MPLS Working Group
Internet Draft
Intended status: Informational

I. Busi (Ed)
Alcatel-Lucent
D. Allan (Ed)
Ericsson

Expires: March 17, 2011

September 17, 2010

Operations, Administration and Maintenance Framework for MPLS-based Transport Networks
draft-ietf-mpls-tp-oam-framework-08.txt

Abstract

The Transport Profile of Multi-Protocol Label Switching (MPLS-TP) is a packet-based transport technology based on the MPLS Traffic Engineering (MPLS-TE) and Pseudowire (PW) data plane architectures.

This document describes a framework to support a comprehensive set of Operations, Administration and Maintenance (OAM) procedures that fulfill the MPLS-TP OAM requirements for fault, performance and protection-switching management and that do not rely on the presence of a control plane.

This document is a product of a joint Internet Engineering Task Force (IETF) / International Telecommunications Union Telecommunication Standardization Sector (ITU-T) effort to include an MPLS Transport Profile within the IETF MPLS and PWE3 architectures to support the capabilities and functionalities of a packet transport network as defined by the ITU-T.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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".

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on March 17, 2011.

Copyright Notice

Table of Contents

Editors' Note:

This Informational Internet-Draft is aimed at achieving IETF Consensus before publication as an RFC and will be subject to an IETF Last Call.

[RFC Editor, please remove this note before publication as an RFC and insert the correct Streams Boilerplate to indicate that the published RFC has IETF Consensus.]

1. Introduction

As noted in the multi-protocol label switching (MPLS-TP) Framework RFCs (RFC 5921 [8] and [9]), MPLS-TP is a packet-based transport technology based on the MPLS Traffic Engineering (MPLS-TE) and Pseudo Wire (PW) data plane architectures defined in RFC 3031 [1], RFC 3985 [2] and RFC 5659 [4].

MPLS-TP supports a comprehensive set of Operations, Administration and Maintenance (OAM) procedures for fault, performance and protection-switching management and that do not rely on the presence of a control plane.

In line with [14], existing MPLS OAM mechanisms will be used wherever possible and extensions or new OAM mechanisms will be defined only where existing mechanisms are not sufficient to meet the requirements. Extensions do not deprecate support for existing MPLS OAM capabilities.

The MPLS-TP OAM framework defined in this document provides a comprehensive set of OAM procedures that satisfy the MPLS-TP OAM requirements of RFC 5860 [11]. In this regard, it defines similar OAM functionality as for existing SONET/SDH and OTN OAM mechanisms (e.g. [18]).

The MPLS-TP OAM framework is applicable to both LSPs and (MS-)PWs and supports co-routed and associated bidirectional p2p transport paths as well as unidirectional p2p and p2mp transport paths.

This document is a product of a joint Internet Engineering Task Force (IETF) / International Telecommunication Union Telecommunication Standardization Sector (ITU-T) effort to include an MPLS Transport Profile within the IETF MPLS and PWE3 architectures to support the capabilities and functionalities of a packet transport network as defined by the ITU-T.

1.1. Contributing Authors

Dave Allan, Italo Busi, Ben Niven-Jenkins, Annamaria Fulignoli, Enrique Hernandez-Valencia, Lieven Levrau, Vincenzo Sestito, Nurit Sprecher, Huub van Helvoort, Martin Vigoureux, Yaacov Weingarten, Rolf Winter

2. Conventions used in this document

2.1. Terminology

AC Attachment Circuit

DBN Domain Border Node

LER Label Edge Router

LME LSP Maintenance Entity

LMEG LSP ME Group

LSP Label Switched Path

LSR Label Switching Router

LSME LSP SPME ME

LSMEG LSP SPME ME Group

ME Maintenance Entity

MEG Maintenance Entity Group

MEP Maintenance Entity Group End Point

MIP Maintenance Entity Group Intermediate Point

PHB Per-hop Behavior

PME PW Maintenance Entity

PMEG PW ME Group

PSME PW SPME ME

PSMEG PW SPME ME Group

PW Pseudowire

SLA Service Level Agreement

SME Section Maintenance Entity Group

SPME Sub-path Maintenance Element

2.2. Definitions

This document uses the terms defined in RFC 5654 [5].

This document uses the term 'Per-hop Behavior' as defined in RFC 2474 [15].

This document uses the term LSP to indicate either a service LSP or a transport LSP (as defined in [8]).

Where appropriate, the following definitions are aligned with ITU-T recommendation Y.1731 [20] in order to have a common, unambiguous terminology. They do not however intend to imply a certain implementation but rather serve as a framework to describe the necessary OAM functions for MPLS-TP.

Adaptation function: The adaptation function is the interface between the client (sub)-layer and the server (sub-layer).

Data plane loopback: An out-of-service test where an interface at either an intermediate or terminating node in a path is placed into a data plane loopback state, such that all traffic (including user data and OAM) received on the looped back interface is sent on the reverse direction of the transport path.

Note - The only way to send an OAM packet to a node set in the data plane loopback mode is via TTL expiry, irrespectively on whether the node is hosting MIPs or MEPs.

Domain Border Node (DBN): An intermediate node in an MPLS-TP LSP that is at the boundary between two MPLS-TP OAM domains. Such a node may be present on the edge of two domains or may be connected by a link to the DBN at the edge of another OAM domain.

Down MEP: A MEP that receives OAM packets from, and transmits them towards, the direction of a server layer.

In-Service: The administrative status of a transport path when it is unlocked.

Intermediate Node: An intermediate node transits traffic for an LSP or a PW. An intermediate node may originate OAM flows directed to downstream intermediate nodes or MEPs.

Loopback: See data plane loopback and OAM loopback definitions.

Maintenance Entity (ME): Some portion of a transport path that requires management bounded by two points (called MEPs), and the relationship between those points to which maintenance and monitoring operations apply (details in section 3.1).

Maintenance Entity Group (MEG): The set of one or more maintenance entities that maintain and monitor a transport path in an OAM domain.

MEP: A MEG end point (MEP) is capable of initiating (MEP Source) and terminating (MEP Sink) OAM messages for fault management and performance monitoring. MEPs define the boundaries of an ME (details in section 3.3).

MEP Source: A MEP acts as MEP source for an OAM message when it originates and inserts the message into the transport path for its associated MEG.

MEP Sink: A MEP acts as a MEP sink for an OAM message when it terminates and processes the messages received from its associated MEG.

MIP: A MEG intermediate point (MIP) terminates and processes OAM messages that are sent to this particular MIP and may generate OAM messages in reaction to received OAM messages. It never generates unsolicited OAM messages itself. A MIP resides within a MEG between MEPs (details in section 3.3).

MPLS-TP Section: As defined in [8], it is the link traversed by an MPLS-TP LSP.

OAM domain: A domain, as defined in [5], whose entities are grouped for the purpose of keeping the OAM confined within that domain.

Note - within the rest of this document the term "domain" is used to indicate an "OAM domain"

OAM flow: Is the set of all OAM messages originating with a specific MEP source that instrument one direction of a MEG.

OAM information element: An atomic piece of information exchanged between MEPs and/or MIPs in MEG used by an OAM application.

OAM loopback: It is the capability of a node to be directed by a received OAM message to generate a reply back to the sender. OAM loopback can work in-service and can support different OAM functions (e.g., bidirectional on-demand connectivity verification).

OAM Message: One or more OAM information elements that when exchanged between MEPs or between MEPs and MIPs performs some OAM functionality (e.g. connectivity verification)

OAM Packet: A packet that carries one or more OAM messages (i.e. OAM information elements).

Out-of-Service: The administrative status of a transport path when it is locked. When a path is in a locked condition, it is blocked from carrying client traffic.

Path Segment: It is either a segment or a concatenated segment, as defined in RFC 5654 [5].

Signal Degrade: A condition declared by a MEP when the data forwarding capability associated with a transport path has deteriorated, as determined by PM. See also ITU-T recommendation G.806 [13].

Signal Fail: A condition declared by a MEP when the data forwarding capability associated with a transport path has failed, e.g. loss of continuity. See also ITU-T recommendation G.806 [13].

Tandem Connection: A tandem connection is an arbitrary part of a transport path that can be monitored (via OAM) independent of the end-to-end monitoring (OAM). The tandem connection may also include the forwarding engine(s) of the node(s) at the boundaries of the tandem connection. Tandem connections may be nested but cannot overlap. See also ITU-T recommendation G.805 [19].

Up MEP: A MEP that transmits OAM packets towards, and receives them from, the direction of the forwarding engine.

3. Functional Components

MPLS-TP is a packet-based transport technology based on the MPLS and PW data plane architectures ([1], [2] and [4]) and is capable of transporting service traffic where the characteristics of information transfer between the transport path endpoints can be demonstrated to comply with certain performance and quality guarantees.

In order to describe the required OAM functionality, this document introduces a set of functional components.

3.1. Maintenance Entity and Maintenance Entity Group

MPLS-TP OAM operates in the context of Maintenance Entities (MEs) that define a relationship between any two points of a transport path to which maintenance and monitoring operations apply. The collection of one or more MEs that belongs to the same transport path and that are maintained and monitored as a group are known as a maintenance entity group (MEG) and the two points that define a maintenance entity are called Maintenance Entity Group (MEG) End Points (MEPs). In between these two points zero or more intermediate points, called Maintenance Entity Group Intermediate Points (MIPs), can exist and can be shared by more than one ME in a MEG.

An abstract reference model for an ME is illustrated in Figure 1 below:

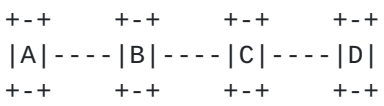

Figure 1 ME Abstract Reference Model

The instantiation of this abstract model to different MPLS-TP entities is described in section 4. In Figure 1, nodes A and D can be LERs for an LSP or the T-PEs for a MS-PW, nodes B and C are LSRs for a LSP or S-PEs for a MS-PW. MEPs reside in nodes A and D while MIPs reside in nodes B and C and may reside in A and D. The links connecting adjacent nodes can be physical links, (sub-)layer LSPs/SPMEs, or serving layer paths.

This functional model defines the relationships between all OAM entities from a maintenance perspective, to allow each Maintenance Entity to monitor and manage the (sub-)layer network under its responsibility and to localize problems efficiently.

An MPLS-TP Maintenance Entity Group may be defined to monitor the transport path for fault and/or performance management.

The MEPs that form a MEG bound the scope of an OAM flows to the MEG (i.e. within the domain of the transport path that is being monitored and managed). There are two exceptions to this:

1) A misbranching fault may cause OAM packets to be delivered to a MEP that is not in the MEG of origin.

2) An out-of-band return path may be used between a MIP or a MEP and the originating MEP.

In case of unidirectional point-to-point transport paths, a single unidirectional Maintenance Entity is defined to monitor it.

In case of associated bi-directional point-to-point transport paths, two independent unidirectional Maintenance Entities are defined to independently monitor each direction. This has implications for transactions that terminate at or query a MIP, as a return path from MIP to source MEP does not necessarily exist in the MEG.

In case of co-routed bi-directional point-to-point transport paths, a single bidirectional Maintenance Entity is defined to monitor both directions congruently.

In case of unidirectional point-to-multipoint transport paths, a single unidirectional Maintenance entity for each leaf is defined to monitor the transport path from the root to that leaf.

