PWE3 Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 17, 2012

F. Zhang, Ed.
B. Wu, Ed.
ZTE Corporation
E. Bellagamba, Ed.
Ericsson

August 16, 2011

Label Distribution Protocol Extensions for Proactive Operations, Administration and Maintenance Configuration of Dynamic MPLS Transport Profile PseudoWire

draft-zhang-mpls-tp-pw-oam-config-06

Abstract

This document specifies extensions to the Label Distribution Protocol (LDP) to configure and control proactive Operations, Adminstration and Maintenance (OAM) functions, suitable for dynamic Single-Segment PseudoWire (SS-PW) and Multi-Segment PseudoWire (MS-PW).

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 February 17, 2012.

Table of Contents

1. Introduction

MPLS Pseudowire (PW) is defined in [RFC3985] and [RFC5659], which provide for emulated services over an MPLS Packet Switched Network (PSN). 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 previously supported by IP/MPLS. The corresponding requirements are defined in [RFC5860].

The MPLS-TP OAM mechanisms that are operated to meet transport requirements are described in [I-D.ietf-mpls-tp-oam-framework], categorized into proactive and on-demand monitoring. Proactive monitoring refers to OAM operations that are either configured to be carried out periodically and continuously or preconfigured to act on certain events such as alarm signals. In contrast, on-demand monitoring is initiated manually and for a limited amount of time, usually for operations such as diagnostics to investigate into a defect condition.

The Network Management System (NMS) or Label Switched Path (LSP) Ping [I-D.ietf-mpls-lsp-ping-mpls-tp-oam-conf] is used to configure these OAM functionalities if a control plane is not instantiated. But if the control plane is used, it MUST support the configuration and modification of OAM maintenance points as well as the activation/ deactivation of OAM when the transport path or transport service is established or modified [RFC5654].

This document specifies the extensions to the LDP protocol to negotiate PW OAM capabilities, configure and bootstrap proactive PW OAM functions, suitable for Point to Point (P2P) SS-PW and MS-PW. The extensions to Point to Multi-Point (P2MP) PW will be studied in the future.

2. Conventions used in this document

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

2.1. Acronyms

AC: Attachment Circuit
AIS: Alarm indication signal
BFD: Bidirectional Forwarding Detection
CC: Continuity Check
CV: Connectivity Verification
DM: Delay Measurement
FEC: Forwarding Equivalence Class
FMS: Fault Management Signal
ICMP: Internet Control Message Protocol
G-ACh: Generic Associated Channel
LDI: Link Down Indication
LDP: Label Distribution Protocol
LKR: Lock Reporting
LM: Loss Measurement
LSP: Label Switched Path
ME: Maintenance Entity
MEG: Maintenance Entity Group
MEP: Maintenance Entity Group End Point
MIP: Maintenance Entity Group Intermediate Point
MPLS-TP: MPLS Transport Profile
MS-PW: Multi-Segment PseudoWire
NMS: Network Management System
OAM: Operations, Adminstration and Maintenance
P2MP: Point to Multi-Point
PE: Provider Edge
PHB: Per-Hop Behavior
PM: Performance Monitoring
PSN: Packet Switched Network
PW: PseudoWire
S-PE: Switching Provider Edge
SPME: Sub-Path Maintenance Entity
SS-PW: Single-Segment Pseudo Wire
T-PE: Terminating Provider Edge
TLV: Type Length Value
VCCV: Virtual Circuit Connectivity Verification

3. Analysis of Existing PW OAM Configuration

Before MPLS-TP standards, PW OAM functions have been implemented by [RFC5085], [RFC5885], [RFC4447] and [I-D.ietf-pwe3-static-pw-status]. [RFC5085] defines Connectivity Verification (CV) function, which belongs to on-demand PW monitoring. Continuity Check (CC), as well as PW and Attachemnt Circuit (AC) status notification, are defined in [RFC5885]. The documents [RFC4447] and [I-D.ietf-pwe3-static-pw-status] give some other ways of PW/AC status notification.

3.1. Virtual Circuit Connectivity Verification

Virtual Circuit Connectivity Verification (VCCV) is used to verify and further diagnose PW forwarding path, and the VCCV capabilities negotiation is defined in [RFC5085].

3.2. VCCV Bidirectional Forwarding Detection

Four CV types based on Bidirectional Forwarding Detection (BFD) are specified in [RFC5885], which describes the VCCV BFD capabilities negotiation and the procedures of selecting one of them when multiple BFD CV types are advertised.

3.3. PW Status

PW status codes provide a mechanism to signal the status of PW and AC failure. When PW control plane exists, the PW Status TLV is carried in the initial Label Mapping message and Notification message to signal all PW status messages [RFC4447]. When an event occurs, an update PW status will be sent.

3.4. Conclusion

In summary, IP/MPLS PW OAM functions and their relationship with LDP/ LSP Ping/NMS are described in table 1. This document will not replace or deprecate these existing functions(e.g., VCCV capability advertisement and PW status negotiation for MPLS networks).

		LDP	LSP Ping	NMS
On-demand OAM	VCCV LSP ping	Capability negotiation		Capability configuration& Bootstrapping
	VCCV ICMP ping	Capability negotiation		Capability configuration& Bootstrapping
Proactive OAM	VCCV BFD	Capability negotiation& Bootstrapping		Capability configuration& Bootstrapping
	PW status	Capability negotiation& Bootstrapping		Capability configuration& Bootstrapping

Table 1: IP/MPLS PW OAM Functions

4. Analysis of PW OAM Configuration Extended by MPLS-TP

4.1. Continuity Check, Connectivity Verification and Remote Defect Indication

The Proactive CC, CV and Remote Defect Indication (RDI) functions of MPLS-TP are based on the extensions to BFD [I-D.ietf-mpls-tp-cc-cv-rdi], which addresses the proactive CV gap that VCCV BFD does not support. The BFD packets can be encapsulated by IP or non-IP in G-ACh, and operate in coordinated or independent mode.

Timer negotiation, such as Transmitter (TX)/Receiver (RX) interval can be performed in subsequent BFD control messages [RFC5880] or directly in the LSP ping configuration messages, but it also can be gotten by control plane signaling [I-D.ietf-mpls-tp-oam-framework].

