MPLS Working Group
Internet Draft
Intended status: Informational

I. Busi (Ed)
Alcatel-Lucent

B. Niven-Jenkins (Ed)
BT

Expires: January 2010

July 13, 2009

MPLS-TP OAM Framework
draft-ietf-mpls-tp-oam-framework-01.txt

Status of this Memo

Copyright Notice

Abstract

Multi-Protocol Label Switching (MPLS) Transport Profile (MPLS-TP) is based on a profile of the MPLS and pseudowire (PW) procedures as specified in the MPLS Traffic Engineering (MPLS-TE), pseudowire (PW) and multi-segment PW (MS-PW) architectures complemented with additional Operations, Administration and Maintenance (OAM) procedures for fault, performance and protection-switching management for packet transport applications that do not rely on the presence of a control plane.

This document provides a framework that supports a comprehensive set of OAM procedures that fulfills the MPLS-TP OAM requirements [11].

Table of Contents

1. Introduction

As noted in [8], MPLS-TP defines a profile of the MPLS-TE and (MS-)PW architectures defined in RFC 3031 [2], RFC 3985 [5] and [7] complemented with additional OAM procedures for fault, performance and protection-switching management for packet transport applications.

[Editor's note - The draft needs to be reviewed to ensure support of OAM for p2mp transport paths]

In line with [12], 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.

The MPLS-TP OAM framework defined in this document provides a comprehensive set of OAM procedures that satisfy the MPLS-TP OAM requirements [11]. In this regard, it is similar to existing SONET/SDH and OTH OAM mechanisms (e.g. [16]).

1.1. Contributing Authors

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

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [1].

2.1. Terminology

DBN Domain Border Node

LME LSP Maintenance Entity

LTCME LSP Tandem Connection Maintenance Entity

[Editor's note - Difference or similarity between tandem connection monitoring (TCM)_and Path Segment Tunnel (PST) need to be defined and agreed]

ME Maintenance Entity

[Editor's note - There is a need to define whether to support OAM on p2mp transport path there is a need to introduce the MEG concept]

MEP Maintenance End Point

MIP Maintenance Intermediate Point

PHB Per-hop Behavior

PME PW Maintenance Entity

PTCME PW Tandem Connection Maintenance Entity

SME Section Maintenance Entity

2.2. Definitions

ME: The collection of MEPs and MIPs and their association (details in section 3.1).

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

MEP Source: A MEP acts as MEP source for the OAM messages that it originates and inserts into its associated ME.

MEP Sink: A MEP acts as a MEP sink for the OAM messages that it terminates and processes from its associated ME.

MIP: A MIP terminates and processes OAM messages and generates OAM messages in reaction to received OAM messages. A MIP resides within a ME between MEPs (details in section 3.4).

OAM domain: A domain, as defined in [10], 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: An OAM flow is a flow of OAM packets between a pair of MEPs or a MEP and a MIP that is used to monitor and maintain a ME [Editor's note - a MEG depending on what we decide for this point]. An OAM flow is associated to a unique ME and contains the OAM monitoring, signalling and notification messages necessary to monitor and maintain that ME. The exact mix of message types in an OAM flow will be dependent on the technology being monitored and the exact deployment scenario of that technology (e.g. some deployments may proactively monitor the connectivity of all transport paths whereas other deployments may only reactively monitor transport paths).

OAM Message: An OAM information element that performs some OAM functionality (e.g. continuity check and connectivity verification)

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

Path: See Transport Path

Signal Fail: A condition when the data associated with a transport path has failed in the sense that a defect condition (not being a degraded defect) is detected.

Tandem Connection: A tandem connection is an arbitrary part of a transport path that can be monitored (via OAM) independently from the end-to-end monitoring (OAM). It may be a monitored segment or a monitored concatenated segment of a transport path. The tandem connection may also include the forwarding engine(s) of the node(s) at the edge(s) of the segment or concatenated segment.

The following terms are defined in [10] as follows:

Associated bidirectional path: A path that supports traffic flow in both directions but which is constructed from a pair of unidirectional paths (one for each direction) which are associated with one another at the path's ingress/egress points. The forward and backward directions are setup, monitored and protected independently. As a consequence they may or may not follow the same route (links and nodes) across the network.

Concatenated Segment: A serial-compound link connection as defined in G.805 [17]. A concatenated segment is a contiguous part of an LSP or multi-segment PW that comprises a set of segments and their interconnecting nodes in sequence. See also "Segment".

Co-routed bidirectional path: A path where the forward and backward directions follow the same route (links and nodes) across the network. Both directions are setup, monitored and protected as a single entity. A transport network path is typically co-routed.

Layer network: Layer network is defined in G.805 [17]. A layer network provides for the transfer of client information and independent operation of the client OAM. A Layer Network may be described in a service context as follows: one layer network may provide a (transport) service to higher client layer network and may, in turn, be a client to a lower layer network. A layer network is a logical construction somewhat independent of arrangement or composition of physical network elements. A particular physical network element may topologically belong to more than one layer network, depending on the actions it takes on the encapsulation associated with the logical layers (e.g. the label stack), and thus could be modeled as multiple logical elements. A layer network may consist of one or more sublayers. Section 1.3 provides a more detailed overview of what constitutes a layer network. For additional explanation of how layer networks relate to the OSI concept of layering see Appendix I of Y.2611 [19].

Section Layer Network: A section layer is a server layer (which may be MPLS-TP or a different technology) which provides for the transfer of the section layer client information between adjacent nodes in the

transport path layer or transport service layer. A section layer may provide for aggregation of multiple MPLS-TP clients. Note that G.805 [17] defines the section layer as one of the two layer networks in a transmission media layer network. The other layer network is the physical media layer network.

Path: See Transport Path

Segment: A link connection as defined in G.805 [17]. A segment is the part of an LSP that traverses a single link or the part of a PW that traverses a single link (i.e. that connects a pair of adjacent {Switching|Terminating} Provider Edges). See also "Concatenated Segment".

Sublayer: Sublayer is defined in G.805 [17]. The distinction between a layer network and a sublayer is that a sublayer is not directly accessible to clients outside of its encapsulating layer network and offers no direct transport service for a higher layer (client) network

Transport Path Layer: A (sub-)layer network that provides point-to-point or point-to-multipoint transport paths. It provides independent (of the client) OAM when transporting its clients.

Unidirectional path: A path that supports traffic flow in only one direction.

The term 'Domain Border Node' is defined in [14] as follows:

Domain Border Node: To be defined

[Editor's note - There is no definition of Domain Border Node in RFC 5151]

The term 'Per-hop Behavior' is defined in [13] as follows:

Per-hop Behavior: a description of the externally observable forwarding treatment applied at a differentiated services-compliant node to a behavior aggregate.

3. Functional Components

MPLS-TP defines a profile of the MPLS and PW architectures ([2], [5] and [7]) that is designed to transport service traffic complying with certain performance and quality requirements. In order to verify and maintain these performance and quality requirements, 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 path.

MPLS-TP OAM operates in the context of Maintenance Entities (MEs).

A Maintenance Entity is the collection of two (or more) Maintenance End Points (MEPs) and their association. The MEPs that form an ME are configured and managed to limit the scope of an OAM flow within the ME the MEPs belong to (i.e. within the domain of the transport path or segment, in the specific layer network, that is being monitored and managed).

An example of an ME with more than two MEPs is a point-to-multipoint ME monitoring a point-to-multipoint transport path (or point-to-multipoint tandem connection).