In all cases, portions of the transport path may be monitored by the instantiation of SPMEs (see section 3.2).

The reference model for the p2mp MEG is represented in Figure 2.

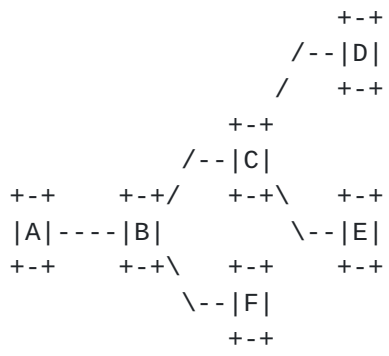

Figure 2 Reference Model for p2mp MEG

In case of p2mp transport paths, the OAM measurements are independent for each ME (A-D, A-E and A-F):

- Fault conditions - some faults may impact more than one ME depending from where the failure is located;
- Packet loss - packet dropping may impact more than one ME depending from where the packets are lost;
- Packet delay - will be unique per ME.

Each leaf (i.e. D, E and F) terminates OAM flows to monitor the ME between itself and the root while the root (i.e. A) generates OAM messages common to all the MEs of the p2mp MEG. All nodes may implement a MIP in the corresponding MEG.

3.2. Nested MEGs: SPMEs and Tandem Connection Monitoring

In order to verify and maintain performance and quality guarantees, there is a need to not only apply OAM functionality on a transport path granularity (e.g. LSP or MS-PW), but also on arbitrary parts of transport paths, defined as Tandem Connections, between any two arbitrary points along a transport path.

Sub-path Maintenance Elements (SPMEs), as defined in [8], are instantiated to provide monitoring of a portion of a set of co-routed transport paths (LSPs or MS-PWs). The operational aspects of instantiating SPMEs are out of scope of this memo.

SPMEs can also be employed to meet the requirement to provide tandem connection monitoring (TCM).

TCM for a given path segment of a transport path is implemented by creating an SPME that has a 1:1 association with the path segment of the transport path that is to be monitored.

In the TCM case, this means that the SPME used to provide TCM can carry only one and only one transport path thus allowing direct correlation between all fault management and performance monitoring information gathered for the SPME and the monitored path segment of the end-to-end transport path. The SPME is monitored using normal LSP monitoring.

Where resiliency is required across an arbitrary portion of a transport path, this may be implemented by more than one diversely routed SPMEs with common end points where only one SPME is active at any given time.

There are a number of implications to this approach:

1) The SPME would use the uniform model of TC code point copying between sub-layers for diffserv such that the E2E markings and PHB treatment for the transport path was preserved by the SPMEs.

2) The SPME normally would use the short-pipe model for TTL handling [6] such that MIP addressing for the E2E entity would be not be impacted by the presence of the SPME, but it should be possible for an operator to specify use of the uniform model.

3) PM statistics need to be adjusted for the encapsulation overhead of the additional SPME sub-layer.

Note that points 1 an 2 above assume that the TTL copying mode and TC copying modes are independently configurable for an LSP.

There are specific issues with the use of the uniform model of TTL copying for an SPME:

1. As any MIP in the SPME sub-layer is not part of the transport path MEG, hence only an out of band return path would be available.

2. The instantiation of a lower level MEG or protection switching actions within a lower level MEG may change the TTL distances to MIPs in the higher level MEGs.

The endpoints of the SPME are MEPs and limit the scope of an OAM flow within each MEG to the MEPs belong to (i.e. within the domain of the SPME that is being monitored and managed).

When considering SPMEs, it is important to consider that the following properties apply to all MPLS-TP MEGs:

- They can be nested but not overlapped, e.g. a MEG may cover a segment or a concatenated segment of another MEG, and may also include the forwarding engine(s) of the node(s) at the edge(s) of the segment or concatenated segment. However when MEGs are nested, the MEPs and MIPs in the nested MEG are no longer part of the encompassing MEG.

- It is possible that MEPs of nested MEGs reside on a single node but again implemented in such a way that they do not overlap.

- Each OAM flow is associated with a single MEG

- OAM packets that instrument a particular direction of a transport path are subject to the same forwarding treatment (i.e. fate share) as the data traffic and in some cases may be required to have common queuing discipline E2E with the class of traffic monitored. OAM packets can be distinguished from the data traffic using the GAL and ACH constructs [7] for LSP and Section or the ACH construct [3]and [7] for (MS-)PW.

- When a SPME is instantiated after the transport path has been instantiated the TTL addressing of the MIPs will change.

3.3. MEG End Points (MEPs)

MEG End Points (MEPs) are the source and sink points of a MEG. In the context of an MPLS-TP LSP, only LERs can implement MEPs while in the context of an SPME LSRs for the MPLS-TP LSP can be LERs for SPMEs that contribute to the overall monitoring infrastructure for the transport path. Regarding PWs, only T-PEs can implement MEPs while for SPMEs supporting one or more PWs both T-PEs and S-PEs can implement SPME MEPs. Any MPLS-TP LSR can implement a MEP for an MPLS-TP Section.

MEPs are responsible for activating and controlling all of the proactive and on-demand monitoring OAM functionality for the MEG. There is a separate class of notifications (such as LKR and AIS) that are originated by intermediate nodes and triggered by server layer events. A MEP is capable of originating and terminating OAM messages for fault management and performance monitoring. These OAM messages are encapsulated into an OAM packet using the G-ACh as defined in RFC 5586 [7]. In this case the G-ACh message is an OAM message and the channel type

indicates an OAM message. A MEP terminates all the OAM packets it receives from the MEG it belongs to and silently discards those that do not (note in the case of a mis-connectivity defect there are further actions taken). The MEG the OAM packet belongs to is inferred from the MPLS or PW label or, in case of an MPLS-TP section, the MEG is inferred from the port on which an OAM packet was received with the GAL at the top of the label stack.

OAM packets may require the use of an available "out-of-band" return path (as defined in [8]). In such cases sufficient information is required in the originating transaction such that the OAM reply packet can be constructed (e.g. IP address).

Each OAM solution will further detail its applicability as a pro-active or on-demand mechanism as well as its usage when:

o The "in-band" return path exists and it is used;

o An "out-of-band" return path exists and it is used;

o Any return path does not exist or is not used.

Once a MEG is configured, the operator can configure which proactive OAM functions to use on the MEG but the MEPs are always enabled. A node at the edge of a MEG always supports a MEP.

MEPs terminate all OAM packets received from the associated MEG. As the MEP corresponds to the termination of the forwarding path for a MEG at the given (sub-)layer, OAM packets never leak outside of a MEG in a properly configured fault-free implementation.

A MEP of an MPLS-TP transport path coincides with transport path termination and monitors it for failures or performance degradation (e.g. based on packet counts) in an end-to-end scope. Note that both MEP source and MEP sink coincide with transport paths' source and sink terminations.

The MEPs of an SPME are not necessarily coincident with the termination of the MPLS-TP transport path and monitor a path segment of the transport path for failures or performance degradation (e.g. based on packet counts) only within the boundary of the MEG for the SPME.

An MPLS-TP MEP sink passes a fault indication to its client (sub-)layer network as a consequent action of fault detection.

A node at the edge of a MEG can either support per-node MEP or per-interface MEP(s). A per-node MEP resides in an unspecified location within the node while a per-interface MEP resides on a specific side of the forwarding engine. In particular a per-interface MEP is called "Up MEP" or "Down MEP" depending on its location relative to the forwarding engine.

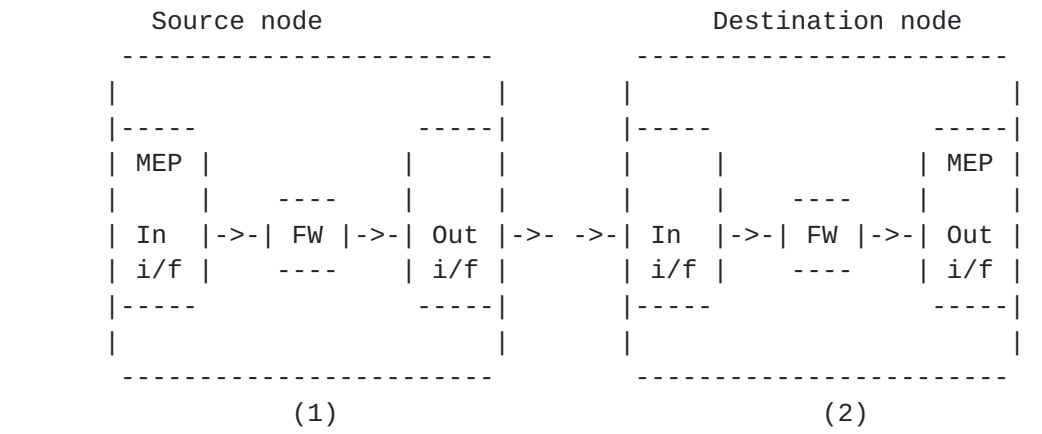

Figure 3 Example of per-interface Up MEPs

Figure 3 describes two examples of per-interface Up MEPs: An Up Source MEP in a source node (case 1) and an Up Sink MEP in a destination node (case 2).

The usage of per-interface Up MEPs extends the coverage of the ME for both fault and performance monitoring closer to the edge of the domain and allows the isolation of failures or performance degradation to being within a node or either the link or interfaces.

Each OAM solution will further detail the implications when used with per-interface or per-node MEPs, if necessary.

It may occur that the Up MEPs of an SPME are set on both sides of the forwarding engine such that the MEG is entirely internal to the node.

It should be noted that a ME may span nodes that implement per node MEPs and per-interface MEPs. This guarantees backward compatibility with most of the existing LSRs that can implement only a per-node MEP as in current implementations label operations are largely performed on the ingress interface, hence the exposure of the GAL as top label will occur at the ingress interface.

Note that a MEP can only exist at the beginning and end of a (sub-)layer in MPLS-TP. If there is a need to monitor some portion of that LSP or PW, a new sub-layer in the form of an SPME is created which permits MEPs and associated MEGs to be created.

In the case where an intermediate node sends a message to a MEP, it uses the top label of the stack at that point.

3.4. MEG Intermediate Points (MIPs)

A MEG Intermediate Point (MIP) is a function located at a point between the MEPs of a MEG for a PW, LSP or SPME.

A MIP is capable of reacting to some OAM packets and forwarding all the other OAM packets while ensuring fate sharing with data plane packets. However, a MIP does not initiate unsolicited OAM packets, but may be addressed by OAM packets initiated by one of the MEPs of the MEG. A MIP can generate OAM packets only in response to OAM packets that are sent on the MEG it belongs to. The OAM messages generated by the MIP are sent in the direction of the source MEP and not forwarded to the sink MEP.

An intermediate node within a MEG can either:

- Support per-node MIP (i.e. a single MIP per node in an unspecified location within the node);

- Support per-interface MIP (i.e. two or more MIPs per node on both sides of the forwarding engine).

