re-negotiation via BFD Control Messages is enabled, when cleared it is disabled.

Symmetric session (S): If set the BFD session MUST use symmetric timing values.

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".

Encapsulation Capability (G): if set, 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. 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".

Encapsulation Capability (U): if set, shows the capability of encapsulating BFD messages into 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 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 and set to 0 on transmission and ignored when received.

The "BFD Configuration sub-TLV" MUST include the following sub-TLVs in the Path message:

- "BFD Identifiers sub-TLV";

- "Negotiation Timer Parameters sub-TLV" if the N flag is cleared.

The "BFD Configuration sub-TLV" MUST include the following sub-TLVs in the Resv message:

- "BFD Identifiers sub-TLV;"

- "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 "Negotiation Timer Parameters sub-TLV" received by the egress contains unsupported values. In this case an updated "Negotiation Timer Parameters sub-TLV", containing values supported by the egress LSR, MUST be returned to the ingress.

3.3.1. BFD Identifiers sub-TLV

The "BFD Identifiers sub-TLV" is carried as a sub-TLV of the "BFD Configuration sub-TLV" and is depicted below.

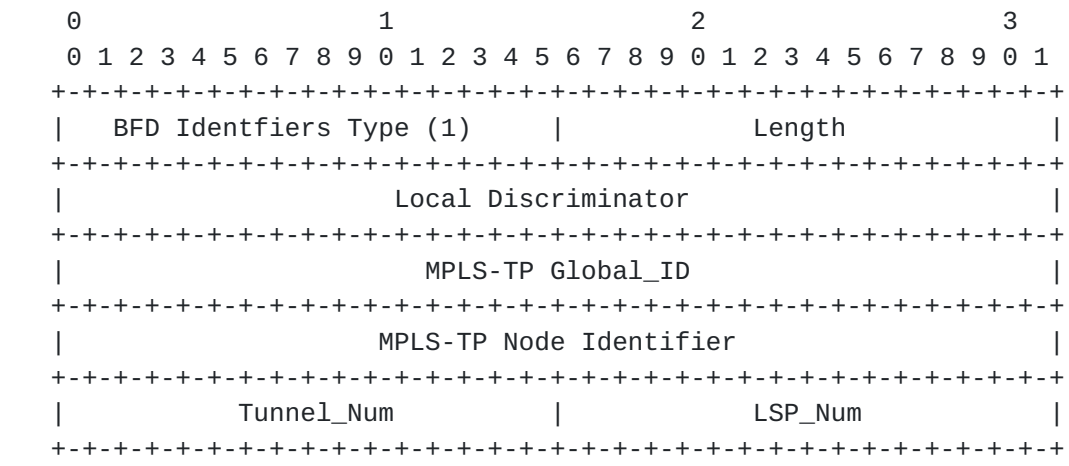

Type: 1, "BFD Identifiers sub-TLV".

Length: indicates the TLV total length in octets, including the Type and Length fields (20).

Local Discriminator: A unique, nonzero discriminator value generated by the transmitting system and referring to itself, used to demultiplex multiple BFD sessions between the same pair of systems as defined in [RFC5880].

MPLS-TP Global_ID, Node Identifier, Tunnel_Num, and LSP_Num: all set as defined in [RFC6370].

3.3.2. Negotiation Timer Parameters sub-TLV

The "Negotiation Timer Parameters sub-TLV" is carried as a sub-TLV of the "BFD Configuration sub-TLV" and is depicted below.

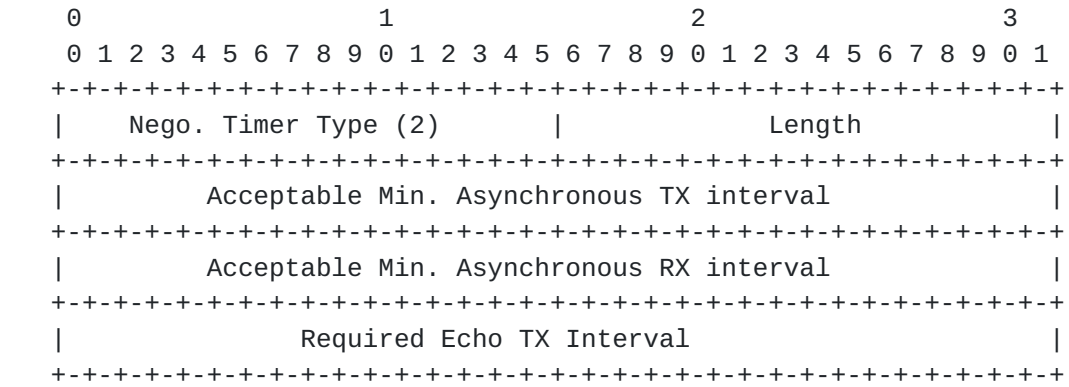

Type: 2, "Negotiation Timer Parameters sub-TLV".