Each MEP resides at the boundaries of the ME that they are part of. An ME may also include a set of zero or more Maintenance Intermediate Points (MIPs), which reside within the Maintenance Entity, between the MEPs.

MEPs and MIPs are associated with only one Maintenance Entity.

The abstract reference model for a ME with MEPs and MIPs is described in Figure 1 below:

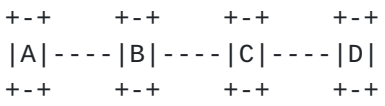

Figure 1 ME Abstract Reference Model

The instantiation of this abstract model to different MPLS-TP entities is described in section 4. In this model, nodes A, B, C and D can be LER/LSR for an LSP or the {S|T}-PEs for a MS-PW. Nodes A and D are MEPs while B and C are MIPs. The links connecting adjacent nodes can be either physical links or lower-level LSPs.

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

When a control plane is not present, the management plane configures MEPs and MIPs. Otherwise they can be configured either by the management plane or by the control plane.

3.1. Maintenance Entity

A Maintenance Entity may be defined to monitor for fault and performance management unidirectional point-to-point or point-to-multipoint transport paths or tandem connections, or co-routed bidirectional point-to-point transport paths and tandem connections in an MPLS-TP layer network.

In case of associated bi-directional paths, two independent Maintenance Entities are define to independently monitor each direction.

An MPLS-TP maintenance entity can be either the whole end-to-end transport path or a Tandem Connection of the transport path.

The following properties apply to all MPLS-TP MEs:

- They can be nested but not overlapped, e.g. a ME may cover a segment or a concatenated segment of another ME, and may also include the forwarding engine(s) of the node(s) at the edge(s) of the segment or concatenated segment, but all its MEPs and MIPs are no longer part of the encompassing ME. It is possible that MEPs of nested MEs reside on a single node.

- Each OAM flow is associated with a single Maintenance Entity.

- OAM packets are subject to the same forwarding treatment (e.g. fate share) as the data traffic, but they can be distinguished from the data traffic using the GAL and ACH constructs [9] for LSP and the ACH construct [6]and [9] for (MS-)PW.

3.2. Point-to-multipoint scenario

The reference model for the p2mp scenario is represented in Figure 4.

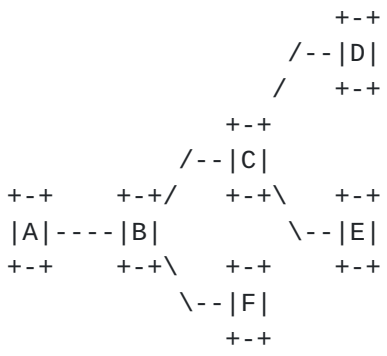

Figure 2 Reference Model for p2mp

In case of p2mp transport paths, the OAM measurements are different for each [Root, Leaf] relationship (A-D, A-E and A-F):

o Fault conditions - depending from where the failure is located

o Packet loss - depending from where the packets are lost

o Packet delay - depending on different paths

Each leaf (i.e. D, E and F) terminates OAM messages to monitor its own [Root, Leaf] relationship while the root (i.e. A) generates OAM messages to monitor all the [Root, Leaf] relationships of p2mp transport path.

[Editor's note - Further considerations regarding the p2mp case will be added in a future version of this document]

3.3. Maintenance End Points (MEPs)

Maintenance End Points (MEPs) are the end points of a ME.

In the context of an MPLS-TP LSP, only LERs can be MEPs while in the context of an LSP Tandem Connection both LERs and LSRs can be MEPs.

In the context of MPLS-TP PW, only T-PEs can be MEPs while in the context of a PW Tandem Connection both T-PEs and S-PEs can be MIPs.

MEPs are responsible for activating and controlling all of the OAM functionality for the ME. A MEP is capable of initiating and terminating OAM messages for fault management and performance monitoring.

MEPs prevent OAM packets corresponding to a ME from leaking outside that ME:

- o A MEP sink terminates all the OAM packets that it receives corresponding to its ME and does not forward them further along the path. If the pro-active CC-V OAM tool detects an unintended connectivity, all traffic on the path is blocked (i.e. all received packets are dropped, including user-data packets).

[Editor's note - Need to discuss whether to keep the last sentence in the bullet above regarding the MEP behavior in case of misconnections]

- o A MEP source tunnels all the OAM packets that it receives, upstream from the associated ME, via label stacking. These packets are not processed within the ME as they belong to another ME.

[Editor's - Need to rephrase the bullet above to clarify what it actually means]

MPLS-TP MEP passes a fault indication to its client (sub-)layer network.

A MEP of a tandem connection is not necessarily coincident with the termination of the MPLS-TP transport path (LSP or PW). Within the boundary of the tandem connection it can monitor the MPLS-TP transport path for failures or performance degradation (e.g. count packets).

[Editor's note - The MEP of a TCM monitors the transport paths' connectivity within the scope of the TCM. This means that failures or performance degradations within the TCM are detected by the TCM MEP while failures or performance degradations outside the TCM are not detected by the TCM MEP.

Improve the text above to explain this concept.]

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.

[Editor's note - Add some text regarding MEP identification as well as about what a MEP should do if it receives an unexpected OAM packet]

3.4. Maintenance Intermediate Points (MIPs)

A Maintenance Intermediate Point (MIP) is a point between the two MEPs in an ME.

In the context of an MPLS-TP LSP and LSP Tandem Connections, LSRs can be MIPs.

In the context of MPLS-TP PW and PW Tandem Connection, S-PEs can be MIPs.

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.

A MIP does not initiate unsolicited OAM packets, but may be addressed by OAM packets initiated by one of the MEPs of the ME. A MIP can generate OAM packets only in response to OAM packets that are sent on the ME it belongs to.

[Editor's note - It is needed to provide an high-level description about how this is achieved (e.g. TTL expiry).]

MIPs are unaware of any OAM flows running between MEPs or between MEPs and other MIPs. MIPs can only receive and process OAM packets addressed to the MIP itself.

A MIP takes no action on the MPLS-TP transport path.

[Editor's note - Add some text regarding MIP identification as well as about what a MIP should do if it receives an unexpected OAM packet]

3.5. Server MEPs

A server MEP is a MEP of an ME that is either:

- defined in a layer network below the MPLS-TP layer network being referenced, or

- defined in a sub-layer of the MPLS-TP layer network that is below the sub-layer being referenced.

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

A server MEP provides also client/server adaptation function between the client (MPLS-TP) layer network and the server layer network. As a consequence the server MEP is aware of the MPLS-TP transport paths that are setup over that server layer's transport path.

For example, a server MEP can be either:

- o A termination point of a physical link (e.g. 802.3), an SDH VC or OTH ODU for the MPLS-TP Section layer network, defined in section 4.1;
- o An MPLS-TP Section MEP for MPLS-TP LSPs, defined in section 4.2;
- o An MPLS-TP LSP MEP for MPLS-TP PWs, defined in section 4.4;
- o An MPLS-TP LSP Tandem Connection MEP for higher-level LTCMEs, defined in section 4.3;
- o An MPLS-TP PW Tandem Connection MEP for higher-level PTCMEs, defined in section 4.5.

The server MEP can run appropriate OAM functions for fault detection within the server (sub-)layer network, and notifies a fault indication to its client MPLS-TP layer network. Server MEP OAM functions are outside the scope of this document.