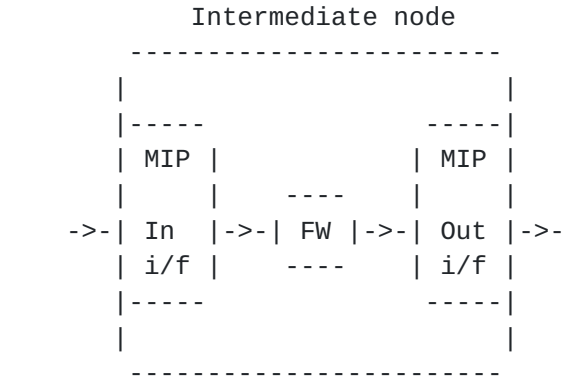

Figure 4 Example of per-interface MIPs

Figure 4 describes an example of two per-interface MIPs at an intermediate node of a point-to-point MEG.

The usage of per-interface MIPs allows the isolation of failures or performance degradation to being within a node or either the link or interfaces.

When sending an OAM packet to a MIP, the source MEP should set the TTL field to indicate the number of hops necessary to reach the node where the MIP resides.

The source MEP should also include Target MIP information in the OAM packets sent to a MIP to allow proper identification of the MIP within the node. The MEG the OAM packet is associated with is inferred from the MPLS label.

A node at the edge of a MEG can also support per-interface Up MEPs and per-interface MIPs on either side of the forwarding engine.

Once a MEG is configured, the operator can enable/disable the MIPs on the nodes within the MEG. All the intermediate nodes and possibly the end nodes host MIP(s). Local policy allows them to be enabled per function and per MEG. The local policy is controlled by the management system, which may delegate it to the control plane.

3.5. Server MEPs

A server MEP is a MEP of a MEG that is either:

- Defined in a layer network that is "below", which is to say encapsulates and transports the MPLS-TP layer network being referenced, or

- Defined in a sub-layer of the MPLS-TP layer network that is "below" which is to say encapsulates and transports the sub-layer being referenced.

A server MEP can coincide with a MIP or a MEP in the client (MPLS-TP) (sub-)layer network.

A server MEP also provides server layer OAM indications to the client/server adaptation function between the client (MPLS-TP) (sub-)layer network and the server (sub-)layer network. The adaptation function maintains state on the mapping of MPLS-TP transport paths that are setup over that server (sub-)layer's transport path.

For example, a server MEP can be either:

- A termination point of a physical link (e.g. 802.3), an SDH VC or OTN ODU, for the MPLS-TP Section layer network, defined in section 4.1;

- An MPLS-TP Section MEP for MPLS-TP LSPs, defined in section 4.2;

- An MPLS-TP LSP MEP for MPLS-TP PWs, defined in section 4.3;

- An MPLS-TP SPME MEP used for LSP path segment monitoring, as defined in section 4.4, for MPLS-TP LSPs or higher-level SPMEs providing LSP path segment monitoring;

- An MPLS-TP SPME MEP used for PW path segment monitoring, as defined in section 4.5, for MPLS-TP PWs or higher-level SPMEs providing PW path segment monitoring.

The server MEP can run appropriate OAM functions for fault detection within the server (sub-)layer network, and provides a fault indication to its client MPLS-TP layer network via the client/server adaptation function. When the server layer is not MPLS-TP, server MEP OAM functions are outside the scope of this document.

3.6. Configuration Considerations

When a control plane is not present, the management plane configures these functional components. Otherwise they can be configured either by the management plane or by the control plane.

Local policy allows disabling the usage of any available "out-of-band" return path, as defined in [8], irrespective of what is requested by the node originating the OAM packet.

SPMEs are usually instantiated when the transport path is created by either the management plane or by the control plane (if present). Sometimes an SPME can be instantiated after the transport path is initially created.

3.7. P2MP considerations

All the traffic sent over a p2mp transport path, including OAM packets generated by a MEP, is sent (multicast) from the root to all the leaves. As a consequence:

- To send an OAM packet to all leaves, the source MEP can send a single OAM packet that will be delivered by the forwarding plane to all the leaves and processed by all the leaves.

- To send an OAM packet to a single leaf, the source MEP sends a single OAM packet that will be delivered by the forwarding plane to all the leaves but contains sufficient information to identify a target leaf, and therefore is processed only by the target leaf and ignored by the other leaves.

- To send an OAM packet to a single MIP, the source MEP sends a single OAM packet with the TTL field indicating the number of hops necessary to reach the node where the MIP resides. This packet will be delivered by the forwarding plane to all intermediate nodes at the same TTL distance of the target MIP and to any leaf that is located at a shorter distance. The OAM message must contain sufficient information to identify the target MIP and therefore is processed only by the target MIP.

- In order to send an OAM packet to M leaves (i.e., a subset of all the leaves), the source MEP sends M different OAM packets targeted to each individual leaf in the group of M leaves. Aggregated or subsetting mechanisms are outside the scope of this document.

P2MP paths are unidirectional, therefore any return path to a source MEP for on-demand transactions will be out-of-band. A mechanism to scope the set of MEPs or MIPs expected to respond to a given "on-demand" transaction is useful as it relieves the source MEP of the requirement to filter and discard undesired responses as normally TTL exhaustion will address all MIPs at a given distance from the source, and failure to exhaust TTL will address all MEPs.

4. Reference Model

The reference model for the MPLS-TP framework builds upon the concept of a MEG, and its associated MEPs and MIPs, to support the functional requirements specified in RFC 5860 [11].

The following MPLS-TP MEGs are specified in this document:

- A Section Maintenance Entity Group (SME), allowing monitoring and management of MPLS-TP Sections (between MPLS LSRs).

- An LSP Maintenance Entity Group (LME), allowing monitoring and management of an end-to-end LSP (between LERs).

- A PW Maintenance Entity Group (PME), allowing monitoring and management of an end-to-end SS/MS-PWs (between T-PEs).

- An LSP SPME ME Group (LSMEG), allowing monitoring and management of an SPME (between any LERs/LSRs along an LSP).

- A PW SPME ME Group (PSMEG), allowing monitoring and management of an SPME (between any T-PEs/S-PEs along the (MS-)PW).

The MEGs specified in this MPLS-TP framework are compliant with the architecture framework for MPLS-TP MS-PWs [4] and LSPs [1].

Hierarchical LSPs are also supported in the form of SPMEs. In this case, each LSP in the hierarchy is a different sub-layer network that can be monitored, independently from higher and lower level LSPs in the hierarchy, on an end-to-end basis (from LER to LER) by a SPME. It is possible to monitor a portion of a hierarchical LSP by instantiating a hierarchical SPME between any LERs/LSRs along the hierarchical LSP.

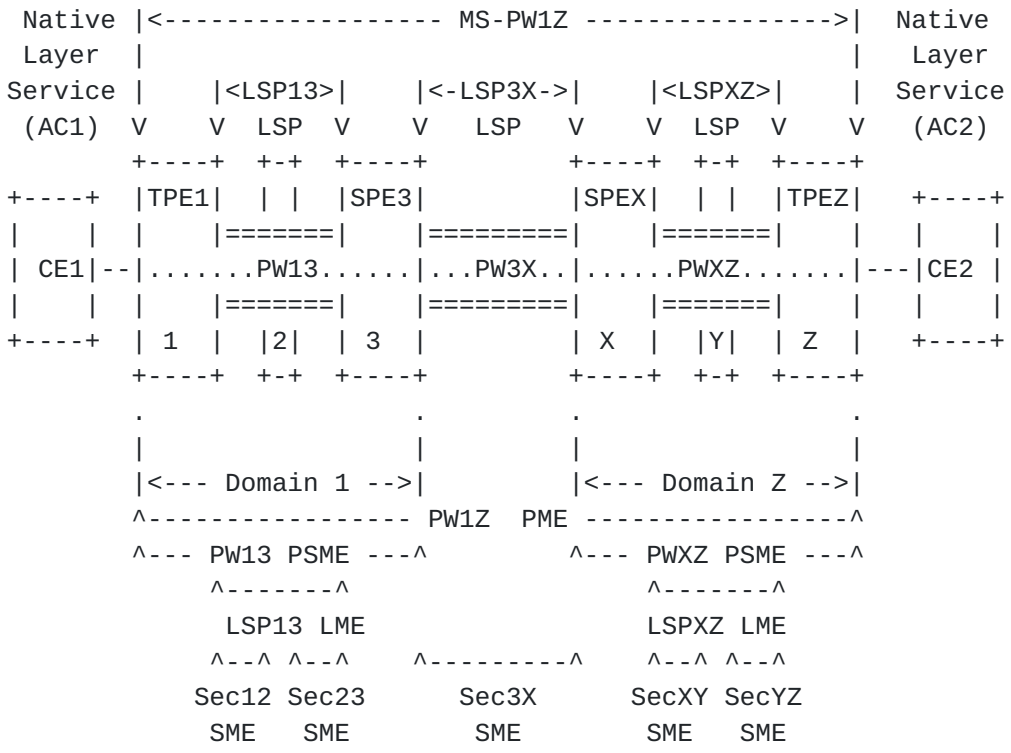

TPE1: Terminating Provider Edge 1 SPE2: Switching Provider Edge 3
TPEX: Terminating Provider Edge X SPEZ: Switching Provider Edge Z

^---^ ME ^ MEP ==== LSP .... PW

Figure 5 Reference Model for the MPLS-TP OAM Framework

Figure 5 depicts a high-level reference model for the MPLS-TP OAM framework. The figure depicts portions of two MPLS-TP enabled network domains, Domain 1 and Domain Z. In Domain 1, LSR1 is adjacent to LSR2 via the MPLS-TP Section Sec12 and LSR2 is adjacent to LSR3 via the MPLS-TP Section Sec23. Similarly, in Domain Z, LSRX is adjacent to LSRY via the MPLS-TP Section SecXY and LSRY is adjacent to LSRZ via the MPLS-TP Section SecYZ. In addition, LSR3 is adjacent to LSRX via the MPLS-TP Section 3X.

Figure 5 also shows a bi-directional MS-PW (PW1Z) between AC1 on TPE1 and AC2 on TPEZ. The MS-PW consists of three bi-directional PW path segments: 1) PW13 path segment between T-PE1 and S-PE3 via the bi-directional LSP13 LSP, 2) PW3X path segment between S-PE3 and S-PEX, via the bi-directional LSP3X LSP, and 3) PWXZ path segment between S-PEX and T-PEZ via the bi-directional LSPXZ LSP.

The MPLS-TP OAM procedures that apply to a MEG are expected to operate independently from procedures on other MEGs. Yet, this does not preclude that multiple MEGs may be affected simultaneously by the same network condition, for example, a fiber cut event.

Note that there are no constrains imposed by this OAM framework on the number, or type (p2p, p2mp, LSP or PW), of MEGs that may be instantiated on a particular node. In particular, when looking at Figure 5, it should be possible to configure one or more MEPs on the same node if that node is the endpoint of one or more MEGs.

Figure 5 does not describe a PW3X PSME because typically SPMEs are used to monitor an OAM domain (like PW13 and PWXZ PSMEs) rather than the segment between two OAM domains. However the OAM framework does not pose any constraints on the way SPMEs are instantiated as long as they are not overlapping.

The subsections below define the MEGs specified in this MPLS-TP OAM architecture framework document. Unless otherwise stated, all references to domains, LSRs, MPLS-TP Sections, LSPs, pseudowires and MEGs in this section are made in relation to those shown in Figure 5.

4.1. MPLS-TP Section Monitoring (SME)

An MPLS-TP Section ME (SME) is an MPLS-TP maintenance entity intended to monitor an MPLS-TP Section as defined in RFC 5654 [5]. An SME may be configured on any MPLS-TP section. SME OAM packets must fate share with the user data packets sent over the monitored MPLS-TP Section.