Length: indicates the TLV total length in octets, including Type and Length fields (16).

Acceptable Min. Asynchronous TX interval: If the S (symmetric) flag is set in the "BFD Configuration sub-TLV", 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 egress LSR cannot support the value, it SHOULD reply with a supported interval.

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

Acceptable Min. Asynchronous RX interval: If the S (symmetric) flag is set in the "BFD Configuration sub-TLV", this field MUST be set equal to "Acceptable Min. Asynchronous TX interval" on transmit and MUST be ignored on receipt since it has no additional meaning with 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 the ingress/egress LSRs can receive periodic BFD control packets. If this value is greater than the "Acceptable Min. Asynchronous TX interval" received from the ingress/egress LSR, the receiving LSR MUST adopt the interval expressed in the "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 [RFC5880] sect. 6.8.9. 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 LSR node cannot support this value it SHOULD reply with a supported value (which may be zero if Echo is not supported).

3.3.3. BFD Authentication sub-TLV

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

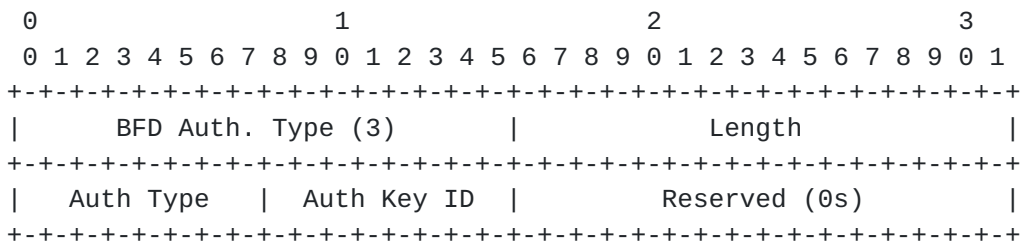

Type: 3, "BFD Authentication sub-TLV".

Length: indicates the TLV total length in octets, including Type and Length fields (8).

Auth Type: indicates which type of authentication to use. The same values as are defined in section 4.1 of [RFC5880] are used. If the egress LSR does not support this type an "OAM Problem/Unsupported BFD Authentication Type" error MUST be generated.

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 and set to 0 on transmission and ignored when received.

3.3.4. Traffic Class sub-TLV

The Traffic Class sub-TLV is carried as a sub-TLV of the "BFD Configuration sub-TLV" or "Fault Management Signal sub-TLV" Section 3.5 and is depicted in Figure 1.

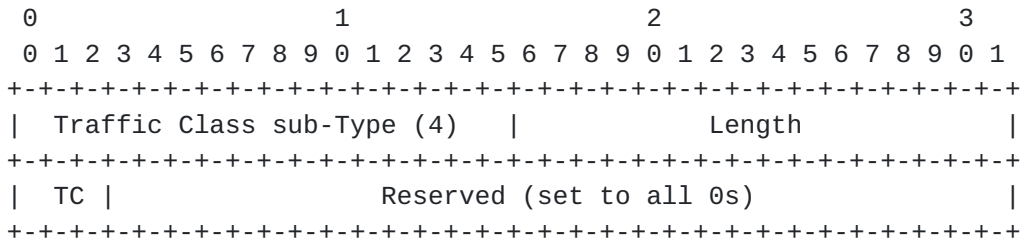

Figure 1: Traffic Class sub-TLV format

Type: 3, "Traffic Class sub-TLV".

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 TC sub-TLV is present, then the value of the TC field MUST be used as the value of the TC field of an MPLS label stack entry. If the TC sub-TLV is absent from "BFD Configuration sub-TLV" or "Fault Management Signal sub-TLV", then selection of the TC value is local decision.