4. Reference Model

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

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

- o A Section Maintenance Entity (SME), allowing monitoring and management of MPLS-TP Sections (between MPLS LSRs).
- o A LSP Maintenance Entity (LME), allowing monitoring and management of an end-to-end LSP (between LERs).
- o A PW Maintenance Entity (PME), allowing monitoring and management of an end-to-end SS/MS-PWs (between T-PEs).
- o An LSP Tandem Connection Maintenance Entity (LTCME), allowing monitoring and management of an LSP Tandem Connection between any LER/LSR along the LSP.

o A MS-PW Tandem Connection Maintenance Entity (PTCME), allows monitoring and management of a SS/MS-PW Tandem Connection between any T-PE/S-PE along the (MS-)PW.

The MEs specified in this MPLS-TP framework are compliant with the architecture framework for MPLS MS-PWs [7] and MPLS LSPs [2].

Hierarchical LSPs are also supported. In this case, each LSP Tunnel in the hierarchy is a different sub-layer network that can be monitored independently from higher and lower level LSP tunnels in the hierarchy, end-to-end (from LER to LER) by an LME. Tandem Connection monitoring via LTCME are applicable on each LSP Tunnel in the hierarchy.

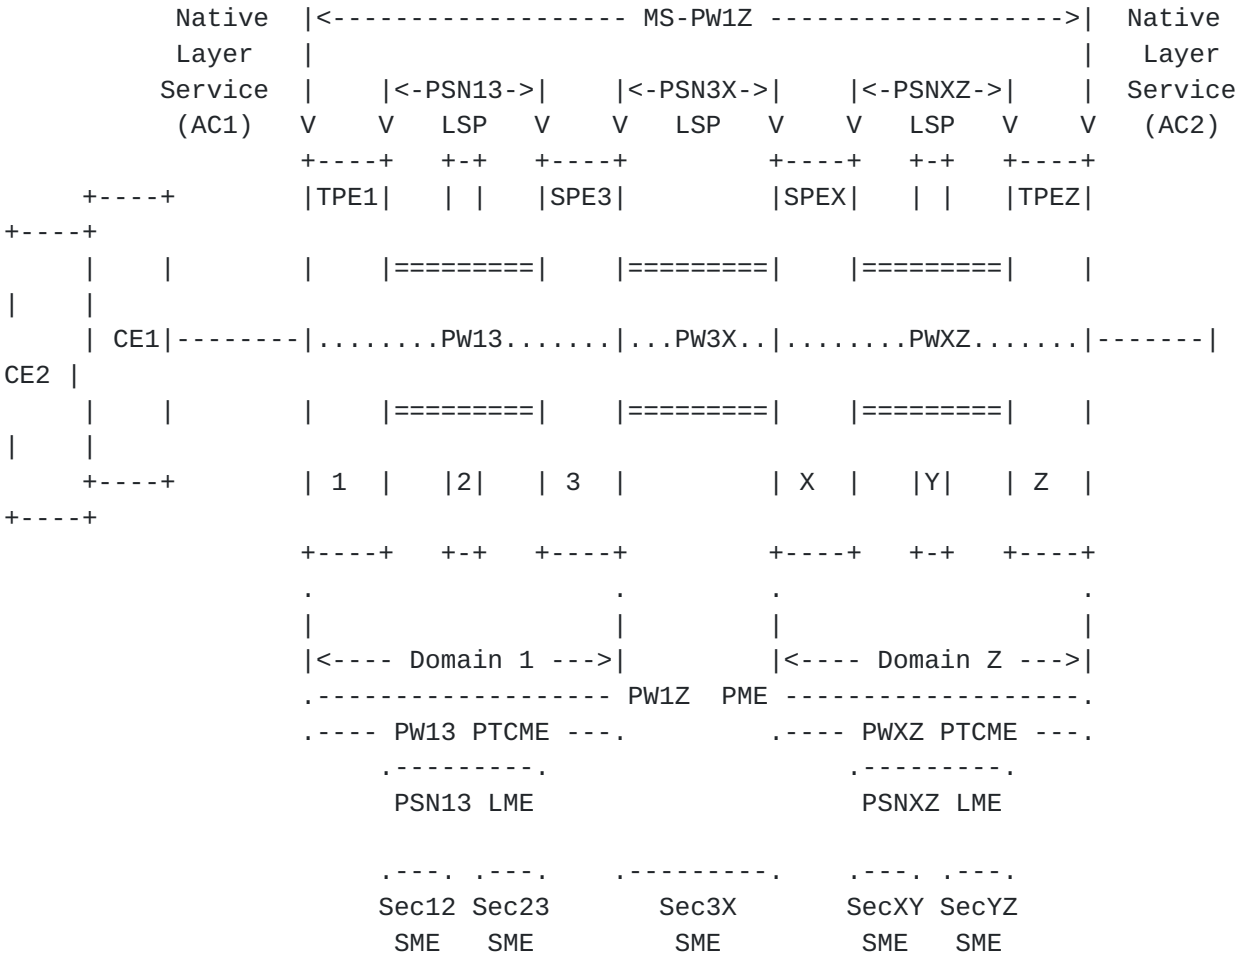

TPE1: Terminating Provider Edge 1 SPE2: Switching Provider Edge 3

TPEX: Terminating Provider Edge X SPEZ: Switching Provider Edge Z

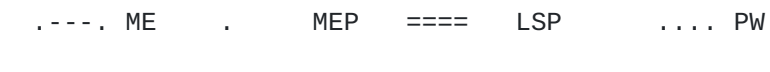

Figure 3 Reference Model for the MPLS-TP OAM Framework

Figure 3 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 Section Sec12 and LSR2 is adjacent to LSR3 via the MPLS Section Sec23. Similarly, in Domain Z, LSRX is adjacent to LSRY via the MPLS Section SecXY and LSRY is adjacent to LSRZ via the MPLS Section SecYZ. In addition, LSR3 is adjacent to LSRX via the MPLS Section 3X.

Figure 3 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 Segments: 1) PW13 segment between T-PE1 and S-PE3 via the bi-directional PSN13 LSP, 2) PW3X segment between S-PE3 and S-PEX, and 3) PWXZ segment between S-PEX and T-PEZ via the bi-directional PSNXZ LSP.

The MPLS-TP OAM procedures that apply to an instance of a given ME are expected to operate independently from procedures on other instances of the same ME and certainly of other MEs. Yet, this does not preclude that multiple MEs may be affected simultaneously by the same network condition, for example, a fibre cut event.

Note that there are no constrains imposed by this OAM framework on the number, or type, of MEs that may be instantiated on a particular node. In particular, when looking at Figure 1, it should be possible to configure one or more MEPs on the same node if that node is the endpoint of one or more MEs.

Figure 3 does not describe a PW3X PTCME because typically TCMs are used to monitor an OAM domain (like PW13 and PWXZ PTCMEs) rather than the segment between two OAM domains. However the OAM framework does not pose any constraints on the way TCM are instantiated as long as they are not overlapping.

The subsections below define the MEs specified in this MPLS-TP OAM architecture framework document. Unless otherwise stated, all references to domains, LSRs, MPLS Sections, LSP, pseudowires and MEs in this Section are made in relation to those shown in Figure 3.

4.1. MPLS-TP Section Monitoring

An MPLS-TP Section ME (SME) is an MPLS-TP maintenance entity intended to monitor the forwarding behaviour of an MPLS Section as defined in [10]. An SME may be configured on any MPLS section. SME OAM packets fate share with the user data packets sent over the monitored MPLS Section.

[Editor's note - Is OAM monitoring only the forwarding behaviour? If not, we need to clarify what it is monitoring]

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 (and MPLS-TP enabled) LSRs rather than monitoring the individual LSP or PW segments traversing the MPLS Section and the server layer technology does not provide adequate OAM capabilities.