An SME is intended to be deployed for applications where it is preferable to monitor the link between topologically adjacent (next hop in this layer network) MPLS-TP LSRs rather than monitoring the individual LSP or PW path segments traversing the MPLS-TP Section and the server layer technology does not provide adequate OAM capabilities.

Figure 5 shows five Section MEs configured in the network between AC1 and AC2:

1. Sec12 ME associated with the MPLS-TP Section between LSR 1 and LSR 2,

2. Sec23 ME associated with the MPLS-TP Section between LSR 2 and LSR 3,

3. Sec3X ME associated with the MPLS-TP Section between LSR 3 and LSR X,

4. SecXY ME associated with the MPLS-TP Section between LSR X and LSR Y, and

5. SecYZ ME associated with the MPLS-TP Section between LSR Y and LSR Z.

4.2. MPLS-TP LSP End-to-End Monitoring (LME)

An MPLS-TP LSP ME (LME) is an MPLS-TP maintenance entity intended to monitor an end-to-end LSP between two LERs. An LME may be configured on any MPLS LSP. LME OAM packets must fate share with user data packets sent over the monitored MPLS-TP LSP.

An LME is intended to be deployed in scenarios where it is desirable to monitor an entire LSP between its LERs, rather than, say, monitoring individual PWs.

Figure 5 depicts two LMEs configured in the network between AC1 and AC2: 1) the LSP13 LME between LER 1 and LER 3, and 2) the LSPXZ LME between LER X and LER Y. Note that the presence of a LSP3X LME in such a configuration is optional, hence, not precluded by this framework. For instance, the SPs may prefer to monitor the MPLS-TP Section between the two LSRs rather than the individual LSPs.

4.3. MPLS-TP PW Monitoring (PME)

An MPLS-TP PW ME (PME) is an MPLS-TP maintenance entity intended to monitor a SS-PW or MS-PW between a pair of T-PEs. A PME can be configured on any SS-PW or MS-PW. PME OAM packets must fate share with the user data packets sent over the monitored PW.

A PME is intended to be deployed in scenarios where it is desirable to monitor an entire PW between a pair of MPLS-TP enabled T-PEs rather than monitoring the LSP aggregating multiple PWs between PEs.

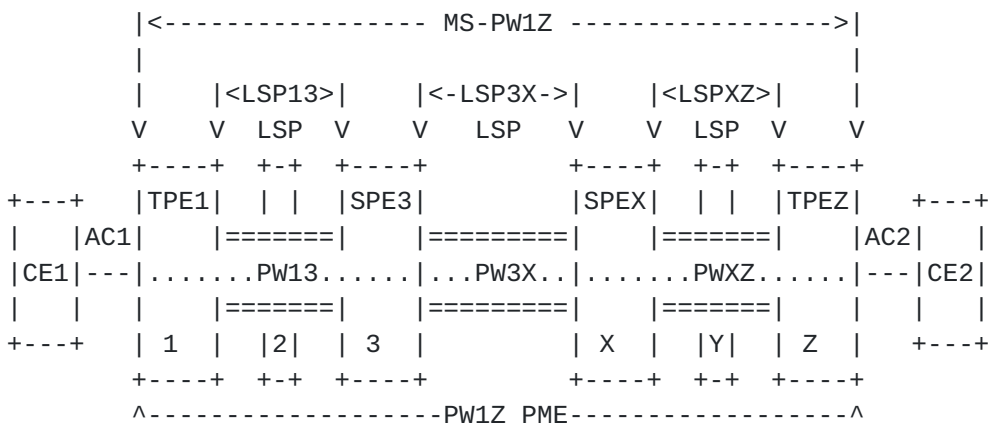

Figure 6 MPLS-TP PW ME (PME)

Figure 6 depicts a MS-PW (MS-PW1Z) consisting of three path segments: PW13, PW3X and PWXZ and its associated end-to-end PME (PW1Z PME).

4.4. MPLS-TP LSP SPME Monitoring (LSME)

An MPLS-TP LSP SPME ME (LSME) is an MPLS-TP SPME with associated maintenance entity intended to monitor an arbitrary part of an LSP between the pair of MEPs instantiated for the SPME independent from the end-to-end monitoring (LME). An LSME can monitor an LSP segment or concatenated segment and it may also include the forwarding engine(s) of the node(s) at the edge(s) of the segment or concatenated segment.

When SPME is established between non-adjacent LSRs, the edges of the SPME becomes adjacent at the LSP sub-layer network and any LSR that were previously in between becomes an LSR for the SPME.

Multiple hierarchical LSMEs can be configured on any LSP. LSME OAM packets must fate share with the user data packets sent over the monitored LSP path segment.

A LSME can be defined between the following entities:

o The end node and any intermediate node of a given LSP.

o Any two intermediate nodes of a given LSP.

An LSME is intended to be deployed in scenarios where it is preferable to monitor the behaviour of a part of an LSP or set of LSPs rather than the entire LSP itself, for example when there is a need to monitor a part of an LSP that extends beyond

the administrative boundaries of an MPLS-TP enabled administrative domain.

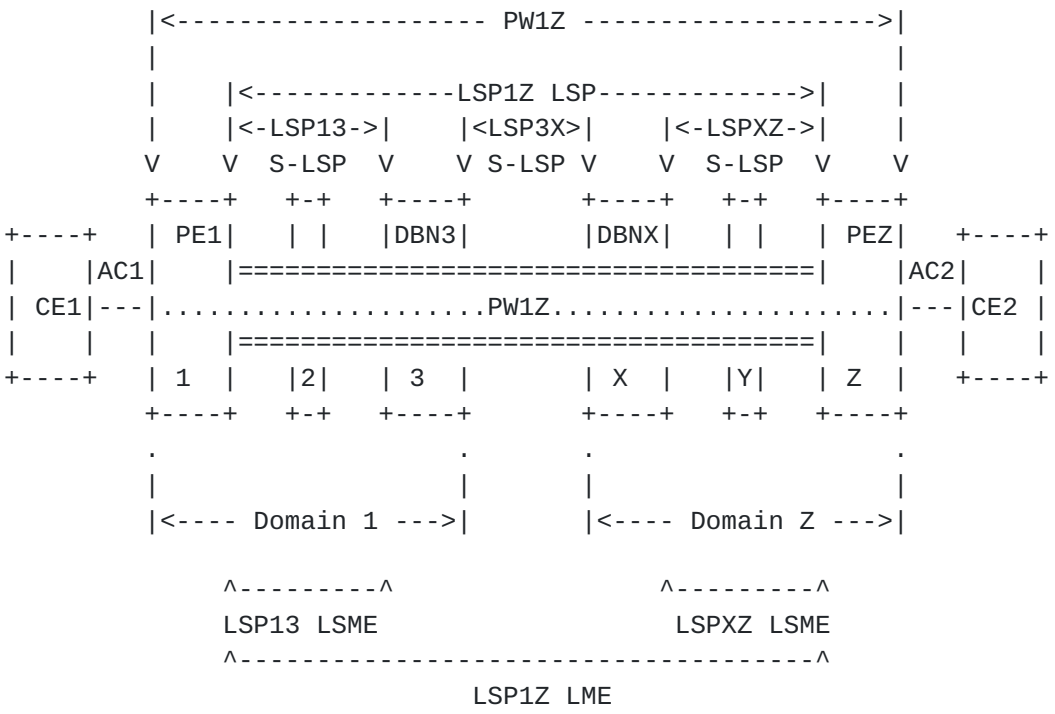

DBN: Domain Border Node

Figure 7 MPLS-TP LSP SPME ME (LSME)

Figure 7 depicts a variation of the reference model in Figure 5 where there is an end-to-end LSP (LSP1Z) between PE1 and PEZ. LSP1Z consists of, at least, three LSP Concatenated Segments: LSP13, LSP3X and LSPXZ. In this scenario there are two separate LSMEs configured to monitor the LSP1Z: 1) a LSME monitoring the LSP13 Concatenated Segment on Domain 1 (LSP13 LSME), and 2) a LSME monitoring the LSPXZ Concatenated Segment on Domain Z (LSPXZ LSME).

It is worth noticing that LSMEs can coexist with the LME monitoring the end-to-end LSP and that LSME MEPs and LME MEPs can be coincident in the same node (e.g. PE1 node supports both the LSP1Z LME MEP and the LSP13 LSME MEP).

4.5. MPLS-TP MS-PW SPME Monitoring (PSME)

An MPLS-TP MS-PW SPME Monitoring ME (PSME) is an MPLS-TP SPME with associated maintenance entity intended to monitor an arbitrary part of an MS-PW between the pair of MEPs instantiated

form the SPME independently from the end-to-end monitoring (PME). A PSME can monitor a PW segment or concatenated segment and it may also include the forwarding engine(s) of the node(s) at the edge(s) of the segment or concatenated segment. A PSME is no different than an SPME, it is simply named as such to discuss SPMEs specifically in a PW context.

When SPME is established between non-adjacent S-PEs, the edges of the SPME becomes adjacent at the MS-PW sub-layer network and any S-PEs that were previously in between becomes an LSR for the SPME.

S-PE placement is typically dictated by considerations other than OAM. S-PEs will frequently reside at operational boundaries such as the transition from distributed (CP) to centralized (NMS) control or at a routing area boundary. As such the architecture would appear not to have the flexibility that arbitrary placement of SPME segments would imply. Support for an arbitrary placement of PSME would require the definition of additional PW sub-layering.
Multiple hierarchical PSMEs can be configured on any MS-PW. PSME OAM packets fate share with the user data packets sent over the monitored PW path Segment.

A PSME can be defined between the following entities:

o T-PE and any S-PE of a given MS-PW

o Any two S-PEs of a given MS-PW.

Note that, in line with the SPME description in section 3.2, when a PW SPME is instantiated after the MS-PW has been instantiated, the TTL addressing of the MIPs may change and MIPs in the nested MEG are no longer part of the encompassing MEG. This means that the S-PE nodes hosting these MIPs are no longer S-PEs but P nodes at the SPME LSP level. The consequences are that the S-PEs hosting the PSME MEPs become adjacent S-PEs. This is no different than the operation of SPMEs in general.

A PSME is intended to be deployed in scenarios where it is preferable to monitor the behaviour of a part of a MS-PW rather than the entire end-to-end PW itself, for example to monitor an MS-PW path segment within a given network domain of an inter-domain MS-PW.