3.4. Performance Monitoring sub-TLV

If the "OAM Function Flags sub-TLV" has either the PM/Loss, PM/Delay or PM/Throughput flag set, the "Performance Monitoring sub-TLV" MUST be present in the "MPLS OAM Configuration sub-TLV". Failure to include the correct sub-TLVs MUST result in an "OAM Problem/ Configuration Error" 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 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" includes the following sub-TLVs:

- "MPLS OAM PM Loss sub-TLV" if the PM/Loss and/or PM/Throughput flag is set in the "OAM Function Flags sub-TLV";

- "MPLS OAM PM Delay sub-TLV" if the PM/Delay flag is set in the "OAM Function Flags sub-TLV";

The "Performance Monitoring sub-TLV" depicted below is carried as a sub-TLV of the "MPLS OAM Configuration sub-TLV".

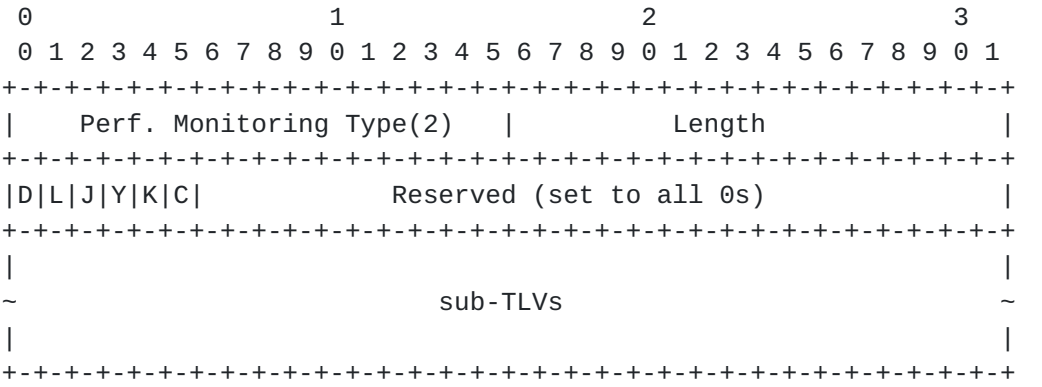

Type: 2, "Performance Monitoring sub-TLV".

Length: indicates the TLV total length in octets, including sub-TLVs as well as Type and Length fields.

Configuration Flags, for the specific function description please refer to [RFC6374]:

- D: Delay inferred/direct (0=INFERRED, 1=DIRECT). If the egress LSR does not support 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 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 and set to 0 on transmission and ignored when received.

3.4.1. MPLS OAM PM Loss sub-TLV

The "MPLS OAM PM Loss sub-TLV" depicted below is carried as a sub-TLV of the "Performance Monitoring sub-TLV".

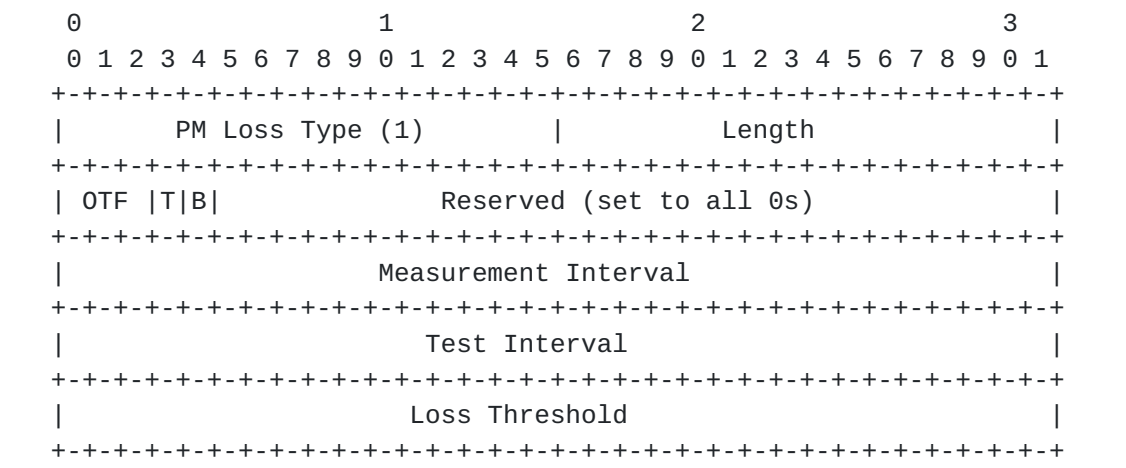

Type: 1,"MPLS OAM PM Loss sub-TLV".