The use of the VCCV control channel provides the context, based on the MPLS-PW label, required to bind and bootstrap the BFD session to a particular PW, so local discriminator values are not exchanged [I-D.ietf-mpls-tp-oam-analysis]. However, in order to identify certain extreme cases of mis-connectivity and fulfill the requirements that the BFD mechanism MUST be the same for LSP, Single Segment Pseudowire (SS-PW), Multi Segment Pseudowire (MS-PW) and Section as well as for Sub-path Maintenance Element (SPME), BFD might still need to use discriminator values to identify the connection being verified at both ends of the PW. The discriminator values can be statically configured, or signaled via LSP Ping or LDP extensions defined in this document.

Per-hop Behavior (PHB), which identifies the per-hop behavior of BFD packet, SHOULD be configured as well. This permits the verification of correct operation of Quality of Serivce (QoS) queuing as well as connectivity.

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 make every BFD Control packet carry an SHA1 hash of itself that can be used to detect errors.

If BFD Authentication using a shared key / password is desired (i.e. actual authentication 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 that should be used and which shared key / password that should be used for this particular session. How the key exchange is performed is out of scope of this document.

4.2. Performance Monitoring Loss/Delay

Performance monitoring (PM) of PWs, especially for packet Loss Measurement (LM) and packet Delay Measurement (DM), are specified in [I-D.ietf-mpls-loss-delay], [I-D.ietf-mpls-tp-loss-delay-profile].

When configuring Performance monitoring functionalities it can be chosen either the default configuration (by only setting the respective flags in the "MPLS-TP PW OAM Configuration TLV") or a customized configuration (by including the respective MPLS-TP PW PM Loss and/or Delay sub-TLVs).

By setting the PM Loss flag in the "MPLS-TP PW OAM Configuration TLV" and including the "MPLS-TP PW 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 "MPLS-TP PW OAM Configuration TLV" and including the "MPLS-TP PW PM Loss sub-TLV" one can configure the measurement interval and the delay threshold values for triggering protection.

4.3. FMS

Fault Management Signals (FMS) are specified in [I-D.ietf-mpls-tp-fault], with which a server PW can notify client PWs about various fault conditions to suppress alarms or to be used as triggers for actions in the client PWs. The following signals are defined: Alarm Indication Signal (AIS), Link Down Indication (LDI) and Lock Reporting (LKR).

For each MEP of each Maintenance Entity Group (MEG), enabling/ disabling the generation of FMS packets, the transmitted period and PHB SHOULD be configured. This can be done independently, and the values of configured parameters can be different, but for easy maintenance, these setting SHOULD be consistent.

4.4. On-demand OAM Functions

The extended on-demand OAM functions MAY need capability negotiation in the LDP Initialization message [RFC5561]. However, On-demand PW OAM functions are expected to be carried out by directly accessing network nodes via a management interface; hence configuration and control of on-demand PW OAM functions are out-of-scope for this document.

4.5. Conclusion

According to the analysis above, LDP needs to be extended to negotiate PW OAM capabilities, configure and bootstrap proactive PW OAM functions, such as, CC-CV-RDI, PM Loss/Delay, FMS. In this way, OAM configuration is bound to PW signaling, avoiding two separate management/configuration steps (PW establishment followed by OAM configuration) which would increases delay, processing and more importantly may be prune to mis-configuration errors.

Furthermore, LSP ping can be used to configure the proactive PW OAM function extended by MPLS-TP also, suitable for dynamic and static PW. For reference, the following table 2 describes the different scope of different proactive OAM bootstrapping schemes of dynamic PW.

		LDP	LSP Ping	NMS
Proactive OAM	CC/CV/RDI	Capability negotiation& Function configuration& Bootstrapping	Function configuration& Bootstrapping	Capability configuration& Function configuration& Bootstrapping
	FMS	Capability negotiation& Function configuration& Bootstrapping	Function configuration& Bootstrapping	Capability configuration& Function configuration& Bootstrapping
	PM Loss/ Delay	Capability negotiation& Function configuration& Bootstrapping	Function configuration& Bootstrapping	Capability configuration& Function configuration& Bootstrapping

Table 2: MPLS-TP PW OAM Functions

5. MPLS-TP PW OAM Capability Advertisement

When a PW is first set up, the PEs MUST attempt to negotiate the usage of OAM functions. At the time of writing this specification, there are PW status negotiation and VCCV capability advertisement. For the proactive OAM functions extended by MPLS-TP, such as CC-CV-RDI, PM loss/delay and FMS, the capability negotiation MAY be also needed, so a PE that supports the MPLS-TP PW OAM capability MUST include MPLS-TP PW OAM Capability TLV in the LDP Initialization message. And if the peer has not advertised this capability, the corresponding PW OAM configuration information will not be sent to the peer.

6. PW OAM Configuration Procedures

A PE may play an active or passive role in the signaling of the PW. There exist two situations:

a) Active/active: both PEs of a PW are active (SS-PW), they select PW OAM configuration parameters and send with the Label Mapping message to each other independently.

b) Active/passive: one PE is active and the others are passive (MS-PW). The active/passive role election is defined in Section 7.2.1 of [RFC6073] and applies here, this document does not define any new role election procedures.

The general rules of OAM configuration procedures are mostly identical between MS-PW and SS-PW, except that SS-PW does not need to configure MIP function and the Mapping message are sent out independently. Section 6.1 takes MS-PW as an example to describe the general OAM configuration procedures. As for SS-PW, the specific differences would be addressed in section 6.2.

6.1. OAM Configuration for MS-PW

6.1.1. Establishment of OAM Entities and Functions

Assuming there is one PW that needs to be setup between T-PE1 and T-PE2, across S-PE1 and S-PE2. OAM functions must be setup and enabled in the appropriate order so that spurious alarms can be avoided.

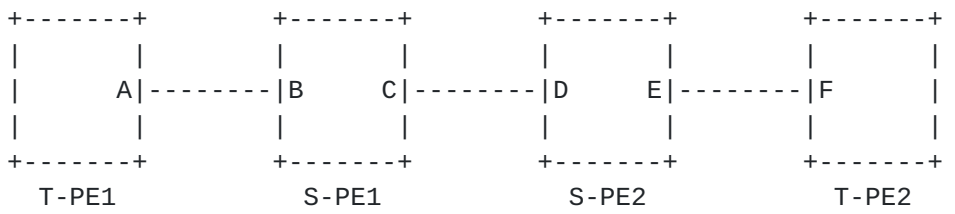

Figure 1: MS-PW OAM Configuration Scheme

Fist of all, T-PE1 MUST setup the OAM sink function to be prepared to receive OAM messages but MUST suppress any OAM alarms (e.g., due to missing or unidentified OAM messages). The Mapping message MUST be sent with the "OAM Alarms Enabled" cleared and "OAM MIP Entities desired" set in the MPLS-TP PW OAM Administation TLV.