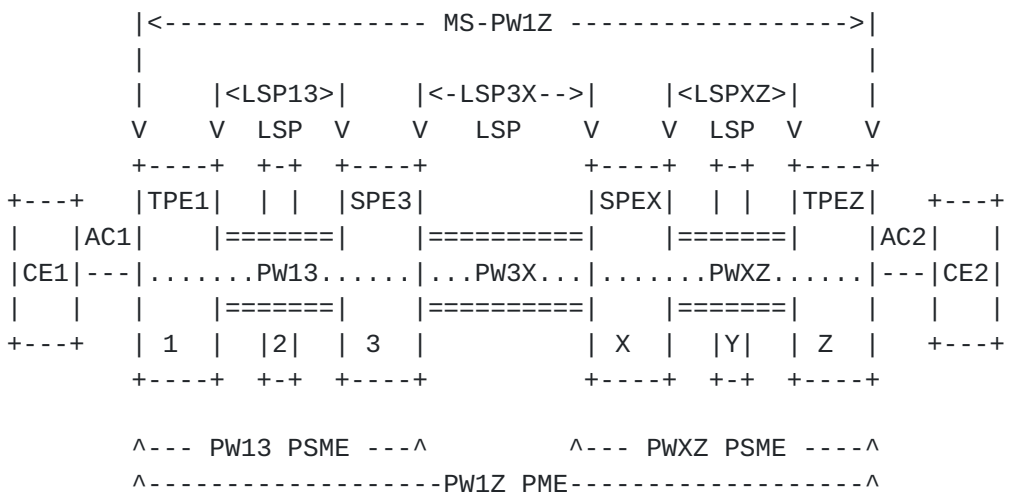

Figure 8 MPLS-TP MS-PW SPME Monitoring (PSME)

Figure 8 depicts the same MS-PW (MS-PW1Z) between AC1 and AC2 as in Figure 6. In this scenario there are two separate PSMEs configured to monitor MS-PW1Z: 1) a PSME monitoring the PW13 MS-PW path segment on Domain 1 (PW13 PSME), and 2) a PSME monitoring the PWXZ MS-PW path segment on Domain Z with (PWXZ PSME).

It is worth noticing that PSMEs can coexist with the PME monitoring the end-to-end MS-PW and that PSME MEPs and PME MEPs can be coincident in the same node (e.g. TPE1 node supports both the PW1Z PME MEP and the PW13 PSME MEP).

4.6. Fate sharing considerations for multilink

Multilink techniques are in use today and are expected to continue to be used in future deployments. These techniques include Ethernet Link Aggregations [21], the use of Link Bundling for MPLS [17] where the option to spread traffic over component links is supported and enabled. While the use of Link Bundling can be controlled at the MPLS-TP layer, use of Link Aggregation (or any server layer specific multilink) is not necessarily under control of the MPLS-TP layer. Other techniques may emerge in the future. These techniques share the characteristic that an LSP may be spread over a set of component links and therefore be reordered but no flow within the LSP is reordered (except when very infrequent and minimally disruptive load rebalancing occurs).

The use of multilink techniques may be prohibited or permitted in any particular deployment. If multilink techniques are used,

the deployment can be considered to be only partially MPLS-TP compliant, however this is unlikely to prevent its use.

The implications for OAM is that not all components of a multilink will be exercised, independent server layer OAM being required to exercise the aggregated link components. This has further implications for MIP and MEP placement, as per-interface MIPs or "down" MEPs on a multilink interface are akin to a layer violation, as they instrument at the granularity of the server layer. The implications for reduced OAM loss measurement functionality is documented in sections 5.5.3 and 6.2.3.

5. OAM Functions for proactive monitoring

In this document, 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.

Proactive monitoring is usually performed "in-service". Such transactions are universally MEP to MEP in operation while notifications emerging from the serving layer are MIP to MEP or can be MIP to MIP. The control and measurement considerations are:

1. Proactive monitoring for a MEG is typically configured at transport path creation time.

2. The operational characteristics of in-band measurement transactions (e.g., CV, LM etc.) are configured at the MEPs.

3. Server layer events are reported by transactions originating at intermediate nodes.

4. The measurements resulting from proactive monitoring are typically only reported outside of the MEG as unsolicited notifications for "out of profile" events, such as faults or loss measurement indication of excessive impairment of information transfer capability.

5. The measurements resulting from proactive monitoring may be periodically harvested by an EMS/NMS.

For statically provisioned transport paths the above information is statically configured; for dynamically established transport paths the configuration information is signaled via the control plane or configured via the management plane.

The operator enables/disables some of the consequent actions defined in section 5.1.1.4.

5.1. Continuity Check and Connectivity Verification

Proactive Continuity Check functions, as required in section 2.2.2 of RFC 5860 [11], are used to detect a loss of continuity defect (LOC) between two MEPs in a MEG.

Proactive Connectivity Verification functions, as required in section 2.2.3 of RFC 5860 [11], are used to detect an unexpected connectivity defect between two MEGs (e.g. mismerging or misconnection), as well as unexpected connectivity within the MEG with an unexpected MEP.

Both functions are based on the (proactive) generation of OAM packets by the source MEP that are processed by the sink MEP. As a consequence these two functions are grouped together into Continuity Check and Connectivity Verification (CC-V) OAM packets.

In order to perform pro-active Connectivity Verification, each CC-V OAM packet also includes a globally unique Source MEP identifier. When used to perform only pro-active Continuity Check, the CC-V OAM packet will not include any globally unique Source MEP identifier. Different formats of MEP identifiers are defined in [10] to address different environments. When MPLS-TP is deployed in transport network environments where IP addressing is not used in the forwarding plane, the ICC-based format for MEP identification is used. When MPLS-TP is deployed in an IP-based environment, the IP-based MEP identification is used.

As a consequence, it is not possible to detect misconnections between two MEGs monitored only for continuity as neither the OAM message type nor OAM message content provides sufficient information to disambiguate an invalid source. To expand:

- For CC leaking into a CC monitored MEG - undetectable
- For CV leaking into a CC monitored MEG - presence of additional Source MEP identifier allows detecting the fault
- For CC leaking into a CV monitored MEG - lack of additional Source MEP identifier allows detecting the fault.
- For CV leaking into a CV monitored MEG - different Source MEP identifier permits fault to be identified.

CC-V OAM packets are transmitted at a regular, operator's configurable, rate. The default CC-V transmission periods are application dependent (see section 5.1.3).

Proactive CC-V OAM packets are transmitted with the "minimum loss probability PHB" within the transport path (LSP, PW) they are monitoring. This PHB is configurable on network operator's basis. PHBs can be translated at the network borders by the same function that translates it for user data traffic. The implication is that CC-V fate shares with much of the forwarding implementation, but not all aspects of PHB processing are exercised. Either on-demand tools are used for finer grained fault finding or an implementation may utilize a CC-V flow per PHB with the entire E-LSP fate sharing with any individual PHB.

In a bidirectional point-to-point transport path, when a MEP is enabled to generate pro-active CC-V OAM packets with a configured transmission rate, it also expects to receive pro-active CC-V OAM packets from its peer MEP at the same transmission rate as a common SLA applies to all components of the transport path. In a unidirectional transport path (either point-to-point or point-to-multipoint), only the source MEP is enabled to generate CC-V OAM packets and only the sink MEP is configured to expect these packets at the configured rate.

MIPs, as well as intermediate nodes not supporting MPLS-TP OAM, are transparent to the pro-active CC-V information and forward these pro-active CC-V OAM packets as regular data packets.

During path setup and tear down, situations arise where CC-V checks would give rise to alarms, as the path is not fully instantiated. In order to avoid these spurious alarms the following procedures are recommended. At initialization, the MEP source function (generating pro-active CC-V packets) should be enabled prior to the corresponding MEP sink function (detecting continuity and connectivity defects). When disabling the CC-V proactive functionality, the MEP sink function should be disabled prior to the corresponding MEP source function.

5.1.1. Defects identified by CC-V

Pro-active CC-V functions allow a sink MEP to detect the defect conditions described in the following sub-sections. For all of the described defect cases, the sink MEP should notify the equipment fault management process of the detected defect.

5.1.1.1. Loss Of Continuity defect

When proactive CC-V is enabled, a sink MEP detects a loss of continuity (LOC) defect when it fails to receive pro-active CC-V OAM packets from the source MEP.

- o Entry criteria: If no pro-active CC-V OAM packets from the source MEP with the correct encapsulation (and in the case of CV, this includes the requirement to have a correct globally unique Source MEP identifier) are received within the interval equal to 3.5 times the receiving MEP's configured CC-V reception period.

- o Exit criteria: A pro-active CC-V OAM packet from the source MEP with the correct encapsulation (and again in the case of CV, with the correct globally unique Source MEP identifier) is received.

5.1.1.2. Mis-connectivity defect

When a pro-active CC-V OAM packet is received, a sink MEP identifies a mis-connectivity defect (e.g. mismerge, misconnection or unintended looping) when the received packet carries an incorrect globally unique Source MEP identifier.

- o Entry criteria: The sink MEP receives a pro-active CC-V OAM packet with an incorrect globally unique Source MEP identifier or receives a CC or CC/CV OAM packet with an unexpected encapsulation.

- o Exit criteria: The sink MEP does not receive any pro-active CC-V OAM packet with an incorrect globally unique Source MEP identifier for an interval equal at least to 3.5 times the longest transmission period of the pro-active CC-V OAM packets received with an incorrect globally unique Source MEP identifier since this defect has been raised. This requires the OAM message to self identify the CC-V periodicity as not all MEPs can be expected to have knowledge of all MEGs.

5.1.1.3. Period Misconfiguration defect

If pro-active CC-V OAM packets are received with a correct globally unique Source MEP identifier but with a transmission period different than the locally configured reception period, then a CV period mis-configuration defect is detected.

- Entry criteria: A MEP receives a CC-V pro-active packet with correct globally unique Source MEP identifier but with a Period field value different than its own CC-V configured transmission period.

- Exit criteria: The sink MEP does not receive any pro-active CC-V OAM packet with a correct globally unique Source MEP identifier and an incorrect transmission period for an interval equal at least to 3.5 times the longest transmission period of the pro-active CC-V OAM packets received with a correct globally unique Source MEP identifier and an incorrect transmission period since this defect has been raised.

5.1.1.4. Unexpected encapsulation defect

If pro-active CC-V OAM packets are received with a correct globally unique Source MEP identifier but with an unexpected encapsulation, then a CV unexpected encapsulation defect is detected.

- Entry criteria: A MEP receives a CC-V pro-active packet with correct globally unique Source MEP identifier but with an unexpected encapsulation.

  It should be noted that there are practical limitations to detecting unexpected encapsulation. It is possible that there are mis-connectivity scenarios where OAM frames can alias as payload if a transport path can carry an arbitrary payload without a pseudo wire. In this case, the mis-connectivity defect can not be detected but a LOC defect may be detected instead.

- Exit criteria: The sink MEP does not receive any pro-active CC-V OAM packet with a correct globally unique Source MEP identifier and an unexpected encapsulation for an interval equal at least to 3.5 times the longest transmission period of the pro-active CC-V OAM packets received with a correct globally unique Source MEP identifier and an unexpected encapsulation since this defect has been raised.

5.1.2. Consequent action

A sink MEP that detects one of the defect conditions defined in section 5.1.1 performs the following consequent actions.

If a MEP detects an unexpected globally unique Source MEP Identifier, it blocks all the traffic (including also the user

data packets) that it receives from the misconnected transport path.

If a MEP detects LOC defect that is not caused by a period mis-configuration, it should block all the traffic (including also the user data packets) that it receives from the transport path, if this consequent action has been enabled by the operator.