Length: indicates the length of the parameters in octets, including Type and Length fields (20).

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 (DSCP 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 and 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 egress LSR cannot support the value, it SHOULD reply with a supported 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]. If the

egress LSR cannot support the value, it SHOULD reply with a supported interval.

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

3.4.2. MPLS OAM PM Delay sub-TLV

The "MPLS OAM PM Delay sub-TLV" depicted below is carried as a sub-TLV of the "Performance Monitoring sub-TLV".

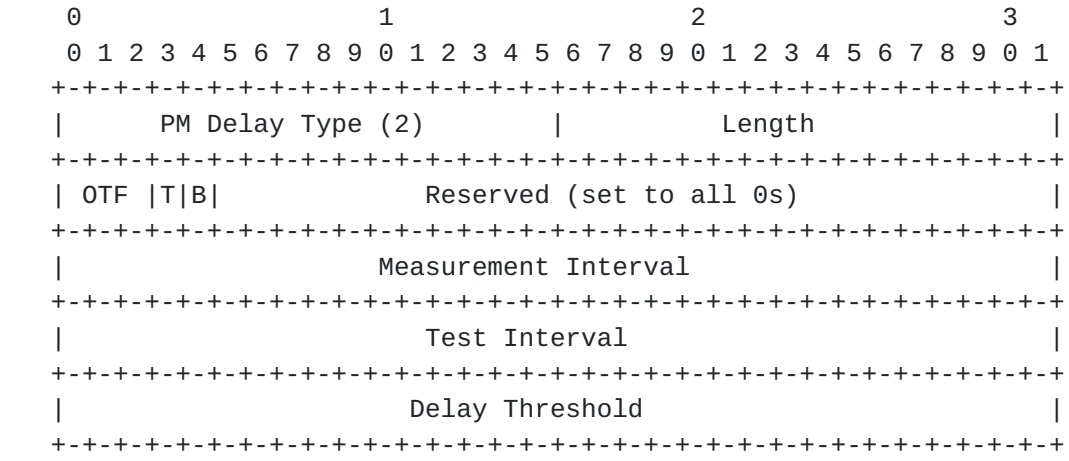

Type: 2,"MPLS OAM PM Delay sub-TLV".

Length: indicates the length of the parameters in octets, including Type and Length fields (20).

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 (DSCP 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 and 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 egress LSR cannot support the value, it SHOULD reply with a supported 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]. If the egress LSR cannot support the value, it SHOULD reply with a supported interval.

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

3.5. MPLS OAM FMS sub-TLV

The "MPLS OAM FMS sub-TLV" depicted below is carried as a sub-TLV of the "MPLS OAM Configuration sub-TLV". When both working and protection paths are signaled, both LSPs SHOULD be signaled with identical settings of the E flag, T flag, and the refresh timer.

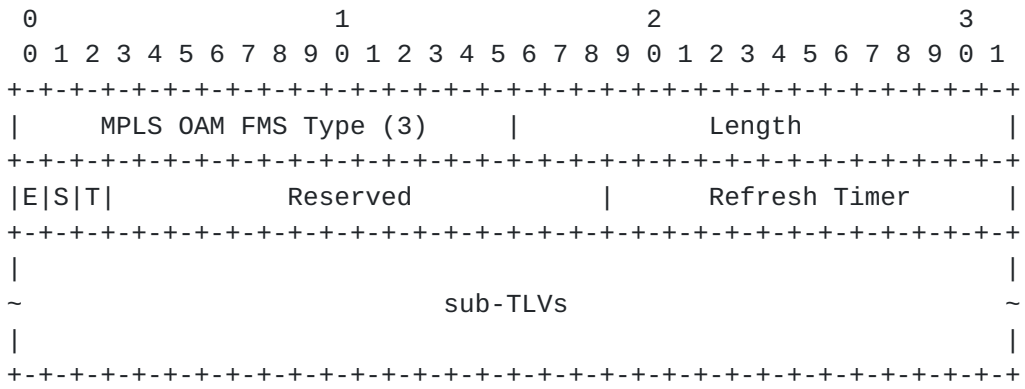

Type: 3, "MPLS OAM FMS sub-TLV".

Length: indicates the TLV total length in octets, including Type and Length fields (8).

FMS Signal Flags are used to enable the FMS signals 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 signal 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 client LSPs. Default value is 0 (disabled). If a Server MEP which 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).