When the Mapping message arrives at the downstream S-PEs, such as S-PE1 and S-PE2, they MUST establish and configure MIP entities according to the set "I"flag in the MPLS-TP PW OAM Administration TLV. If failure, a Notification message SHOULD be sent, with a Status Code set to "MIP Configuration Failure". If OAM entities are established successfully, the middle points (S-PE1 and S-PE2) MUST forward the Mapping message downstream, the endpoint (T-PE2) MUST set the OAM Source function and MUST be prepared to Send OAM messages.

The same rules are applied to the reverse direction (from T-PE2 to T-PE1), that is to say, T-PE2 needs to setup the OAM sink function to be prepared to receive OAM messages but MUST suppress any OAM alarms (e.g., due to missing or unidentified OAM messages). The Mapping message MUST be sent with the "OAM Alarms Enabled" cleared, "OAM MIP Entities desired" set in the MPLS-TP PW OAM Administration TLV. When T-PE1 receives the Mapping message, it completes any pending OAM configuration and enables the OAM source function to send OAM messages.

After this round, OAM entities are established and configured for the PW and OAM messages MAY already be exchanged, and OAM alarms can now be enabled. The T-PE nodes (T-PE1 and T-PE2), while still keeping OAM alarms disabled send a Notification message with "OAM Alarms Enabled" PW status flag set, and enable the OAM alarms after processing the Notification message. At this point, data-plane OAM is fully functional, and the MPLS-TP OAM PW configuration TLV MAY be omitted in subsequent Notification messages

The PW MAY be setup with OAM entities right away with the first signaling, as described above, but a PW MAY be signaled and established without OAM configuration first, and OAM entities may be added later. This can be done by sending a Notification message with

the related configuration parameters subsequently.

6.1.2. Adjustment of OAM Parameters

There may be a need to change the parameters of an already established and configured OAM function during the lifetime of the PW. To do so the T-PE nodes need to send a Notification message with the updated parameters. OAM parameters that influence the content and timing of OAM messages and identify the way OAM defects and alarms are derived and generated. Hence, to avoid spurious alarms, it is important that both sides, OAM sink and source, are updated in a synchronized way. Firstly, the alarms of the OAM sink function should be suppressed and only then should expected OAM parameters be adjusted. Subsequently, the parameters of the OAM source function can be updated. Finally, the alarms of the OAM sink side can be enabled again.

In accordance with the above operation, T-PE1 MUST send a Notification message with "OAM Alarms Enabled" cleared and including the updated MPLS-TP PW OAM Configuration TLV corresponding to the new parameter settings. The initiator (T-PE1) MUST keep its OAM sink and source functions running unmodified, but it MUST suppress OAM alarms after the updated Notification message is sent. The receiver (T-PE2) MUST firstly disable all OAM alarms, then update the OAM parameters according to the information in the Notification message and reply with a Notification message acknowledging the changes by including the MPLS-TP PW OAM Configuration TLV. Note that the receiving side has the possibility to adjust the requested OAM configuration parameters and reply with and updated MPLS-TP PW OAM Configuration TLV in the Notification message, reflecting the actually configured values. However, in order to avoid an extensive negotiation phase, in the case of adjusting already configured OAM functions, the receiving side SHOULD NOT update the parameters requested in the Notification message to an extent that would provide lower performance than what has been configured previously.

The initiator (T-PE1) MUST only update its OAM sink and source functions when it has received the Notification message from the peer. After the OAM parameters are updated and OAM is running according the new parameter settings, OAM alarms are still disabled, so a subsequent Notification messages exchanges with "OAM Alarms Enabled" flag set are needed to enable OAM alarms again.

6.1.3. Deleting OAM Entities

In some cases it may be useful to remove some or all OAM entities and functions from one PW without actually tearing down the connection. To avoid any spurious alarm, the following procedure should be

followed:

The T-PE nodes disable OAM alarms and SHOULD send Notification message to each other with "OAM Alarms Enabled" cleared but unchanged OAM configuration and without the MPLS-TP PW OAM Configuration TLV. After that, T-PE1 (T-PE2) SHOULD delete OAM source functions, then send a Notification message with "OAM MIP Entities desired" cleared. While T-PE2 (T-PE1) deletes OAM sink function, S-PE1 and S-PE2 delete MIP configuration when they receive the Notification message with "OAM MIP Entities desired" cleared.

Alternatively, if only some OAM functions need to be removed, the T-PE node sends the Notification message with the updated OAM Configuration TLV. Changes between the contents of the previously signaled OAM Configuration TLV and the currently received TLV represent which functions SHOULD be removed/added.

6.2. OAM Configuration for SS-PW

Assuming there is one PW that needs to be setup between T-PE1 and T-PE2.

If the receiving PE (T-PE2) have initiated the MPLS-TP PW OAM configuration request to the other PE (T-PE1), it MUST compare its AII against T-PE1's. If it is numerically lower, will reply a Notification message with the updated "MPLS-TP PW OAM Configuration TLV", and the Status Code set to "Wrong MPLS-TP PW OAM Configuration TLV".

On the other hand, if the T-PE2's AII is numerically higher than T-PE1's, it MUST reply a Notification message with Status Code set to "Rejected MPLS-TP PW OAM Configuration TLV".

7. LDP extensions

Below, LDP extensions to configure proactive MPLS-TP PW OAM functions are defined.