It is worth noticing that the OAM requirements document [11] recommends that CC-V proactive monitoring be enabled on every MEG in order to reliably detect connectivity defects. However, CC-V proactive monitoring can be disabled by an operator for a MEG. In the event of a misconnection between a transport path that is pro-actively monitored for CC-V and a transport path which is not, the MEP of the former transport path will detect a LOC defect representing a connectivity problem (e.g. a misconnection with a transport path where CC-V proactive monitoring is not enabled) instead of a continuity problem, with a consequent wrong traffic delivering. For these reasons, the traffic block consequent action is applied even when a LOC condition occurs. This block consequent action can be disabled through configuration. This deactivation of the block action may be used for activating or deactivating the monitoring when it is not possible to synchronize the function activation of the two peer MEPs.

If a MEP detects a LOC defect (section 5.1.1.1), a mis-connectivity defect (section 5.1.1.2) it declares a signal fail condition at the transport path level.

It is a matter if local policy if a MEP detecting a period misconfiguration defect (section 5.1.1.3) declares a signal fail condition at the transport path level.

5.1.3. Configuration considerations

At all MEPs inside a MEG, the following configuration information needs to be configured when a proactive CC-V function is enabled:

o MEG ID; the MEG identifier to which the MEP belongs;

o MEP-ID; the MEP's own identity inside the MEG;

- list of the other MEPs in the MEG. For a point-to-point MEG the list would consist of the single MEP ID from which the OAM packets are expected. In case of the root MEP of a p2mp MEG, the list is composed by all the leaf MEP IDs inside the MEG. In case of the leaf MEP of a p2mp MEG, the list is composed by the root MEP ID (i.e. each leaf needs to know the root MEP ID from which it expect to receive the CC-V OAM packets).

- PHB; it identifies the per-hop behaviour of CC-V packet. Proactive CC-V packets are transmitted with the "minimum loss probability PHB" previously configured within a single network operator. This PHB is configurable on network operator's basis. PHBs can be translated at the network borders.

- transmission rate; the default CC-V transmission periods are application dependent (depending on whether they are used to support fault management, performance monitoring, or protection switching applications):

  - Fault Management: default transmission period is 1s (i.e. transmission rate of 1 packet/second).

  - Performance Monitoring: default transmission period is 100ms (i.e. transmission rate of 10 packets/second). Performance monitoring is only relevant when the transport path is defect free. CC-V contributes to the accuracy of PM statistics by permitting the defect free periods to be properly distinguished.

  - Protection Switching: default transmission period is 3.33ms (i.e. transmission rate of 300 packets/second), in order to achieve sub-50ms the CC-V defect entry criteria should resolve in less than 10msec, and complete a protection switch within a subsequent period of 50 msec. It is also possible to lengthen the transmission period to 10ms (i.e. transmission rate of 100 packets/second): in this case the CC-V defect entry criteria is reached later (i.e. 30msec).

It should be possible for the operator to configure these transmission rates for all applications, to satisfy his internal requirements.

Note that the reception period is the same as the configured transmission rate.

For statically provisioned transport paths the above parameters are statically configured; for dynamically established transport paths the configuration information are signaled via the control plane.

The operator should be able to enable/disable some of the consequent actions. Which consequent action can be enabled/disabled are described in section 5.1.1.4.

5.2. Remote Defect Indication

The Remote Defect Indication (RDI) function, as required in section 2.2.9 of RFC 5860 [11], is an indicator that is transmitted by a sink MEP to communicate to its source MEP that a signal fail condition exists. RDI is only used for bidirectional connections and is associated with proactive CC-V. The RDI indicator is piggy-backed onto the CC-V packet.

When a MEP detects a signal fail condition (e.g. in case of a continuity or connectivity defect), it should begin transmitting an RDI indicator to its peer MEP. The RDI information will be included in all pro-active CC-V packets that it generates for the duration of the signal fail condition's existence.

A MEP that receives packets from a peer MEP (as best can be validated with the CC or CV tool in use) with the RDI information should determine that its peer MEP has encountered a defect condition associated with a signal fail.

MIPs as well as intermediate nodes not supporting MPLS-TP OAM are transparent to the RDI indicator and forward these proactive CC-V packets that include the RDI indicator as regular data packets, i.e. the MIP should not perform any actions nor examine the indicator.

When the signal fail defect condition clears, the MEP should clear the RDI indicator from subsequent transmission of pro-active CC-V packets. A MEP should clear the RDI defect upon reception of a pro-active CC-V packet from the source MEP with the RDI indicator cleared.

5.2.1. Configuration considerations

In order to support RDI indication, this may be a unique OAM message or an OAM information element embedded in a CV message. In this case the RDI transmission rate and PHB of the OAM packets carrying RDI should be the same as that configured for CC-V.

5.3. Alarm Reporting

The Alarm Reporting function, as required in section 2.2.8 of RFC 5860 [11], relies upon an Alarm Indication Signal (AIS) message to suppress alarms following detection of defect conditions at the server (sub-)layer.

When a server MEP asserts signal fail, it notifies that to the co-located MPLS-TP client/server adaptation function which then generates packets with AIS information in the downstream direction to allow the suppression of secondary alarms at the MPLS-TP MEP in the client (sub-)layer.

The generation of packets with AIS information starts immediately when the server MEP asserts signal fail. These periodic packets, with AIS information, continue to be transmitted until the signal fail condition is cleared. It is assumed that to avoid spurious alarm generation a MEP detecting loss of continuity will wait for a hold off interval prior to asserting an alarm to the management system.

Upon receiving a packet with AIS information an MPLS-TP MEP enters an AIS defect condition and suppresses loss of continuity alarms associated with its peer MEP but does not block traffic received from the transport path. A MEP resumes loss of continuity alarm generation upon detecting loss of continuity defect conditions in the absence of AIS condition.

MIPs, as well as intermediate nodes, do not process AIS information and forward these AIS OAM packets as regular data packets.

For example, let's consider a fiber cut between LSR 1 and LSR 2 in the reference network of Figure 5. Assuming that all the MEGs described in Figure 5 have pro-active CC-V enabled, a LOC defect is detected by the MEPs of Sec12 SME, LSP13 LME, PW1 PSME and PW1Z PME, however in a transport network only the alarm associated to the fiber cut needs to be reported to an NMS while all secondary alarms should be suppressed (i.e. not reported to the NMS or reported as secondary alarms).

If the fiber cut is detected by the MEP in the physical layer (in LSR2), LSR2 can generate the proper alarm in the physical layer and suppress the secondary alarm associated with the LOC defect detected on Sec12 SME. As both MEPs reside within the same node, this process does not involve any external protocol exchange. Otherwise, if the physical layer has not enough OAM

capabilities to detect the fiber cut, the MEP of Sec12 SME in LSR2 will report a LOC alarm.

In both cases, the MEP of Sec12 SME in LSR 2 notifies the adaptation function for LSP13 LME that then generates AIS packets on the LSP13 LME in order to allow its MEP in LSR3 to suppress the LOC alarm. LSR3 can also suppress the secondary alarm on PW13 PSME because the MEP of PW13 PSME resides within the same node as the MEP of LSP13 LME. The MEP of PW13 PSME in LSR3 also notifies the adaptation function for PW1Z PME that then generates AIS packets on PW1Z PME in order to allow its MEP in LSRZ to suppress the LOC alarm.

The generation of AIS packets for each MEG in the MPLS-TP client (sub-)layer is configurable (i.e. the operator can enable/disable the AIS generation).

AIS packets are transmitted with the "minimum loss probability PHB" within a single network operator. This PHB is configurable on network operator's basis.

AIS condition is cleared if no AIS message has been received in 3.5 times the AIS transmission period.

5.4. Lock Reporting

The Lock Reporting function, as required in section 2.2.7 of RFC 5860 [11], relies upon a Locked Report (LKR) message used to suppress alarms following administrative locking action in the server (sub-)layer.

When a server MEP is locked, the MPLS-TP client (sub-)layer adaptation function generates packets with LKR information in both directions to allow the suppression of secondary alarms at the MEPs in the client (sub-)layer. Again it is assumed that there is a hold off for any loss of continuity alarms in the client layer MEPs downstream of the node originating the locked report.

The generation of packets with LKR information starts immediately when the server MEP is locked. These periodic packets, with LKR information, continue to be transmitted until the locked condition is cleared.

Upon receiving a packet with LKR information an MPLS-TP MEP enters an LKR defect condition and suppresses loss of continuity alarm associated with its peer MEP but does not block traffic received from the transport path. A MEP resumes loss of

continuity alarm generation upon detecting loss of continuity defect conditions in the absence of LKR condition.

MIPs, as well as intermediate nodes, do not process the LKR information and forward these LKR OAM packets as regular data packets.

For example, let's consider the case where the MPLS-TP Section between LSR 1 and LSR 2 in the reference network of Figure 5 is administrative locked at LSR2 (in both directions).

Assuming that all the MEGs described in Figure 5 have pro-active CC-V enabled, a LOC defect is detected by the MEPs of LSP13 LME, PW1 PSME and PW1Z PME, however in a transport network all these secondary alarms should be suppressed (i.e. not reported to the NMS or reported as secondary alarms).

The MEP of Sec12 SME in LSR 2 notifies the adaptation function for LSP13 LME that then generates LKR packets on the LSP13 LME in order to allow its MEPs in LSR1 and LSR3 to suppress the LOC alarm. LSR3 can also suppress the secondary alarm on PW13 PSME because the MEP of PW13 PSME resides within the same node as the MEP of LSP13 LME. The MEP of PW13 PSME in LSR3 also notifies the adaptation function for PW1Z PME that then generates AIS packets on PW1Z PME in order to allow its MEP in LSRZ to suppress the LOC alarm.

The generation of LKR packets for each MEG in the MPLS-TP client (sub-)layer is configurable (i.e. the operator can enable/disable the LKR generation).

LKR packets are transmitted with the "minimum loss probability PHB" within a single network operator. This PHB is configurable on network operator's basis.

Locked condition is cleared if no LKR packet has been received for 3.5 times the transmission period.

5.5. Packet Loss Measurement

Packet Loss Measurement (LM) is one of the capabilities supported by the MPLS-TP Performance Monitoring (PM) function in order to facilitate reporting of QoS information for a transport path as required in section 2.2.11 of RFC 5860 [11]. LM is used to exchange counter values for the number of ingress and egress packets transmitted and received by the transport path monitored by a pair of MEPs.

Proactive LM is performed by periodically sending LM OAM packets from a MEP to a peer MEP and by receiving LM OAM packets from the peer MEP (if a bidirectional transport path) during the life time of the transport path. Each MEP performs measurements of its transmitted and received packets. These measurements are then correlated with the peer MEP in the ME to derive the impact of packet loss on a number of performance metrics for the ME in the MEG. The LM transactions are issued such that the OAM packets will experience the same queuing discipline as the measured traffic while transiting between the MEPs in the ME.

For a MEP, near-end packet loss refers to packet loss associated with incoming data packets (from the far-end MEP) while far-end packet loss refers to packet loss associated with egress data packets (towards the far-end MEP).

MIPs, as well as intermediate nodes, do not process the LM information and forward these pro-active LM OAM packets as regular data packets.

5.5.1. Configuration considerations

In order to support proactive LM, the transmission rate and PHB class associated with the LM OAM packets originating from a MEP need be configured as part of the LM provisioning. LM OAM packets should be transmitted with the PHB that yields the lowest discard probability within the measured PHB Scheduling Class (see RFC 3260 [16]).

If that PHB class is not an ordered aggregate where the ordering constraint is all packets with the PHB class being delivered in order, LM can produce inconsistent results.