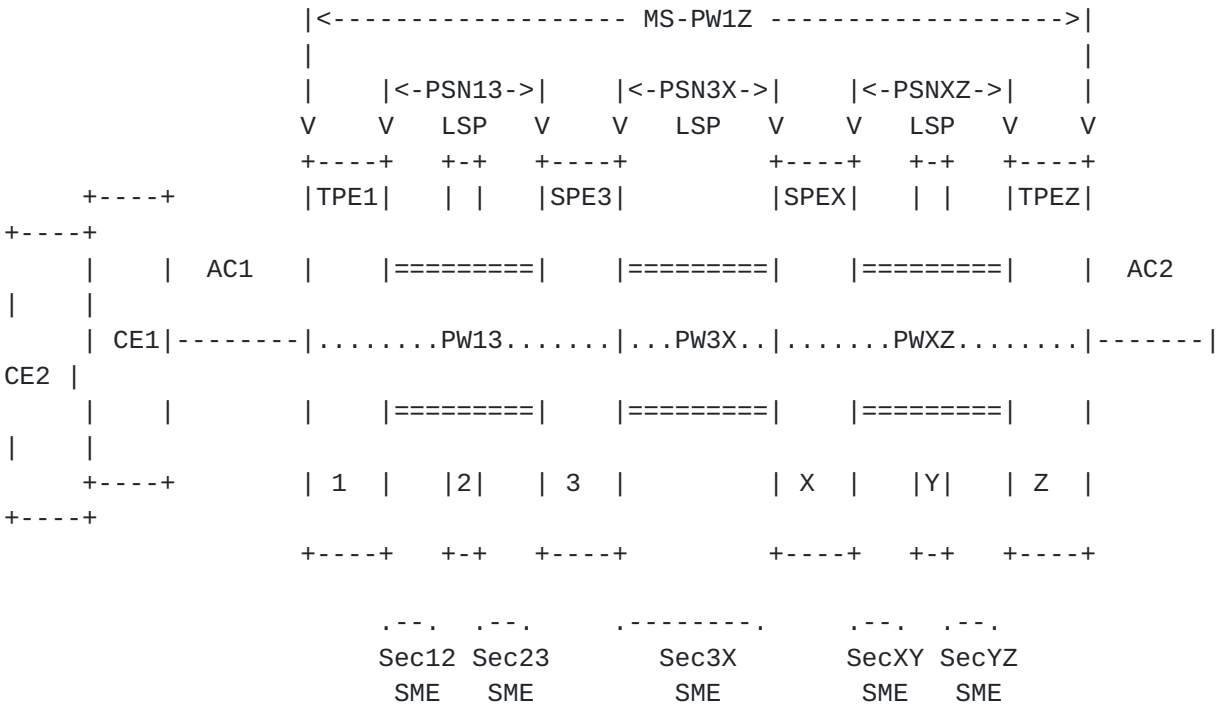

Figure 4 Reference Example of MPLS-TP Section MEs (SME)

Figure 4 shows 5 Section MEs configured in the path between AC1 and AC2: 1) Sec12 ME associated with the MPLS Section between LSR 1 and LSR 2, 2) Sec23 ME associated with the MPLS Section between LSR 2 and LSR 3, 3) Sec3X ME associated with the MPLS Section between LSR 3 and LSR X, 4) SecXY ME associated with the MPLS Section between LSR X and LSR Y, and 5) SecYZ ME associated with the MPLS Section between LSR Y and LSR Z.

4.2. MPLS-TP LSP End-to-End Monitoring

An MPLS-TP LSP ME (LME) is an MPLS-TP maintenance entity intended to monitor the forwarding behaviour of an end-to-end LSP between two (e.g., a point-to-point LSP) or more (e.g., a point-to-multipoint LSP) LERs. An LME may be configured on any MPLS LSP. LME OAM packets 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 the forwarding behaviour of an entire LSP between its LERs, rather than, say, monitoring individual PWs.

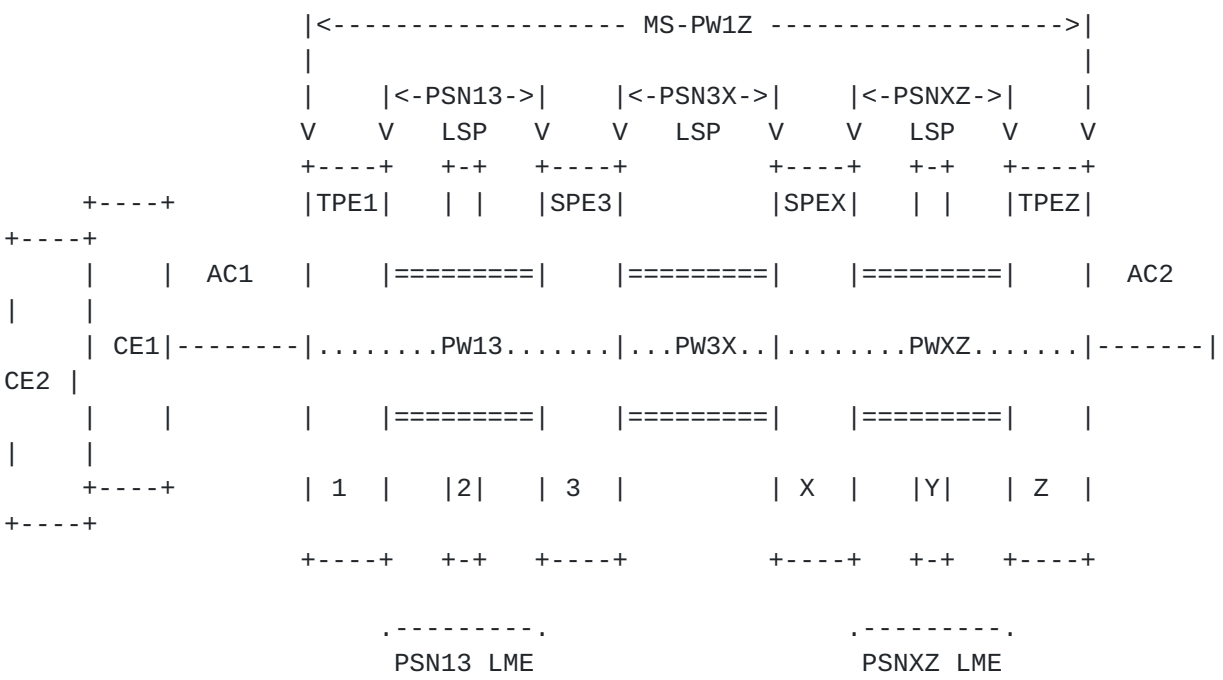

Figure 5 Examples of MPLS-TP LSP MEs (LME)

Figure 5 depicts 2 LMEs configured in the path between AC1 and AC2: 1) the PSN13 LME between LER 1 and LER 3, and 2) the PSNXZ LME between LER X and LER Y. Note that the presence of a PSN3X 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 LSP Tandem Connection Monitoring

An MPLS-TP LSP Tandem Connection Monitoring ME (LTCME) is an MPLS-TP maintenance entity intended to monitor the forwarding behaviour of an arbitrary part of an LSP between a given pair of LSRs independently from the end-to-end monitoring (LME). An LTCME 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.

Multiple LTCMEs MAY be configured on any LSP. The LSRs that terminate the LTCME may or may not be immediately adjacent at the MPLS-TP layer. LTCME OAM packets fate share with the user data packets sent over the monitored LSP segment.