7.1. MPLS-TP PW OAM Capability TLV

A new Capability Parameter TLV called the MPLS-TP PW OAM Capability TLV is defined, and the format is as follows:

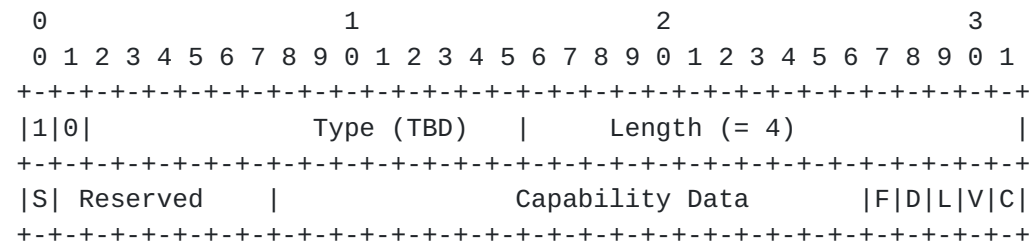

MPLS-TP PW OAM Capability TLV

The value of the U-bit for the MPLS-TP PW OAM Capability TLV MUST be set to 1 so that a receiver MUST silently ignore this TLV if unknown to it, and continue processing the rest of the message[RFC5036]. Currently defined specific OAM Capability Flags in the "Capability Data" field from right to left are:

One bit "C" (31, IANA to assign)	CC mode supported
One bit "V" (30, IANA to assign)	CV mode supported
One bit "L" (29, IANA to assign)	PM Loss supported
One bit "D" (28, IANA to assign)	PM Delay supported
One bit "F" (27, IANA to assign)	FMS supported

Bits 8-26: This field MUST be set to zero on transmission and MUST be ignored on receipt.

The above bits can be set individually to indicate more than one kind of OAM capabilities at once, and the other reserved bits MUST be set to zero on transmission and MUST be ignored on receipt. Moreover, if CV flag is set, the CC flag MUST be set at the same time.

The MPLS-TP PW OAM Capability TLV MAY be included by a PE in an Initialization message to signal its peer that it supports the MPLS-TP PW OAM Capability. If the remote peer does not support the MPLS-TP PW OAM Capability TLV or the Initialization message sent by the remote peer does not include the MPLS-TP PW OAM Capability TLV, the resulting negotiation does not support MPLS-TP PW OAM capability. If instead the negotiation supports the MPLS-TP PW OAM capability, then the subsequent LDP Mapping message will carry the information of the MPLS-TP PW OAM configuration.

7.1.1. Backward Compatibility

If both the two T-PEs can recognize the MPLS-TP PW OAM Capability TLV,and CC or CV mode is supported, the BFD configuration procedure described in this document is adopted. Otherwise, if at least one of the two T-PEs do not support the CC or CV mode, the old VCCV BFD [RFC5885] will be performed. In this situation, the procedure

described in [RFC5885] MUST be followed: the C and V flags of MPLS-TP PW OAM Configuration TLV MUST NOT be set and the BFD Configuration sub-TLV MUST NOT be carried as a sub-TLV of MPLS-TP PW OAM Configuration TLV also.

The described behavior ensures full compatibility with the existing implementations.

7.2. MPLS-TP PW OAM Administration TLV

The format of the MPLS-TP PW OAM Administration TLV is as follows:

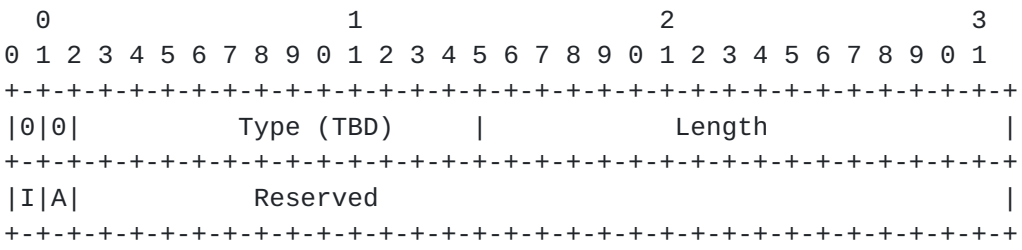

MPLS-TP PW OAM Administration TLV

One bit "I" (0, IANA to assign): "OAM MIP Entities Desired" is allocated. If the "OAM MIP entities desired" bit is set, it is indicating that the establishment of OAM MIP entities is required at every transit node of the signaled PW. If the establishment of a MIP is not supported, a Notification message MUST be sent with Status Code set to "MIP Configuration Failure".

One bit "A" (1, IANA to assign): "OAM Alarms Enabled" is allocated. If the "OAM Alarms Enabled" bit is set, it is indicating that the T-PE needs to enable OAM alarms.

Reserved (2-31 bits): This field MUST be set to zero on transmission and MUST be ignored on receipt.

7.3. MPLS-TP PW OAM Configuration TLV

The "MPLS-TP PW OAM Configuration TLV" is depicted in the following figure. It may be carried in the Mapping and Notification messages, just following the PW Status TLV.

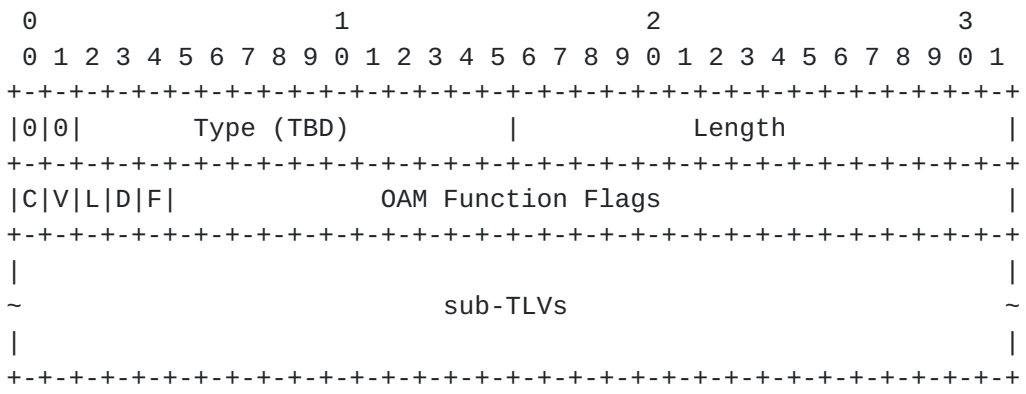

MPLS-TP PW OAM Configuration TLV

The "MPLS-TP PW OAM Configuration TLV" contains a number of flags indicating which OAM functions should be activated as well as OAM function specific sub-TLVs with configuration parameters for the particular functions.

Type: indicates a new type: the MPLS-TP PW OAM Configuration TLV (IANA to assign).