5.5.2. Sampling skew

If an implementation makes use of a hardware forwarding path which operates in parallel with an OAM processing path, whether hardware or software based, the packet and byte counts may be skewed if one or more packets can be processed before the OAM processing samples counters. If OAM is implemented in software this error can be quite large.

5.5.3. Multilink issues

If multilink is used at the LSP ingress or egress, there may be no single packet processing engine where to inject or extract a LM packet as an atomic operation to which accurate packet and byte counts can be associated with the packet.

In the case where multilink is encountered in the LSP path, the reordering of packets within the LSP can cause inaccurate LM results.

5.6. Packet Delay Measurement

Packet Delay Measurement (DM) is one of the capabilities supported by the MPLS-TP PM function in order to facilitate reporting of QoS information for a transport path as required in section 2.2.12 of RFC 5860 [11]. Specifically, pro-active DM is used to measure the long-term packet delay and packet delay variation in the transport path monitored by a pair of MEPs.

Proactive DM is performed by sending periodic DM OAM packets from a MEP to a peer MEP and by receiving DM OAM packets from the peer MEP (if a bidirectional transport path) during a configurable time interval.

Pro-active DM can be operated in two ways:

- One-way: a MEP sends DM OAM packet to its peer MEP containing all the required information to facilitate one-way packet delay and/or one-way packet delay variation measurements at the peer MEP. Note that this requires synchronized precision time at either MEP by means outside the scope of this framework.

- Two-way: a MEP sends DM OAM packet with a DM request to its peer MEP, which replies with a DM OAM packet as a DM response. The request/response DM OAM packets containing all the required information to facilitate two-way packet delay and/or two-way packet delay variation measurements from the viewpoint of the source MEP.

MIPs, as well as intermediate nodes, do not process the DM information and forward these pro-active DM OAM packets as regular data packets.

5.6.1. Configuration considerations

In order to support pro-active DM, the transmission rate and PHB associated with the DM OAM packets originating from a MEP need be configured as part of the DM provisioning. DM OAM packets should be transmitted with the PHB that yields the lowest discard probability within the measured PHB Scheduling Class (see RFC 3260 [16]).

5.7. Client Failure Indication

The Client Failure Indication (CFI) function, as required in section 2.2.10 of RFC 5860 [11], is used to help process client defects and propagate a client signal defect condition from the process associated with the local attachment circuit where the defect was detected (typically the source adaptation function for the local client interface) to the process associated with the far-end attachment circuit (typically the source adaptation function for the far-end client interface) for the same transmission path in case the client of the transport path does not support a native defect/alarm indication mechanism, e.g. AIS.

A source MEP starts transmitting a CFI indication to its peer MEP when it receives a local client signal defect notification via its local CSF function. Mechanisms to detect local client signal fail defects are technology specific. Similarly mechanisms to determine when to cease originating client signal fail indication are also technology specific.

A sink MEP that has received a CFI indication report this condition to its associated client process via its local CFI function. Consequent actions toward the client attachment circuit are technology specific.

Either there needs to be a 1:1 correspondence between the client and the MEG, or when multiple clients are multiplexed over a transport path, the CFI message requires additional information to permit the client instance to be identified.

MIPs, as well as intermediate nodes, do not process the CFI information and forward these pro-active CFI OAM packets as regular data packets.

5.7.1. Configuration considerations

In order to support CFI indication, the CFI transmission rate and PHB of the CFI OAM message/information element should be configured as part of the CFI configuration.

6. OAM Functions for on-demand monitoring

In contrast to proactive monitoring, on-demand monitoring is initiated manually and for a limited amount of time, usually for operations such as e.g. diagnostics to investigate into a defect condition.

On-demand monitoring covers a combination of "in-service" and "out-of-service" monitoring functions. The control and measurement implications are:

1. A MEG can be directed to perform an "on-demand" functions at arbitrary times in the lifetime of a transport path.

2. "out-of-service" monitoring functions may require a-priori configuration of both MEPs and intermediate nodes in the MEG (e.g., data plane loopback) and the issuance of notifications into client layers of the transport path being removed from service (e.g., lock-reporting)

3. The measurements resulting from on-demand monitoring are typically harvested in real time, as these are frequently initiated manually. These do not necessarily require different harvesting mechanisms that for harvesting proactive monitoring telemetry.

The functions that are exclusive out-of-service are those described in section 6.3. The remainder are applicable to both in-service and out-of-service transport paths.

6.1. Connectivity Verification

In order to preserve network resources, e.g. bandwidth, processing time at switches, it may be preferable to not use proactive CC-V. In order to perform fault management functions, network management may invoke periodic on-demand bursts of on-demand CV packets, as required in section 2.2.3 of RFC 5860 [11].

On demand connectivity verification is a transaction that flows from the source MEP to a target MIP or MEP.

Use of on-demand CV is dependent on the existence of either a bi-directional ME, or an associated return ME, or the availability of an out-of-band return path because it requires the ability for target MIPs and MEPs to direct responses to the originating MEPs.

An additional use of on-demand CV would be to detect and locate a problem of connectivity when a problem is suspected or known based on other tools. In this case the functionality will be triggered by the network management in response to a status signal or alarm indication.

On-demand CV is based upon generation of on-demand CV packets that should uniquely identify the MEG that is being checked. The on-demand functionality may be used to check either an entire MEG (end-to-end) or between a source MEP and a specific MIP. This functionality may not be available for associated bidirectional transport paths or unidirectional paths, as the MIP may not have a return path to the source MEP for the on-demand CV transaction.

On-demand CV may generate a one-time burst of on-demand CV packets, or be used to invoke periodic, non-continuous, bursts of on-demand CV packets. The number of packets generated in each burst is configurable at the MEPs, and should take into account normal packet-loss conditions.

When invoking a periodic check of the MEG, the source MEP should issue a burst of on-demand CV packets that uniquely identifies the MEG being verified. The number of packets and their transmission rate should be pre-configured at the source MEP. The source MEP should use the mechanisms defined in sections 3.3 and 3.4 when sending an on-demand CV packet to a target MEP or target MIP respectively. The target MEP/MIP shall return a reply on-demand CV packet for each packet received. If the expected number of on-demand CV reply packets is not received at source MEP, this is an indication that a connectivity problem may exist.

On-demand CV should have the ability to carry padding such that a variety of MTU sizes can be originated to verify the MTU transport capability of the transport path.

MIPs that are not target by on-demand CV packets, as well as intermediate nodes, do not process the CV information and forward these on-demand CV OAM packets as regular data packets.

6.1.1. Configuration considerations

For on-demand CV the source MEP should support the configuration of the number of packets to be transmitted/received in each burst of transmissions and their packet size.

In addition, when the CV packet is used to check connectivity toward a target MIP, the number of hops to reach the target MIP should be configured.

The PHB of the on-demand CV packets should be configured as well. This permits the verification of correct operation of QoS queuing as well as connectivity.

6.2. Packet Loss Measurement

On-demand Packet Loss Measurement (LM) is one of the capabilities supported by the MPLS-TP Performance Monitoring function in order to facilitate diagnostic of QoS performance for a transport path, as required in section 2.2.11 of RFC 5860 [11]. As proactive LM, on-demand LM is used to exchange counter values for the number of ingress and egress packets transmitted and received by the transport path monitored by a pair of MEPs. LM is not performed MEP to MIP or between a pair of MIPs.

On-demand LM is performed by periodically sending LM OAM packets from a MEP to a peer MEP and by receiving LM OAM packets from the peer MEP (if a bidirectional transport path) during a pre-defined monitoring period. Each MEP performs measurements of its transmitted and received packets. These measurements are then correlated to evaluate the packet loss performance metrics of the transport path.

Use of packet loss measurement in an out-of-service transport path requires a traffic source such as a tester.

MIPs, as well as intermediate nodes, do not process the LM information and forward these on-demand LM OAM packets as regular data packets.

6.2.1. Configuration considerations

In order to support on-demand LM, the beginning and duration of the LM procedures, the transmission rate and PHB associated with the LM OAM packets originating from a MEP must be configured as part of the on-demand LM provisioning. LM OAM packets should be transmitted with the PHB that yields the lowest discard probability within the measured PHB Scheduling Class (see RFC 3260 [16]).

6.2.2. Sampling skew

If an implementation makes use of a hardware forwarding path which operates in parallel with an OAM processing path, whether hardware or software based, the packet and byte counts may be skewed if one or more packets can be processed before the OAM processing samples counters. If OAM is implemented in software this error can be quite large.

6.2.3. Multilink issues

Multi-link Issues are as described in section 5.5.3.

6.3. Diagnostic Tests

Diagnostic tests are tests performed on a MEG that has been taken out-of-service.

6.3.1. Throughput Estimation

Throughput estimation is an on-demand out-of-service function, as required in section 2.2.5 of RFC 5860 [11], that allows verifying the bandwidth/throughput of an MPLS-TP transport path (LSP or PW) before it is put in-service.

Throughput estimation is performed between MEPs and can be performed in one-way or two-way modes.

According to RFC 2544 [12], this test is performed by sending OAM test packets at increasing rate (up to the theoretical maximum), graphing the percentage of OAM test packets received and reporting the rate at which OAM test packets begin to drop. In general, this rate is dependent on the OAM test packet size.

When configured to perform such tests, a MEP source inserts OAM test packets with a specified packet size and transmission pattern at a rate to exercise the throughput.

For a one-way test, the remote MEP sink receives the OAM test packets and calculates the packet loss. For a two-way test, the remote MEP loopbacks the OAM test packets back to original MEP and the local MEP sink calculates the packet loss. However, a two-way test will return the minimum of available throughput in the two directions. Alternatively it is possible to run two individual one-way tests to get a distinct measurement in the two directions.

It is worth noting that two-way throughput estimation can only evaluate the minimum of available throughput of the two directions. In order to estimate the throughput of each direction uniquely, two one-way throughput estimation sessions have to be setup.

MIPs, as well as intermediate nodes, do not process the throughput test information and forward these on-demand test OAM packets as regular data packets.

6.3.1.1. Configuration considerations

Throughput estimation is an out-of-service tool. The diagnosed MEG should be put into a Lock status before the diagnostic test is started.

A MEG can be put into a Lock status either via an NMS action or using the Lock Instruct OAM tool as defined in section 7.

At the transmitting MEP, provisioning is required for a test signal generator, which is associated with the MEP. At a receiving MEP, provisioning is required for a test signal detector which is associated with the MEP.

6.3.1.2. Limited OAM processing rate

If an implementation is able to process payload at much higher data rates than OAM packets, then accurate measurement of throughput using OAM packets is not achievable. Whether OAM packets can be processed at the same rate as payload is implementation dependent.

6.3.1.3. Multilink considerations

If multilink is used, then it may not be possible to perform throughput measurement, as the throughput test may not have a mechanism for utilizing more than one component link of the aggregated link.

6.3.2. Data plane Loopback

Data plane loopback is an out-of-service function, as required in section 2.2.5 of RFC 5860 [11], that permits all traffic (including user data and OAM, with the exception of the disable loopback command) originated at the ingress of a transport path or inserted by the test equipment to be looped back unmodified (other than normal per hop processing such as TTL decrement) in the direction of the point of origin by an interface at either an intermediate node or a terminating node. TTL is decremented normally during this process. It is also normal to disable proactive monitoring of the path as the source MEP will see all source MEP originated OAM messages returned to it.