A LTCME can be defined between the following entities:

o LER and any LSR of a given LSP.

o Any two LSRs of a given LSP.

An LTCME is intended to be deployed in scenarios where it is preferable to monitor the behaviour of a part of an LSP 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.

Note that LTCMEs are equally applicable to hierarchical LSPs.

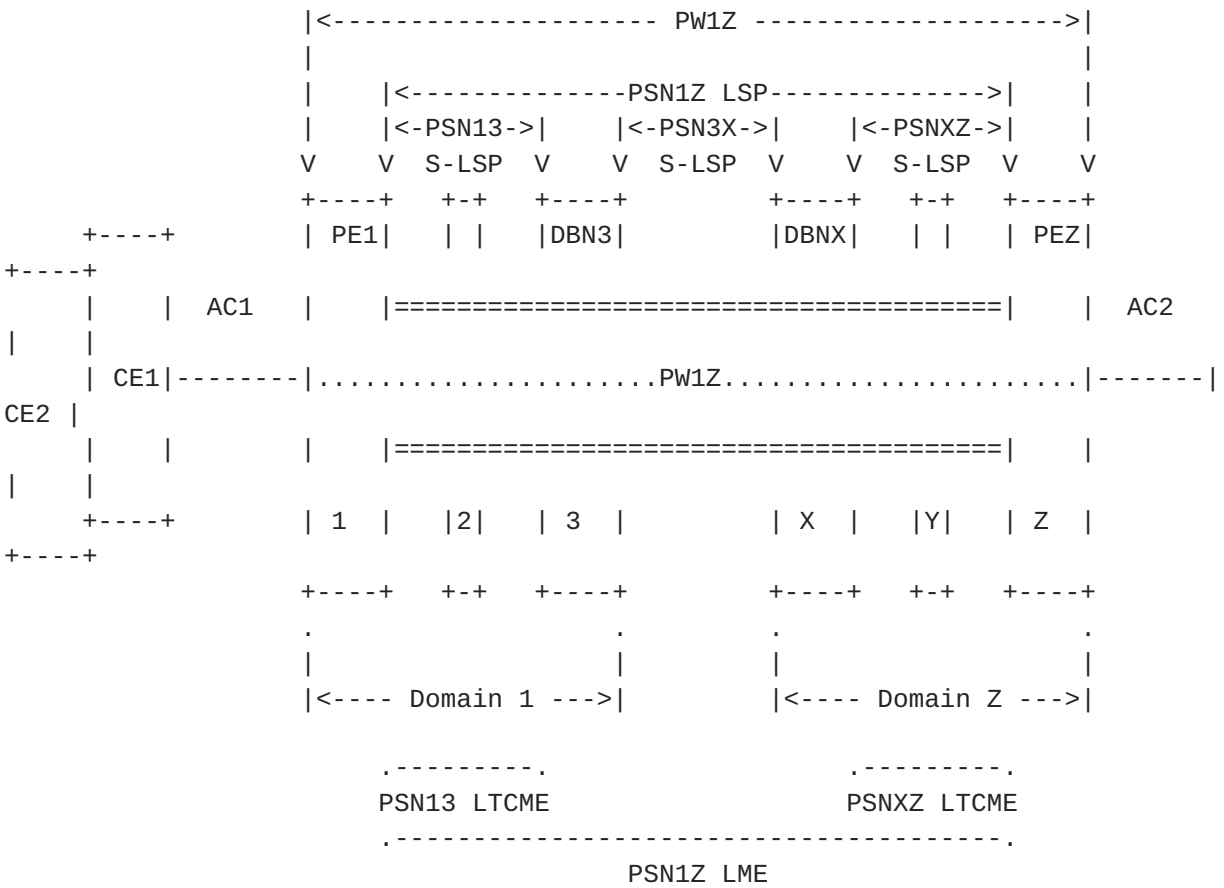

DBN: Domain Border Node

Figure 6 MPLS-TP LSP Tandem Connection Monitoring ME (LTCME)

Figure 6 depicts a variation of the reference model in Figure 3 where there is an end-to-end PSN LSP (PSN1Z LSP) between PE1 and PEZ. PSN1Z LSP consists of, at least, three stitched LSP Segments: PSN13, PSN3X and PSNXZ. In this scenario there are two separate LTCMEs configured to monitor the forwarding behaviour of the PSN1Z LSP: 1) a LTCME monitoring the PSN13 LSP Segment on Domain 1 (PSN13 LTCME), and 2) a LTCME monitoring the PSNXZ LSP Segment on Domain Z (PSNXZ LTCME).

It is worth noticing that LTCMEs can coexist with the LME monitoring the end-to-end LSP and that LTCME MEPs and LME MEPs can be coincident

in the same node (e.g. PE1 node supports both the PSN1Z LME MEP and the PSN13 LTCME MEP).

4.4. MPLS-TP PW Monitoring

An MPLS-TP PW ME (PME) is an MPLS-TP maintenance entity intended to monitor the end-to-end forwarding behaviour of a SS-PW or MS-PW between a pair of T-PEs. A PME MAY be configured on any SS-PW or MS-PW. PME OAM packets 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 the forwarding behaviour of an entire PW between a pair of MPLS-TP enabled T-PEs rather than monitoring the LSP aggregating multiple PWs between PEs.

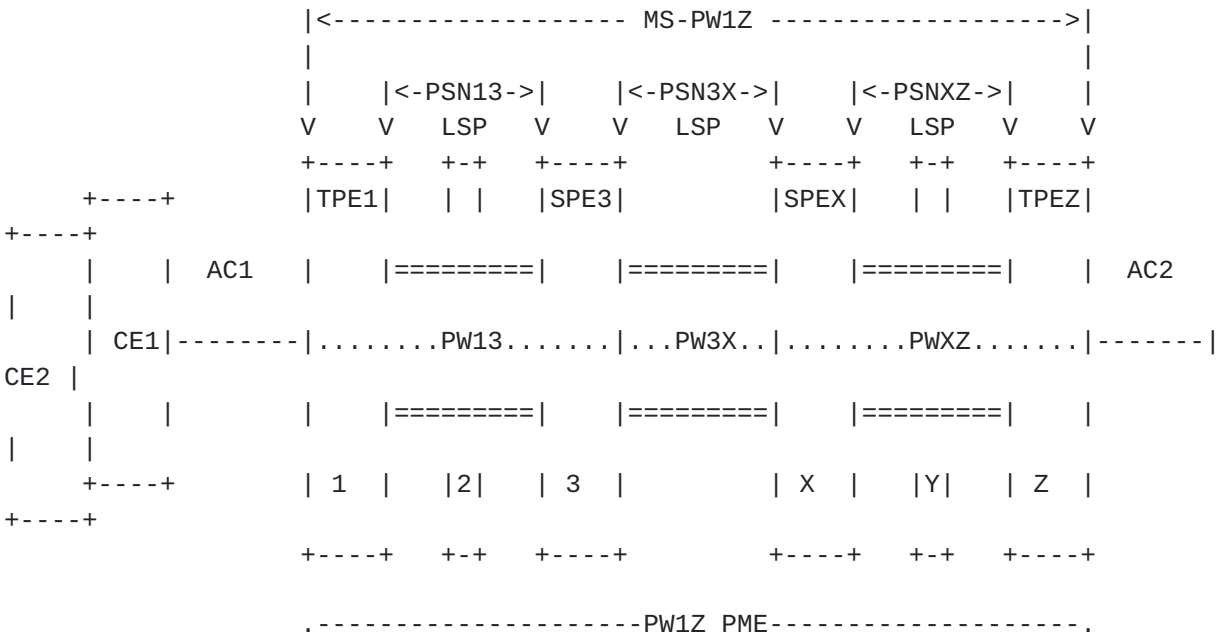

Figure 7 MPLS-TP PW ME (PME)

Figure 7 depicts a MS-PW (MS-PW1Z) consisting of three segments: PW13, PW3X and PWXZ and its associated end-to-end PME (PW1Z PME).