- Remaining bits: Reserved for future specification and set to 0.

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 egress LSR cannot support the value it SHOULD reply with a supported timer value.

FMS sub-TLV MAY include Traffic Class sub-TLV Section 3.3.4. If 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 local decision.

4. Summary of MPLS OAM configuration errors

In addition to error values specified in [RFC7260] this document defines the following values for the "OAM Problem" Error Code:

- If an egress LSR does not support the specified BFD version, an error MUST be generated: "OAM Problem/Unsupported BFD Version".

- If an egress LSR does not support the specified BFD Encapsulation format, an error MUST be generated: "OAM Problem/ Unsupported BFD Encapsulation format".

- If an egress LSR does not support BFD Authentication, and it is requested, an error MUST be generated: "OAM Problem/BFD Authentication unsupported".

- If an egress LSR does not support the specified BFD Authentication Type, an error MUST be generated: "OAM Problem/ Unsupported BFD Authentication Type".

- 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".

- If an egress LSR does not support the specified Timestamp Format, an error MUST be generated: "OAM Problem/Unsupported Timestamp Format".

- If an egress LSR does not support specified Delay mode, an "OAM Problem/Unsupported Delay Mode" error MUST be generated.

- If an egress LSR does not support specified Loss mode, an "OAM Problem/Unsupported Loss Mode" error MUST be generated.

- If an egress LSR does not support Delay variation measurements, and it is requested, an "OAM Problem/Delay variation unsupported" error MUST be generated.

- If an egress LSR does not support Dyadic mode, and it is requested, an "OAM Problem/Dyadic mode unsupported" error MUST be generated.

- If an egress LSR does not support Loopback mode, and it is requested, an "OAM Problem/Loopback mode unsupported" error MUST be generated.

- If an egress LSR does not support Combined mode, and it is requested, an "OAM Problem/Combined mode unsupported" error MUST be generated.

- 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.

- 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.

5. IANA Considerations

5.1. MPLS OAM Type

This document specifies new the "MPLS OAM Type". IANA is requested to allocate a new type (TBA1) from the OAM Type space of the RSVP-TE OAM Configuration Registry.

Type	Description	Reference
TBA1	MPLS OAM	This document

Table 1: OAM MPLS Type

5.2. MPLS OAM Configuration sub-TLV

This document specifies the "MPLS OAM Configuration sub-TLV", IANA is requested to allocate a new type (TBA2) from the technology-specific sub-TLV space of the RSVP-TE OAM Configuration Registry.

Type	Description	Reference
TBA2	MPLS OAM Configuration sub-TLV	This document

Table 2: MPLS OAM Configuration sub-TLV Type

5.3. MPLS OAM Configuration Sub-TLV Types

IANA is requested to create an "MPLS OAM sub-TLV Types" sub-registry in the "RSVP-TE OAM Configuration Registry" for the sub-TLVs carried in the "MPLS OAM Configuration sub-TLV". Values from this new sub-registry to be allocated through IETF Review except for the Reserved for Experimental Use range. This document defines the following types:

Type	Description	Reference
0	Reserved	This document
1	BFD Configuration sub-TLV	This document
2	Performance Monitoring sub-TLV	This document
3	MPLS OAM FMS sub-TLV	This document
4-65532	Unassigned	
65533-65534	Reserved for Experimental Use	This document
65535	Reserved	This document

Table 3: MPLS OAM Configuration sub-TLV Types

5.4. BFD Configuration Sub-TLV Types

IANA is requested to create a "BFD Configuration sub-TLV Types" sub-registry in the "RSVP-TE OAM Configuration Registry" for the sub-TLV types carried in the "BFD Configuration sub-TLV". Values from this new sub-registry to be allocated through IETF Review except for the Reserved for Experimental Use range. This document defines the following types:

Type	Description	Reference
0	Reserved	This document
1	BFD Identifiers sub-TLV	This document
2	Negotiation Timer Parameters sub-TLV	This document
3	BFD Authentication sub-TLV	This document
4	Traffic Class sub-TLV	This document
5-65532	Unassigned	
65533-65534	Reserved for Experimental Use	This document
65535	Reserved	This document

Table 4: BFD Configuration Sus-TLV Types

5.5. Performance Monitoring sub-TLV Types

IANA is requested to create a "Performance Monitoring sub-TLV Type" sub-registry in the "RSVP-TE OAM Configuration Registry" for the sub-TLV types carried in the "Performance Monitoring sub-TLV". Values from this new sub-registry to be allocated through IETF Review except for the Reserved for Experimental Use range. This document defines the following types:

Type	Description	Reference
0	Reserved	This document
1	MPLS OAM PM Loss sub-TLV	This document
2	MPLS OAM PM Delay sub-TLV	This document
3-65532	Unassigned	
65533-65534	Reserved for Experimental Use	This document
65535	Reserved	This document

Table 5: Performance Monitoring sub-TLV Types

5.6. New RSVP-TE error codes

The following values need to be assigned under the "OAM Problem" Error Code [RFC7260] by IETF Review process:

Error Value Sub-codes	Description	Reference
TBA3	Unsupported BFD Version	This document
TBA4	Unsupported BFD Encapsulation format	This document
TBA5	Unsupported BFD Authentication Type	This document
TBA6	Mismatch of BFD Authentication Key ID	This document
TBA7	Unsupported Timestamp Format	This document
TBA8	Unsupported Delay Mode	This document
TBA9	Unsupported Loss Mode	This document
TBA10	Delay variation unsupported	This document
TBA11	Dyadic mode unsupported	This document
TBA12	Loopback mode unsupported	This document
TBA13	Combined mode unsupported	This document
TBA14	Fault management signaling unsupported	This document
TBA15	Unable to create fault management association	This document

Table 6: MPLS OAM Configuration Error Codes

The "Sub-codes - 40, OAM Problem" sub-registry is located in the "Error Codes and Globally-Defined Error Value Sub-Codes" registry.

6. Contributing Authors

This document is the result of a large team of authors and contributors. The following is a list of the co-authors:

John Drake

Benoit Tremblay

7. Acknowledgements

The authors would like to thank David Allan, Lou Berger, Annamaria Fulignoli, Eric Gray, Andras Kern, David Jocha and David Sinicrope for their useful comments.

8. Security Considerations

The signaling of OAM related parameters and the automatic establishment of OAM entities introduces additional security considerations to those discussed in [RFC3473]. 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 as discussed in [RFC7260] and [RFC6060].

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

[RFC5654] Niven-Jenkins, B., Brungard, D., Betts, M., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, September 2009.

[RFC5860] Vigoureux, M., Ward, D., and M. Betts, "Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks", RFC 5860, May 2010.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010.

[RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, June 2010.

[RFC6060] Fedyk, D., Shah, H., Bitar, N., and A. Takacs, "Generalized Multiprotocol Label Switching (GMPLS) Control of Ethernet Provider Backbone Traffic Engineering (PBB-TE)", RFC 6060, March 2011.

[RFC6370] Bocci, M., Swallow, G., and E. Gray, "MPLS Transport Profile (MPLS-TP) Identifiers", RFC 6370, September 2011.

[RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, September 2011.

[RFC6427] Swallow, G., Fulignoli, A., Vigoureux, M., Boutros, S., and D. Ward, "MPLS Fault Management Operations, Administration, and Maintenance (OAM)", RFC 6427, November 2011.

[RFC6428] Allan, D., Swallow Ed. , G., and J. Drake Ed. , "Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile", RFC 6428, November 2011.

[RFC7260] Takacs, A., Fedyk, D., and J. He, "GMPLS RSVP-TE Extensions for Operations, Administration, and Maintenance (OAM) Configuration", RFC 7260, June 2014.

9.2. Informative References

[LSP-PING-CONF] Bellagamba, E., Andersson, L., Ward, D., Drake, J., and P. Skoldstrom, "Configuration of pro-active MPLS-TP Operations, Administration, and Maintenance (OAM) Functions Using LSP Ping", 2014, <draft-ietf-mpls-lsp-ping-mpls-tp-oam-conf>.

[RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, February 2009.

[RFC6371] Busi, I. and D. Allan, "Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks", RFC 6371, September 2011.

[RFC6375] Frost, D. and S. Bryant, "A Packet Loss and Delay Measurement Profile for MPLS-Based Transport Networks", RFC 6375, September 2011.

Authors' Addresses

Elisa Bellagamba
Ericsson

Email: elisa.bellagamba@ericsson.com

Attila Takacs
Ericsson

Email: attila.takacs@ericsson.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 70 7957731
Email: pontus.skoldstrom@acreo.se

Dave Ward
Cisco

Email: dward@cisco.com