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

OAM Function Flags: a bitmap numbered from left to right as shown in the figure.

These flags are defined in this document:

OAM Function Flag bit#	Description
0 (C)	Continuity Check (CC)
1 (V)	Connectivity Verification (CV)
2 (L)	Performance Monitoring/Loss (PM/Loss)
3 (D)	Performance Monitoring/Delay (PM/Delay)
4 (F)	Fault Management Signals (FMS)
5-31	Reserved (set all to 0s)

Sub-TLVs corresponding to the different flags are as follows.

- "BFD Configuration sub-TLV", which MUST be included if the CC and/or the CV OAM Function flag is set. Furthermore, if the CV flag is set, the CC flag MUST be set at the same time.
- "Performance Monitoring sub-TLV", which MUST be included if the PM/Loss OAM Function flag is set.
- "MPLS-TP PW FMS sub-TLV", which MAY be included if the FMS OAM Function flag is set. If the "MPLS-TP PW FMS sub-TLV" is not included, default configuration values are used.

7.3.1. BFD Configuration sub-TLV

The "BFD Configuration sub-TLV is defined for BFD specific configuration parameters, which accommodates generic BFD OAM information and carries sub-TLVs.

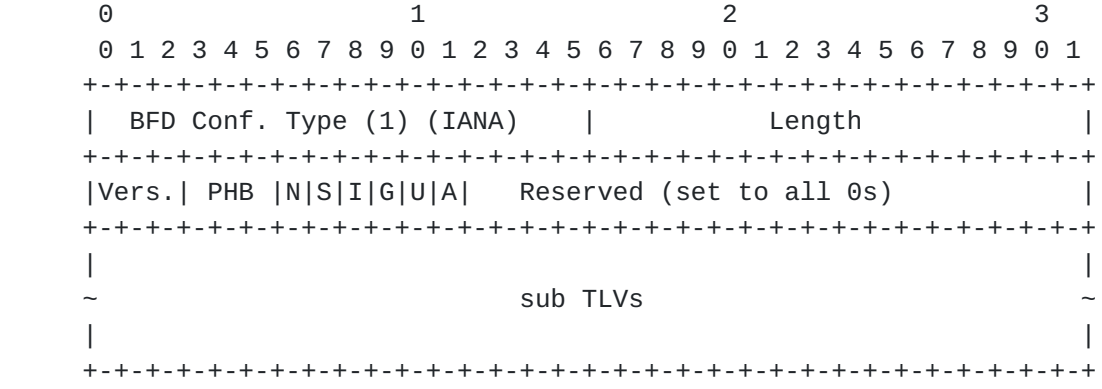

BFD Configuration sub-TLV

Type: indicates a new type, the "BFD Configuration sub-TLV" (IANA to define, suggested value 1).

Length: indicates the length of the TLV including sub-TLVs but excluding the Type and Length field, in octets.

Version: identifies the BFD protocol version. If a node does not support a specific BFD version, a Notification message MUST be generated with Status Code set to "Unsupported OAM Version".

PHB: Identifies the Per-Hop Behavior (PHB) to be used for periodic continuity monitoring messages.

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).

Encapsulation Capability (G): if set, it shows the capability of encapsulating BFD messages into G-Ach channel without IP/UDP headers. 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 G-Ach channel with IP/UDP headers. If both the G bit and U bit are set, configuration gives precedence to the G bit.

Operation mode (A): if set, it configures BFD in the associated mode. If it is not set it configures BFD in independent mode.

Reserved: Reserved for future specification and set to 0.

The "BFD Configuration sub-TLV" MUST include the following sub-TLVs in the Mapping message:
o "Local Discriminator sub-TLV".
o "Negotiation Timer Parameters sub-TLV" if the N flag is cleared.

7.3.1.1. Local Discriminator sub-TLV

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

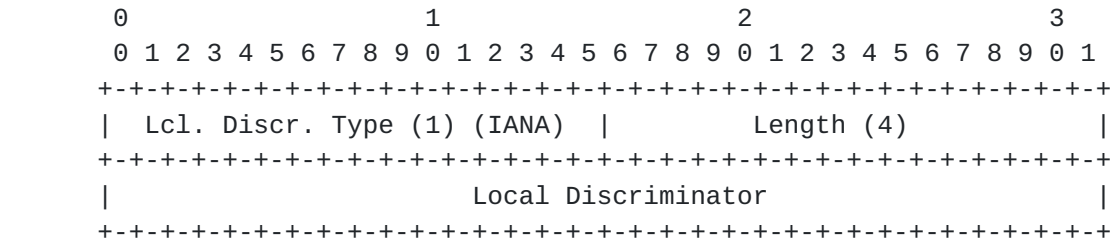

Local Discriminator sub-TLV

Type: indicates a new type, the "Local Discriminator sub-TLV" (IANA to define, suggested value 1).

Length: indicates the TLV total length in octets (4).

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.

7.3.1.2. 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.

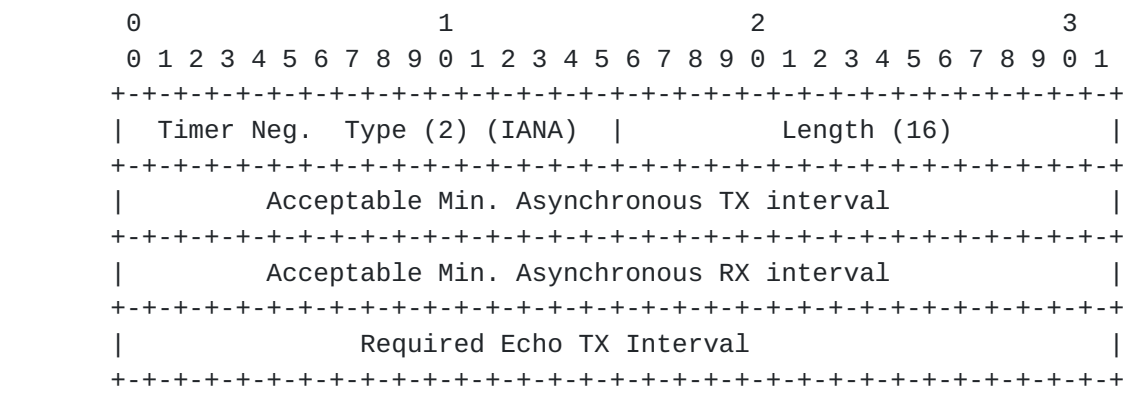

Negotiation Timer Parameters sub-TLV

Type: indicates a new type, the "Negotiation Timer Parameters sub-TLV" (IANA to define, suggested value 2).