4.5. MPLS-TP MS-PW Tandem Connection Monitoring

An MPLS-TP MS-PW Tandem Connection Monitoring ME (PTCME) is an MPLS-TP maintenance entity intended to monitor the forwarding behaviour of an arbitrary part of an MS-PW between a given pair of PEs independently from the end-to-end monitoring (PME). An PTCME can monitor a PW segment or concatenated segment and it may alos include the forwarding engine(s) of the node(s) at the edge(s) of the segment or concatenated segment.

Multiple PTCMEs MAY be configured on any MS-PW. The PEs may or may not be immediately adjacent at the MS-PW layer. PTCME OAM packets

fate share with the user data packets sent over the monitored PW Segment.

A PTCME 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. It can span several PW segments.

A PTCME 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 Segment within a given network domain of an inter-domain MS-PW.

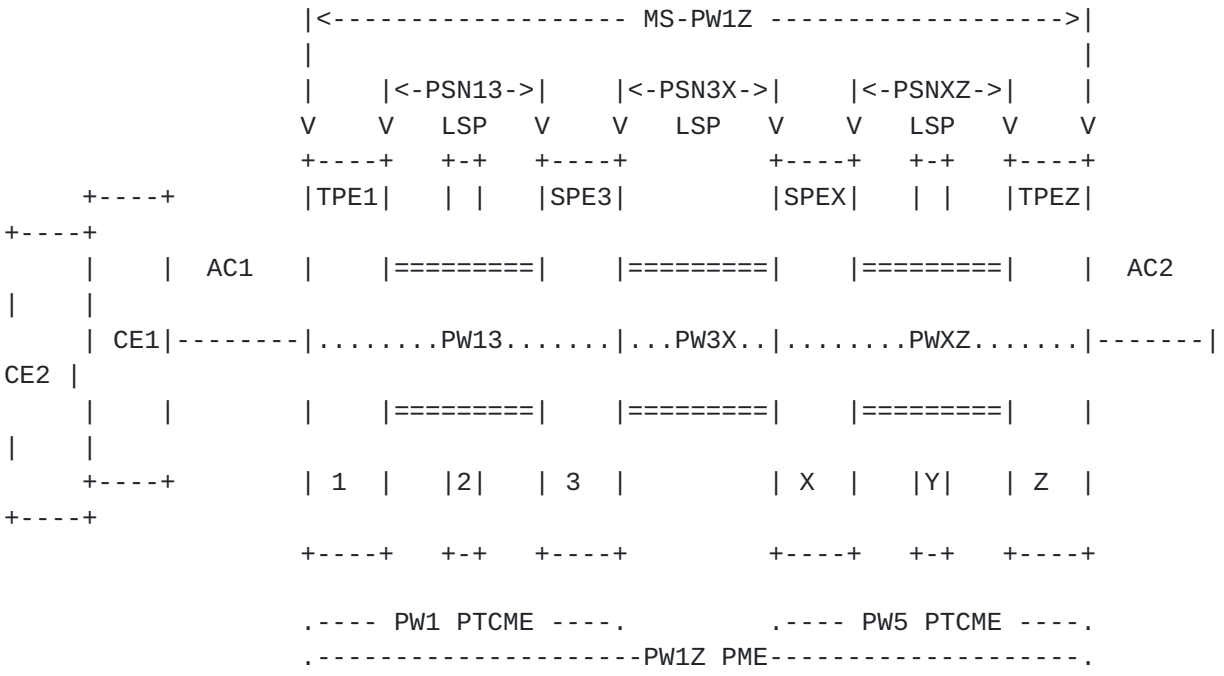

Figure 8 MPLS-TP MS-PW Tandem Connection Monitoring (PTCME)

Figure 8 depicts the same MS-PW (MS-PW1Z) between AC1 and AC2 as in Figure 7. In this scenario there are two separate PTCMEs configured to monitor the forwarding behaviour of MS-PW1Z: 1) a PTCME monitoring the PW13 MS-PW Segment on Domain 1 (PW13 PTCME), and 2) a PTCME monitoring the PWXZ MS-PW Segment on Domain Z with (PWXZ PTCME).

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

5. OAM Functions for pro-active monitoring

5.1. Continuity Check and Connectivity Verification

[Editor's note - There is a need to decide whether pro-active CC and CV functions need to be combined in the same tool or separated into different tools. This version of the document combines the two functions. Future versions will be aligned with the decision taken about whether to combine or not CC and CV.]

Proactive Continuity Check and Connectivity Verification (CC-V) functions are used to detect a loss of continuity (LOC), an unexpected connectivity between two MEs (e.g. mismerging or misconnection), as well as unexpected connectivity within the ME with an unexpected MEP.

Execution of proactive CC-V is based on the (proactive) generation of CC-V OAM packets by the source MEP that are processed by the sink MEP. Each CC-V OAM packet MUST include a globally unique ME identifier, and MUST be transmitted at a regular, operator's configurable, rate. The default CC-V transmission periods are application dependent (see section 5.1.4).

Proactive CC-V OAM packets are transmitted with the "minimum loss probability PHB" within a single network operator. This PHB is configurable on network operator's basis. PHBs can be translated at the network borders.

[Editor's note - Describe the relation between the previous paragraph and the fate sharing requirement. Need to clarify also in the requirement document that for proactive CC-V the fate sharing is related to the forwarding behavior and not to the QoS behavior]

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. 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.

To initialize the proactive CC-V monitoring on a configured ME without affecting traffic, 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 following defect conditions.

For all of these defect cases, the sink MEP SHOULD notify the equipment fault management process of the detected defect.

[Editor's note - Investigate whether in case there is a misconfiguration on the minimum loss probability PHB on the two MEPs a defect can be useful to notify the misconfiguration to the operator.]

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 peer MEP.

- Entry criteria: if no pro-active CC-V OAM packets from the peer MEP (i.e. with the correct ME and peer MEP identifiers) are received within the interval equal to 3.5 times the receiving MEP's configured CC-V transmission period.

- Exit criteria: a pro-active CC-V OAM packet from the peer MEP (i.e. with the correct ME and peer MEP identifiers) 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 or misconnection) with its peer source MEP when the received packet carries an incorrect ME identifier.

- Entry criteria: the sink MEP receives a pro-active CC-V OAM packet with an incorrect ME ID.

- o Exit criteria: the sink MEP does not receive any pro-active CC-V OAM packet with an incorrect ME ID 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 ME ID since this defect has been raised.

5.1.1.3. MEP misconfiguration defect

When a pro-active CC-V packet is received, a sink MEP identifies a MEP misconfiguration defect with its peer source MEP when the received packet carries a correct ME Identifier but an unexpected peer MEP Identifier which includes the MEP's own MEP Identifier.

- o Entry criteria: the sink MEP receives a CC-V pro-active packet with correct ME ID but with unexpected MEP ID.

- o Exit criteria: the sink MEP does not receive any pro-active CC-V OAM packet with a correct ME ID and unexpected MEP ID 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 ME ID and unexpected MEP ID since this defect has been raised.