If the loopback function is to be performed at an intermediate node it is only applicable to co-routed bi-directional paths. If the loopback is to be performed end to end, it is applicable to both co-routed bi-directional or associated bi-directional paths.

Where a node implements data plane loopback capability and whether it implements more than one point is implementation dependent.

6.4. Route Tracing

It is often necessary to trace a route covered by a MEG from a source MEP to the sink MEP including all the MIPs in-between, and may be conducted after provisioning an MPLS-TP transport path for, e.g., trouble shooting purposes such as fault localization.

The route tracing function, as required in section 2.2.4 of RFC 5860 [11], is providing this functionality. Based on the fate sharing requirement of OAM flows, i.e. OAM packets receive the same forwarding treatment as data packet, route tracing is a basic means to perform connectivity verification and, to a much lesser degree, continuity check. For this function to work properly, a return path must be present.

Route tracing might be implemented in different ways and this document does not preclude any of them.

Route tracing should always discover the full list of MIPs and of the peer MEPs. In case a defect exist, the route trace function will only be able to tract up to the defect, and needs to be able to return the incomplete list of OAM entities that it was able to trace such that the fault can be localized.

6.4.1. Configuration considerations

The configuration of the route trace function must at least support the setting of the number of trace attempts before it gives up.

6.5. Packet Delay Measurement

Packet Delay Measurement (DM) is one of the capabilities supported by the MPLS-TP PM function in order to facilitate reporting of QoS information for a transport path, as required in section 2.2.12 of RFC 5860 [11]. Specifically, on-demand DM is used to measure packet delay and packet delay variation in the transport path monitored by a pair of MEPs during a pre-defined monitoring period.

On-Demand DM is performed by sending periodic DM OAM packets from a MEP to a peer MEP and by receiving DM OAM packets from

the peer MEP (if a bidirectional transport path) during a configurable time interval.

On-demand DM can be operated in two ways:

- o One-way: a MEP sends DM OAM packet to its peer MEP containing all the required information to facilitate one-way packet delay and/or one-way packet delay variation measurements at the peer MEP. Note that this requires synchronized precision time at either MEP by means outside the scope of this framework.

- o Two-way: a MEP sends DM OAM packet with a DM request to its peer MEP, which replies with an DM OAM packet as a DM response. The request/response DM OAM packets containing all the required information to facilitate two-way packet delay and/or two-way packet delay variation measurements from the viewpoint of the source MEP.

MIPs, as well as intermediate nodes, do not process the DM information and forward these on-demand DM OAM packets as regular data packets.

6.5.1. Configuration considerations

In order to support on-demand DM, the beginning and duration of the DM procedures, the transmission rate and PHB associated with the DM OAM packets originating from a MEP need be configured as part of the DM provisioning. DM OAM packets should be transmitted with the PHB that yields the lowest discard probability within the measured PHB Scheduling Class (see RFC 3260 [16]).

In order to verify different performances between long and short packets (e.g., due to the processing time), it should be possible for the operator to configure the packet size of the on-demand OAM DM packet.

7. OAM Functions for administration control

7.1. Lock Instruct

Lock Instruct (LKI) function, as required in section 2.2.6 of RFC 5860 [11], is a command allowing a MEP to instruct the peer MEP(s) to put the MPLS-TP transport path into a locked condition.

This function allows single-side provisioning for administratively locking (and unlocking) an MPLS-TP transport path.

Note that it is also possible to administratively lock (and unlock) an MPLS-TP transport path using two-side provisioning, where the NMS administratively put both MEPs into ad administrative lock condition. In this case, the LKI function is not required/used.

MIPs, as well as intermediate nodes, do not process the lock instruct information and forward these on-demand LKI OAM packets as regular data packets.

7.1.1. Locking a transport path

A MEP, upon receiving a single-side administrative lock command from an NMS, sends an LKI request OAM packet to its peer MEP(s). It also puts the MPLS-TP transport path into a locked state and notifies its client (sub-)layer adaptation function upon the locked condition.

A MEP, upon receiving an LKI request from its peer MEP, can accept or not the instruction and replies to the peer MEP with an LKI reply OAM packet indicating whether it has accepted or not the instruction. This requires either an in-band or out-of-band return path.

If the lock instruction has been accepted, it also puts the MPLS-TP transport path into a locked and notifies its client (sub-)layer adaptation function upon the locked condition.

Note that if the client (sub-)layer is also MPLS-TP, Lock Reporting (LKR) generation at the client MPLS-TP (sub-)layer is started, as described in section 5.4.

7.1.2. Unlocking a transport path

A MEP, upon receiving a single-side administrative unlock command from NMS, sends an LKI removal request OAM packet to its peer MEP(s).

The peer MEP, upon receiving an LKI removal request, can accept or not the removal instruction and replies with an LKI removal reply OAM packet indicating whether it has accepted or not the instruction.

If the lock removal instruction has been accepted, it also clears the locked condition on the MPLS-TP transport path and notifies this event to its client (sub-)layer adaptation function.

The MEP that has initiated the LKI clear procedure, upon receiving a positive LKI removal reply, also clears the locked condition on the MPLS-TP transport path and notifies this event to its client (sub-)layer adaptation function.

Note that if the client (sub-)layer is also MPLS-TP, Lock Reporting (LKR) generation at the client MPLS-TP (sub-)layer is terminated, as described in section 5.4.

8. Security Considerations

A number of security considerations are important in the context of OAM applications.

OAM traffic can reveal sensitive information such as passwords, performance data and details about e.g. the network topology. The nature of OAM data therefore suggests to have some form of authentication, authorization and encryption in place. This will prevent unauthorized access to vital equipment and it will prevent third parties from learning about sensitive information about the transport network. However it should be observed that the combination of all permutations of unique MEP to MEP, MEP to MIP, and intermediate system originated transactions mitigates against the practical establishment and maintenance of a large number of security associations per MEG.

For this reason it is assumed that the network is physically secured against man-in-the-middle attacks. Further, this document describes OAM functions that, if a man-in-the-middle attack was possible, could be exploited to significantly disrupt proper operation of the network.

Mechanisms that the framework does not specify might be subject to additional security considerations.

9. IANA Considerations

No new IANA considerations.

10. Acknowledgments

The authors would like to thank all members of the teams (the Joint Working Team, the MPLS Interoperability Design Team in

IETF and the Ad Hoc Group on MPLS-TP in ITU-T) involved in the definition and specification of MPLS Transport Profile.

The editors gratefully acknowledge the contributions of Adrian Farrel, Yoshinori Koike, Luca Martini, Yuji Tochio and Manuel Paul for the definition of per-interface MIPs and MEPs.

The editors gratefully acknowledge the contributions of Malcolm Betts, Yoshinori Koike, Xiao Min, and Maarten Vissers for the lock report and lock instruction description.

The authors would also like to thank Alessandro D'Alessandro, Loa Andersson, Malcolm Betts, Stewart Bryant, Rui Costa, Xuehui Dai, John Drake, Adrian Farrel, Dan Frost, Xia Liang, Liu Gouman, Peng He, Feng Huang, Su Hui, Yoshionori Koike, George Swallow, Yuji Tochio, Curtis Villamizar, Maarten Vissers and Xuequin Wei for their comments and enhancements to the text.

This document was prepared using 2-Word-v2.0.template.dot.

11. References

11.1. Normative References

[1] Rosen, E., Viswanathan, A., Callon, R., "Multiprotocol Label Switching Architecture", RFC 3031, January 2001

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

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

[4] Bocci, M., Bryant, S., "An Architecture for Multi-Segment Pseudo Wire Emulation Edge-to-Edge", RFC 5659, October 2009

[5] Niven-Jenkins, B., Brungard, D., Betts, M., sprecher, N., Ueno, S., "MPLS-TP Requirements", RFC 5654, September 2009

[6] Agarwal, P., Akyol, B., "Time To Live (TTL) Processing in Multiprotocol Label Switching (MPLS) Networks", RFC 3443, January 2003

[7] Vigoureux, M., Bocci, M., Swallow, G., Ward, D., Aggarwal, R., "MPLS Generic Associated Channel", RFC 5586, June 2009

[8] Bocci, M., et al., "A Framework for MPLS in Transport Networks", RFC 5921, July 2010

[9] Bocci, M., et al., " MPLS Transport Profile User-to-Network and Network-to-Network Interfaces", draft-ietf-mpls-tp-uni-nni-00 (work in progress), August 2010

[10] Swallow, G., Bocci, M., "MPLS-TP Identifiers", draft-ietf-mpls-tp-identifiers-02 (work in progress), July 2010

[11] Vigoureux, M., Betts, M., Ward, D., "Requirements for OAM in MPLS Transport Networks", RFC 5860, May 2010

[12] Bradner, S., McQuaid, J., "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999

[13] ITU-T Recommendation G.806 (01/09), "Characteristics of transport equipment - Description methodology and generic functionality ", January 2009

11.2. Informative References

[14] Sprecher, N., Nadeau, T., van Helvoort, H., Weingarten, Y., "MPLS-TP OAM Analysis", draft-ietf-mpls-tp-oam-analysis-02 (work in progress), July 2010

[15] Nichols, K., Blake, S., Baker, F., Black, D., "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998

[16] Grossman, D., "New terminology and clarifications for Diffserv", RFC 3260, April 2002.

[17] Kompella, K., Rekhter, Y., Berger, L., "Link Bundling in MPLS Traffic Engineering (TE)", RFC 4201, October 2005

[18] ITU-T Recommendation G.707/Y.1322 (01/07), "Network node interface for the synchronous digital hierarchy (SDH)", January 2007

[19] ITU-T Recommendation G.805 (03/00), "Generic functional architecture of transport networks", March 2000

[20] ITU-T Recommendation Y.1731 (02/08), "OAM functions and mechanisms for Ethernet based networks", February 2008

[21] IEEE Standard 802.1AX-2008, "IEEE Standard for Local and Metropolitan Area Networks - Link Aggregation", November 2008

Authors' Addresses

Dave Allan
Ericsson

Email: david.i.allan@ericsson.com

Italo Busi
Alcatel-Lucent

Email: Italo.Busi@alcatel-lucent.com

Ben Niven-Jenkins
Velocix

Email: ben@niven-jenkins.co.uk

Annamaria Fulignoli
Ericsson

Email: annamaria.fulignoli@ericsson.com

Enrique Hernandez-Valencia
Alcatel-Lucent

Email: Enrique.Hernandez@alcatel-lucent.com

Lieven Levrau
Alcatel-Lucent

Email: Lieven.Levrau@alcatel-lucent.com

Vincenzo Sestito
Alcatel-Lucent

Email: Vincenzo.Sestito@alcatel-lucent.com

Nurit Sprecher
Nokia Siemens Networks

Email: nurit.sprecher@nsn.com

Huub van Helvoort
Huawei Technologies

Email: hhelvoort@huawei.com

Martin Vigoureux
Alcatel-Lucent

Email: Martin.Vigoureux@alcatel-lucent.com

Yaacov Weingarten
Nokia Siemens Networks

Email: yaacov.weingarten@nsn.com

Rolf Winter
NEC

Email: Rolf.Winter@nw.neclab.eu