Length: indicates the TLV total length in octets (16).

Acceptable Min. Asynchronous TX interval: in case of S (symmetric) flag set in the "BFD Configuration" TLV, it expresses the desired time interval (in microseconds) at which the T-PE initiating the signaling intends to both transmit and receive BFD periodic control packets. If the receiving T-PE can not support such value, it is allowed to reply back with an interval greater than the one proposed.

In case of S (symmetric) flag cleared in the "BFD Configuration sub-TLV", this field expresses the desired time interval (in microseconds) at which T-PE intends to transmit BFD periodic control packets in its transmitting direction.

Acceptable Min. Asynchronous RX interval: in case of S (symmetric) flag set in the "BFD Configuration sub-TLV", 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".

In case of S (symmetric) flag cleared in the "BFD Configuration sub-TLV", it expresses the minimum time interval (in microseconds) at which T-PE can receive BFD periodic control packets. In case this value is greater than the "Acceptable Min. Asynchronous TX interval" received from the other T-PE, such T-PE 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 [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 receiving system can not support this value a Notification MUST be generated with Status Code set to "Unsupported BFD TX Echo rate interval". By default the value is set to 0.

7.3.1.3. 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.

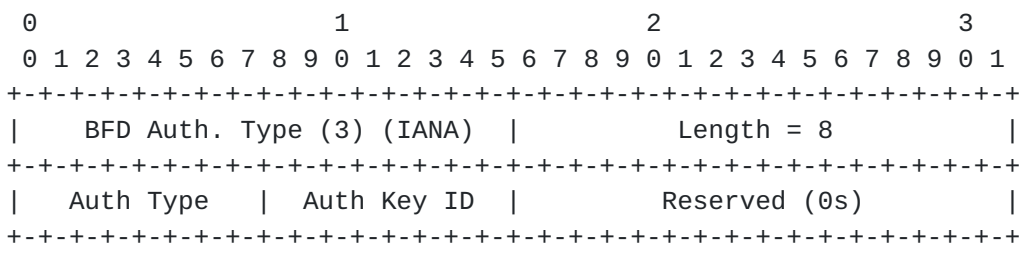

BFD Authentication sub-TLV

Type: indicates a new type, the "BFD Authentication sub-TLV" (IANA to define, suggested value 3).

Length: indicates the TLV total length in octets (8).

Auth Type: indicates which type of authentication to use. The same values as are defined in section 4.1 of [RFC5880] are used.

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.

Reserved: Reserved for future specification and set to 0.

7.3.2. Performance Monitoring sub-TLV

If the "MPLS-TP PW OAM Configuration TLV" has either the L (Loss), D (Delay) flag set, the "Performance Monitoring sub-TLV" MUST be present.

In case the values need to be different than the default ones the "MPLS-TP PW PM Loss sub-TLV, "MPLS-TP PW PM Delay sub-TLV" MAY be included:

- "MPLS-PW PM Loss sub-TLV" if the L flag is set in the "MPLS-TP PW OAM Configuration TLV";

- "MPLS-PW PM Dealy sub-TLV" if the D flag is set in the "MPLS-TP PW OAM Configuration TLV ".

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

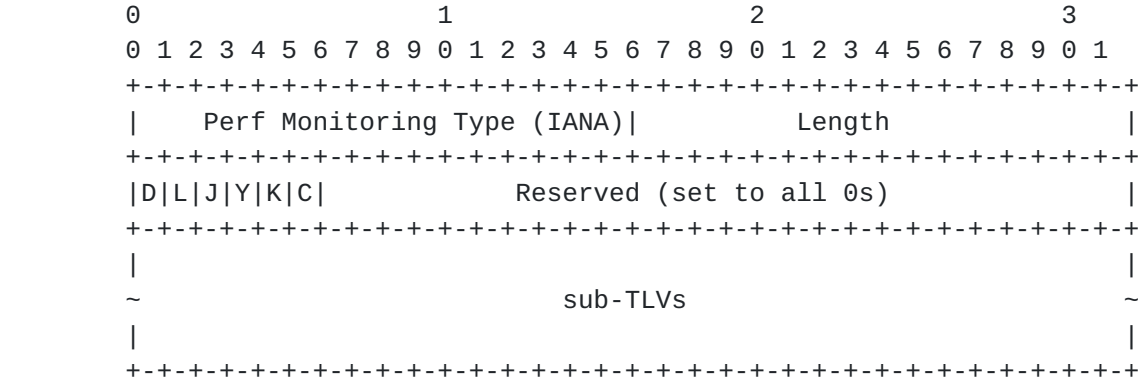

Performance Monitoring sub-TLV

- D: Delay inferred/direct (0=INFERRED, 1=DIRECT)
- L: Loss inferred/direct (0=INFERRED, 1=DIRECT)
- J: Delay variation/jitter (1=ACTIVE, 0=NOT ACTIVE)
- Y: Dyadic (1=ACTIVE, 0=NOT ACTIVE)
- K: Loopback (1=ACTIVE, 0=NOT ACTIVE)
- C: Combined (1=ACTIVE, 0=NOT ACTIVE)