5.1.1.4. Period Misconfiguration defect

If pro-active CC-V OAM packets is received with a correct ME and MEP identifiers but with a transmission period different than its own configured transmission period, then a CC-V period mis-configuration defect is detected

- o Entry criteria: a MEP receives a CC-V pro-active packet with correct ME ID and MEP ID but with a Period field value different than its own CC-V configured transmission period.

- o Exit criteria: the sink MEP does not receive any pro-active CC-V OAM packet with a correct ME and MEP IDs 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 ME and MEP IDs and an incorrect transmission period 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 MUST perform the following consequent actions. Some of these consequent actions SHOULD be enabled/disabled by the operator depending upon the application used (see section 5.1.4).

If a MEP detects an unexpected ME Identifier, or an unexpected MEP, it MUST block all the traffic (including also the user data packets) that it receives from the misconnected transport path.

If a MEP detects LOC defect and the CC-V monitoring is enabled 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 is enabled on every ME in order to reliably detect connectivity defects. However, CC-V proactive monitoring MAY be disabled by an operator on a ME. 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 MAY 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, an unexpected ME Identifier, or an unexpected MEP it MUST declare a signal fail condition at the transport path level.

If a MEP detects an Unexpected Period defect it SHOULD declare a signal fail condition at the transport path level.

[Editor's note - Transport equipment also performs defect correlation (as defined in G.806) in order to properly report failures to the transport NSM. The current working assumption, to be further investigated, is that defect correlations are outside the scope of this document and to be defined in ITU-T documents.]

5.1.3. Configuration considerations

At all MEPs inside a ME, the following configuration information need to be configured when pro-active CC-V function is enabled:

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

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

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

- transmission rate; the default CC-V transmission periods are application dependent (see section 5.1.4)

- 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.

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

5.1.4. Applications for proactive CC-V

CC-V is applicable for 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)

- Protection Switching: in order to achieve sub-50ms recovery time the default transmission period is 3.33ms (i.e. transmission rate of 300 packets/second) although a transmission period of 10ms can also be used. In some cases, when a slower recovery time is acceptable, it is also possible to lengthen the transmission rate.

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

In addition, the operator should be able to define the consequent action to be performed for each of these applications.

5.2. Remote Defect Indication

The Remote Defect Indication (RDI) is an indicator that is transmitted by a MEP to communicate to its peer MEPs that a signal fail condition exists. RDI is only used for bidirectional connections and is associated with proactive CC-V activation. 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 or an AIS condition is detected), 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.

[Editor's note - Add some forward compatibility information to cover the case where future OAM mechanisms that contributes to the signal fail detection (and RDI generation) are defined.]

A MEP that receives the packets 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, the RDI transmission rate and PHB of the MEP should be configured as part of the CC-V configuration.

5.2.2. Applications for Remote Defect Indication

RDI is applicable for the following applications:

- Single-ended fault management - A MEP that receives an RDI indication from its peer MEP, can report a far-end defect condition (i.e. the peer MEP has detected a signal fail condition in the traffic direction from the MEP that receives the RDI indication to the peer MEP that has sent the RDI information).

- Contribution to far-end performance monitoring - The indication of the far-end defect condition is used as a contribution to the bidirectional performance monitoring process.

5.3. Alarm Reporting

Alarm Reporting function relies upon an Alarm Indication Signal (AIS) message used to suppress alarms following detection of defect conditions at the server (sub) layer.

- A server MEP that detects a signal fail conditions in the server (sub-)layer, can generate packets with AIS information in a direction opposite to its peers MEPs to allow the suppression of secondary alarms at the MEP in the client (sub-)layer.

A server MEP is responsible for notifying the MPLS-TP layer network MEP upon fault detection in the server layer network to which the server MEP is associated.

Only Server MEPs can issue MPLS-TP packets with AIS information. Upon detection of a signal fail condition the Server MEP can immediately start transmitting periodic packets with AIS information. These periodic packets, with AIS information, continue to be transmitted until the signal fail condition is cleared.

Upon receiving a packet with AIS information an MPLS-TP MEP detects an AIS defect condition and suppresses loss of continuity alarms associated with all its peer MEPs. A MEP resumes loss of continuity alarm generation upon detecting loss of continuity defect conditions in the absence of AIS condition.

For example, let's consider a fiber cut between LSR 1 and LSR 2 in the reference network of Figure 3. Assuming that all the MEs described in Figure 3 have pro-active CC-V enabled, a LOC defect is detected by the MEPs of Sec12 SME, PSN13 LME, PW1 PTCME and PW1Z PME, however in transport network only the alarm associate to the fiber cut needs to be reported to NMS while all these 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 generates AIS packets on the PSN13 LME in order to allow its MEP in LSR3 to suppress the LOC alarm.

LSR3 can also suppress the secondary alarm on PW1 PTCME because the MEP of PW1 PTCME resides within the same node as the MEP of PSN13 LME.

The MEP of PW1 PTCME in LSR3 also 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 MEs in the 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.

A MIP is transparent to packets with AIS information and therefore does not require any information to support AIS functionality.

5.4. Lock Reporting

[Editor's note - Requirements for Lock Indication and Lock Reporting are still under discussion in draft-ietf-mpls-tp-oam-requirement-02.

Lock Indication is not required by draft-ietf-mpls-tp-oam-requirement-02 This section will be aligned according to the final decision regarding the requirement.]

To be incorporated in a future revision of this document

5.5. Lock Indication

[Editor's note - Requirements for Lock Indication and Lock Reporting are still under discussion in draft-ietf-mpls-tp-oam-requirement-02.

Lock Indication is not required by draft-ietf-mpls-tp-oam-requirement-02 This section will be aligned according to the final decision regarding the requirement.]

The Locked Indication Signal (LIS) is used to propagate an administrative locking of a source MEP and consequential interruption of data forwarding towards the sink MEP. It allows a sink MEP receiving LIS to differentiate between a defect condition and an administrative locking action at the source MEP. An example application that requires administrative locking of a MEP is the out-of-service test.

5.6. Packet Loss

Packet Loss (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. 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.

Pro-active 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 to derive the impact of packet loss on a number of performance metrics for the transport path.

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).

5.6.1. Configuration considerations

In order to support pro-active LM, the transmission rate and PHB associated with the LM OAM packets originating from a MEP need be configured as part of the LM provisioning procedures. LM OAM packets should be transmitted with the PHB that yields the lowest packet loss

performance among the PHB Scheduling Classes or Ordered Aggregates (see RFC 3260 [14]) in the monitored transport path for the relevant network domain(s).

5.6.2. Applications for Packet Loss

LM is relevant for the following applications:

- Single or double-end performance monitoring: determination of the packet loss performance of a transport path for SLS verification purposes.

- Single or double-end performance monitoring: determination of the packet loss performance of a PHB Scheduling Class or Ordered Aggregate within a transport path

- Contribution to service unable time. Both near-end and far-end packet loss measurements contribute to performance metrics such as near-end severely errored seconds (Near-End SES) and far-end severely errored seconds (Far-End SES) respectively, which together contribute to unavailable time, in a manner similar to Recommendation G.826 [19] and Recommendation G.7710 [20].

5.7. Client Failure Indication

The Client Failure Indication (CSF) function 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 transmission path does not support a native defect/alarm indication mechanism, e.g. FDI/AIS.

A source MEP starts transmitting a CSF 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.