7.3.2.1. MPLS-TP PW PM Loss TLV

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

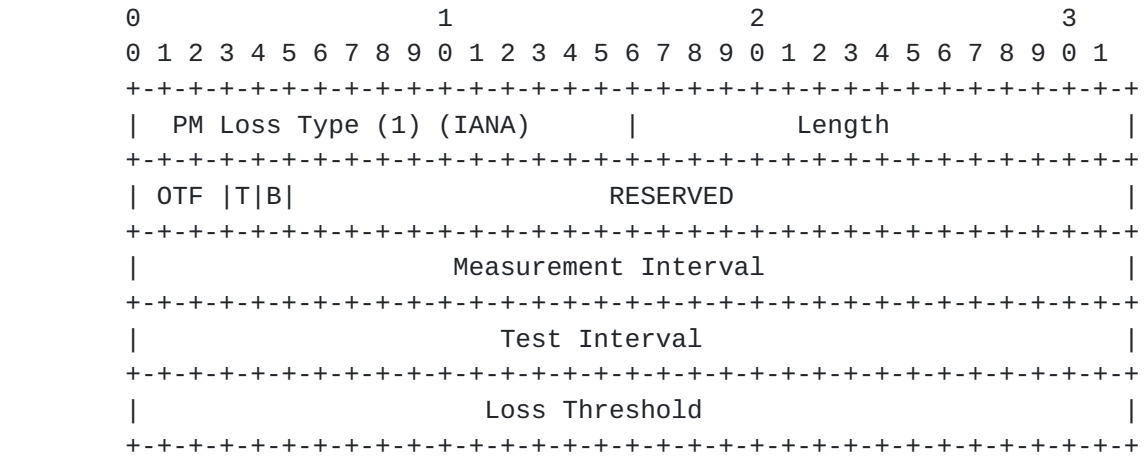

MPLS-TP PW PM Loss sub-TLV

Type: indicates a new type, the "MPLS-TP PW PM Loss sub-TLV" (IANA to define, suggested value 1).

Length: indicates the length of the parameters in octets.

OTF: Origin Timestamp Format of the Origin Timestamp field described in [I-D.ietf-mpls-loss-delay]. By default it is set to IEEE 1588 version 1.

Configuration Flags, please refer to [I-D.ietf-mpls-loss-delay] 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.

Measurement Interval: the time interval (in microseconds) at which LM query messages MUST be sent on both directions. If the T-PE receiving the Mapping message can not support such value, it can reply back with a higher interval. By default it is set to (TBD).

Test Interval: test messages interval as described in [I-D.ietf-mpls-loss-delay]. By default it is set to (TBD).

Loss Threshold: the threshold value of lost packets over which protections MUST be triggered. By default it is set to (TBD).

7.3.2.2. MPLS-TP PW PM Delay TLV

The "MPLS-TP PW PM Delay sub-TLV" depicted below is carried as a sub-TLV of the "MPLS-TP PW OAM Configuration TLV"

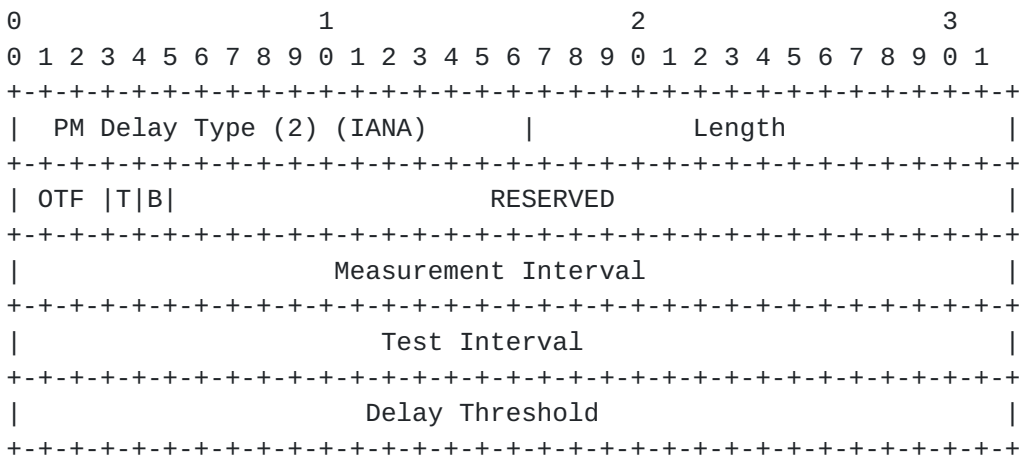

MPLS-TP PW PM Delay sub-TLV

Type: indicates a new type, the "MPLS-TP PW PM Delay sub-TLV" (IANA to define, suggested value 2).

Length: indicates the length of the parameters in octets.

OTF: Origin Timestamp Format of the Origin Timestamp field described in [I-D.ietf-mpls-loss-delay]. By default it is set to IEEE 1588 version 1.

Configuration Flags, please refer to [I-D.ietf-mpls-loss-delay] 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.

Measurement Interval: the time interval (in microseconds) at which LM query messages MUST be sent on both directions. If the T-PE receiving the Mapping message can not support such value, it can reply back with a higher interval. By default it is set to (TBD).

Test Interval: test messages interval as described in [I-D.ietf-mpls-loss-delay]. By default it is set to (TBD).

Delay Threshold: the threshold value of packet delay time over which protections MUST be triggered. By default it is set to (TBD).

7.3.3. MPLS-TP PW FMS TLV

The "MPLS-TP PW FMS sub-TLV" depicted below is carried as a sub-TLV of the "MPLS-TP PW OAM Configuration TLV".

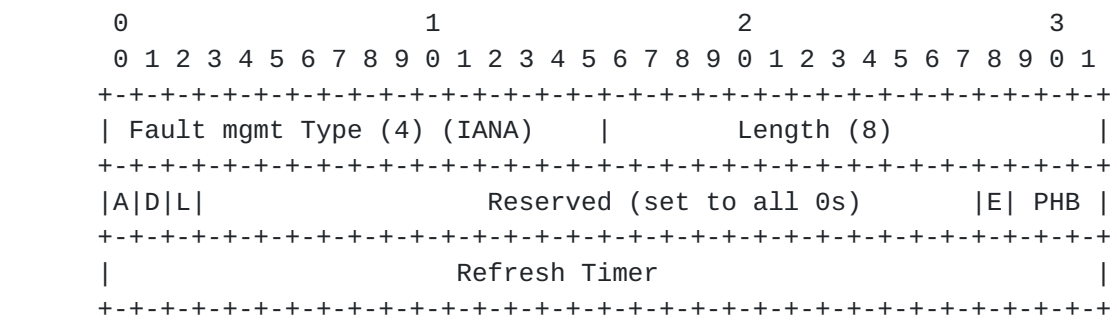

MPLS-TP PW FMS sub-TLV

Type: indicates a new type, the "MPLS-TP PW FMS sub-TLV" (IANA to define, suggested value 4).

Length: indicates the length of the parameters in octets (8).

Signal Flags: are used to enable the following signals:

- A: Alarm Indication Signal (AIS) as described in [I-D.ietf-mpls-tp-fault]
- D: Link Down Indication (LDI) as described in [I-D.ietf-mpls-tp-fault]
- L: Locked Report (LKR) as described in [I-D.ietf-mpls-tp-fault]
- Remaining bits: Reserved for future specification and set to 0.

Configuration Flags:

- E: used to enable/disable explicitly clearing faults
- PHB: identifies the per-hop behavior of packets with fault management information

Refresh Timer: indicates the refresh timer (in microseconds) of fault indication messages. If the T-PE receiving the Path message can not support such value, it can reply back with a higher interval.

8. IANA Considerations

This document specifies the following new LDP TLV types:

- MPLS-TP PW OAM Capability TLV;
- MPLS-TP PW OAM Administration TLV;
- MPLS-TP PW OAM Configuration TLV;

Sub-TLV types to be carried in the "MPLS-TP PW OAM Configuration

TLV":
- BFD Configuration sub-TLV;
- Performance Monitoring sub-TLV
- MPLS-TP PW FMS sub-TLV;

Sub-TLV types to be carried in the "BFD Configuration sub-TLV":
- Local Discriminator sub-TLV;
- Negotiation Timer Parameters sub-TLV.
- BFD Authentication sub-TLV

Sub-TLV types to be carried in the "Performance Monitoring sub-TLV":
- MPLS-TP PW PM Loss sub-TLV;
- MPLS-TP PW PM Loss sub-TLV.

8.1. OAM Configuration Errors

This document defines several new LDP status codes, IANA already maintains the registry "STATUS CODE NAME SPACE" defined by [RFC5036]. The following values are required to be assigned:

Range/Value	E	Description
0x0000003C	0	"MIP Configuration Failure"
0x0000003D	0	"Rejected MPLS-TP PW OAM Configuration TLV"
0x0000003E	0	"Wrong MPLS-TP PW OAM Configuration TLV"
0x0000003F	0	"Unsupported OAM Version"
0x0000004A	0	"Unsupported BFD TX Echo rate interval"

9. Security Considerations

Security considerations relating to LDP are described in section 5 of [RFC5036] and section 11 of [RFC5561]. Security considerations relating to use of LDP in setting up PWs is described in section 8 of [RFC4447].

This document defines new TLV/sub-TLV types, and OAM configuration procedures intended for use with MPLS-TP, which do not raise any additional security issues.

10. Acknowledgement

The authors would like to thank Andew Malis, Greg Mirsky, Luca Martini, Matthew Bocci, Thomas Nadeau for their valuable comments and discussions, especially would like to thank Eric Gray for his review of this document.

11. References

11.1. Normative references

[I-D.ietf-mpls-lsp-ping-mpls-tp-oam-conf]
Bellagamba, E., Andersson, L., Skoldstrom, P., Ward, D., and J. Drake, "Configuration of Pro-Active Operations, Administration, and Maintenance (OAM) Functions for MPLS-based Transport Networks using LSP Ping", draft-ietf-mpls-lsp-ping-mpls-tp-oam-conf-02 (work in progress), July 2011.

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

[RFC4447] Martini, L., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.

[RFC5085] Nadeau, T. and C. Pignataro, "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, December 2007.

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

[RFC6073] Martini, L., Metz, C., Nadeau, T., Bocci, M., and M. Aissaoui, "Segmented Pseudowire", RFC 6073, January 2011.

11.2. Informative References

[I-D.ietf-mpls-loss-delay]
Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", draft-ietf-mpls-loss-delay-04 (work in progress), July 2011.

[I-D.ietf-mpls-tp-cc-cv-rdi]
Allan, D., Swallow, G., and J. Drake, "Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile", draft-ietf-mpls-tp-cc-cv-rdi-06 (work in progress), August 2011.

[I-D.ietf-mpls-tp-fault]
Swallow, G., Fulignoli, A., Vigoureux, M., Boutros, S., and D. Ward, "MPLS Fault Management OAM",

draft-ietf-mpls-tp-fault-06 (work in progress), August 2011.

[I-D.ietf-mpls-tp-loss-delay-profile] Frost, D. and S. Bryant, "A Packet Loss and Delay Measurement Profile for MPLS-based Transport Networks", draft-ietf-mpls-tp-loss-delay-profile-04 (work in progress), July 2011.

[I-D.ietf-mpls-tp-oam-analysis] Sprecher, N. and L. Fang, "An Overview of the OAM Tool Set for MPLS based Transport Networks", draft-ietf-mpls-tp-oam-analysis-04 (work in progress), June 2011.

[I-D.ietf-mpls-tp-oam-framework] Allan, D., Busi, I., Niven-Jenkins, B., Fulignoli, A., Hernandez-Valencia, E., Levrau, L., Sestito, V., Sprecher, N., Helvoort, H., Vigoureux, M., Weingarten, Y., and R. Winter, "Operations, Administration and Maintenance Framework for MPLS-based Transport Networks", draft-ietf-mpls-tp-oam-framework-11 (work in progress), February 2011.

[I-D.ietf-pwe3-static-pw-status] Martini, L., Swallow, G., Heron, G., and M. Bocci, "Pseudowire Status for Static Pseudowires", draft-ietf-pwe3-static-pw-status-06 (work in progress), July 2011.

[RFC3985] Bryant, S. and P. Pate, "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.

[RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", RFC 5036, October 2007.

[RFC5561] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and JL. Le Roux, "LDP Capabilities", RFC 5561, July 2009.

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

[RFC5659] Bocci, M. and S. Bryant, "An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge", RFC 5659, October 2009.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection

(BFD)", RFC 5880, June 2010.

[RFC5885] Nadeau, T. and C. Pignataro, "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)", RFC 5885, June 2010.

Authors' Addresses

Fei Zhang (editor)
ZTE Corporation

Email: zhang.fei3@zte.com.cn

Bo Wu (editor)
ZTE Corporation

Email: wu.bo@zte.com.cn

Elisa Bellagamba (editor)
Ericsson
Farogatan 6
Kista, 164 40
Sweden

Phone: +46 761440785
Email: elisa.bellagamba@ericsson.com

Attila Takacs
Ericsson
Laborc u. 1.
Budapest, 1037
Hungary

Email: attila.takacs@ericsson.com

Xuehui Dai
ZTE Corporation

Email: dai.xuehui@zte.com.cn

Min Xiao
ZTE Corporation

Email: xiao.min2@zte.com.cn