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

5.7.1. Configuration considerations

In order to support RCSF indication, the CSF transmission rate and PHB of the CSF OAM packet should be configured as part of the CSF configuration.

5.7.2. Applications for Remote Defect Indication

CSF is applicable for the following applications:

- Single-ended fault management - A MEP that receives a CSF indication from its peer MEP, can report a far-end client defect condition (i.e. the peer MEP has been informed of local client signal fail condition in the traffic direction from the client to the peer MEP that transmitted the CSF).

- Contribution to far-end performance monitoring - The indication of the far-end defect condition may be used to account on network operator contribution to the bidirectional performance monitoring process.

CSF supports the application described in Appendix VIII of ITU-T G.806 [18].

5.8. Delay Measurement

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. 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.

Pro-active 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.

- 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.

5.8.1. 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 LM provisioning procedures. DM OAM packets should be transmitted with the PHB that yields the lowest packet loss performance among the PHB Scheduling Classes or Ordered Aggregates (see RFC 3260 [14]) in the monitored transport path for the relevant network domain(s).

5.8.2. Applications for Packet Loss

DM is relevant for the following applications:

- Single or double-end performance monitoring: determination of the delay performance of a transport path for SLS verification purposes.

- Single or double-end performance monitoring: determination of the delay performance of a PHB Scheduling Class or Ordered Aggregate within a transport path

6. OAM Functions for on-demand monitoring

6.1. Connectivity Verification

In order to preserve network resources, e.g. bandwidth, processing time at switches, it may be preferable to not use pro-active CC-V. In order to perform fault management functions network management may invoke periodic on-demand bursts of on-demand CV packets. Use of on-demand CV is dependent on the existence of a bi-directional connection ME, because it requires the presence of a return path in the data plane.

[Editor's note - Clarify in the sentence above and within the paragraph that on-demand CV requires a return path to send back the reply to on-demand CV packets]

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 ME that is being checked. The on-demand functionality may be used to check either an entire ME (end-to-end) or between a MEP to a specific MIP.

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 ME, the source MEP should issue a burst of on-demand CV packets that uniquely identifies the ME being verified. The number of packets and their transmission rate should be pre-configured and known to both the source MEP and the target MEP or MIP. The source MEP should use the TTL field to indicate the number of hops necessary, when targeting a MIP and use the default value when performing an end-to-end check [IB => This is quite generic for addressing packets to MIPs and MEPs so it is better to move this text in section 2]. 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, a LOC state is detected.

[Editor's note - We need to add some text for the usage of on-demand CV with different packet sizes, e.g. to discover MTU problems.]

When a connectivity problem is detected (e.g. via a pro-active CC-V OAM tool), on demand CV tool can be used to check the path. The series should check CV from MEP to peer MEP on the path, and if a fault is discovered, by lack of response, then additional checks may be performed to each of the intermediate MIP to locate the fault.

6.1.1. Configuration considerations

For on-demand CV the MEP should support configuration of number of packets to be transmitted/received in each burst of transmissions and their packet size. The transmission rate should be either pre-configured or negotiated between the different nodes.

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.

[Editor's note - We need to be better define the reason for such configuration]

6.2. Packet Loss

On-demand Packet Loss (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 Pro-active 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.

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 evaluate the packet loss performance metrics of the transport path.

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 procedures. LM OAM packets should be transmitted with the PHB that yields the lowest packet loss performance among the PHB Scheduling Classes or Ordered Aggregates (see RFC 3260 [14]) in the monitored transport path for the relevant network domain(s).

6.2.2. Applications for On-demand Packet Loss

On-demand LM is relevant for the following applications:

- Single-end performance monitoring: diagnostic of the packet loss performance of a transport path for SLS trouble shooting purposes.

- Single-end performance monitoring: diagnostic of the packet loss performance of a PHB Scheduling Class or Ordering Aggregate within a transport path for QoS trouble shooting purposes.

6.3. Diagnostic

To be incorporated in a future revision of this document

6.4. Route Tracing

To be incorporated in a future revision of this document

6.5. Delay Measurement

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. 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:

- 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.

- 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.

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 LM provisioning procedures. DM OAM packets should be transmitted with the PHB that yields the lowest packet delay performance among the PHB Scheduling Classes or Ordering Aggregates (see RFC 3260 [14]) in the monitored transport path for the relevant network domain(s).

6.5.2. Applications for Delay Measurement

DM is relevant for the following applications:

- o Single or double-end performance monitoring: determination of the packet delay and/or delay variation performance of a transport path for SLS verification purposes.

- o Single or double-end performance monitoring: determination of the packet delay and/or delay variation a PHB Scheduling Class or Ordering Aggregate within a transport path

- o Contribution to service unable time. Packet delay measurements may contribute to performance metrics such as near-end severely errored seconds (Near-End SES) and far-end severely errored seconds (Far-End SES), which together contribute to unavailable time.

6.6. Lock Instruct

To be incorporated in a future revision of this document

7. Security Considerations

A number of security considerations 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.

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

8. IANA Considerations

No new IANA considerations.

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

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

10. References

10.1. Normative References

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

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

[3] Rosen, E., et al., "MPLS Label Stack Encoding", RFC 3032, January 2001

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

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

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

[7] Bocci, M., Bryant, S., "An Architecture for Multi-Segment Pseudo Wire Emulation Edge-to-Edge", draft-ietf-pwe3-ms-pw-arch-05 (work in progress), September 2008

[8] Bocci, M., et al., "A Framework for MPLS in Transport Networks", draft-ietf-mpls-tp-framework-01 (work in progress), June 2009

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

10.2. Informative References

[10] Niven-Jenkins, B., Brungard, D., Betts, M., sprecher, N., Ueno, S., "MPLS-TP Requirements", draft-ietf-mpls-tp-requirements-09 (work in progress), June 2009

[11] Vigoureux, M., Betts, M., Ward, D., "Requirements for OAM in MPLS Transport Networks", draft-ietf-mpls-tp-oam-requirements-02 (work in progress), June 2009

[12] Sprecher, N., Nadeau, T., van Helvoort, H., Weingarten, Y., "MPLS-TP OAM Analysis", draft-sprecher-mpls-tp-oam-analysis-04 (work in progress), May 2009

[13] 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

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

[15] Farrel, A., Ayyangar, A., Vasseur, JP., "Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 5151, February 2008

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

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

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

[19] ITU-T Recommendation G.826 (12/02), "End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections", December 2002

[20] ITU-T Recommendation G.7710 (07/07), "Common equipment management function requirements", July 2007

[21] ITU-T Recommendation Y.2611 (06/12), " High-level architecture of future packet-based networks", 2006

Authors' Addresses

Italo Busi (Editor)
Alcatel-Lucent

Email: Italo.Busi@alcatel-lucent.it

Ben Niven-Jenkins (Editor)
BT

Email: benjamin.niven-jenkins@bt.com

Contributing Authors' Addresses

Annamaria Fulignoli
Ericsson

Email: annamaria.fulignoli@ericsson.com

Enrique Hernandez-Valencia
Alcatel-Lucent

Email: enrique@alcatel-lucent.com

Lieven Levrau
Alcatel-Lucent

Email: llevrau@alcatel-lucent.com

Dinesh Mohan
Nortel

Email: mohand@nortel.com

Vincenzo Sestito
Alcatel-Lucent

Email: vincenzo.sestito@alcatel-lucent.it

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.fr

Yaacov Weingarten
Nokia Siemens Networks

Email: yaacov.weingarten@nsn.com

Rolf Winter
NEC

Email: Rolf.Winter@nw.neclab.eu