Internet Engineering Task Force (IETF)
Request for Comments: 5921
Category: Informational
ISSN: 2070-1721

M. Bocci, Ed.
Alcatel-Lucent
S. Bryant, Ed.
D. Frost, Ed.
Cisco Systems
L. Levrau
Alcatel-Lucent
L. Berger
LabN

July 2010

A Framework for MPLS in Transport Networks

Abstract

This document specifies an architectural framework for the application of Multiprotocol Label Switching (MPLS) to the construction of packet-switched transport networks. It describes a common set of protocol functions -- the MPLS Transport Profile (MPLS-TP) -- that supports the operational models and capabilities typical of such networks, including signaled or explicitly provisioned bidirectional connection-oriented paths, protection and restoration mechanisms, comprehensive Operations, Administration, and Maintenance (OAM) functions, and network operation in the absence of a dynamic control plane or IP forwarding support. Some of these functions are defined in existing MPLS specifications, while others require extensions to existing specifications to meet the requirements of the MPLS-TP.

This document defines the subset of the MPLS-TP applicable in general and to point-to-point transport paths. The remaining subset, applicable specifically to point-to-multipoint transport paths, is outside the scope of this document.

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 Pseudowire Emulation Edge-to-Edge (PWE3) architectures to support the capabilities and functionalities of a packet transport network as defined by the ITU-T.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

1.1. Motivation and Background

This document describes an architectural framework for the application of MPLS to the construction of packet-switched transport networks. It specifies the common set of protocol functions that meet the requirements in [RFC5654], and that together constitute the MPLS Transport Profile (MPLS-TP) for point-to-point transport paths. The remaining MPLS-TP functions, applicable specifically to point-to-multipoint transport paths, are outside the scope of this document.

Historically, the optical transport infrastructure -- Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) and Optical Transport Network (OTN) -- has provided carriers with a high benchmark for reliability and operational simplicity. To achieve this, transport technologies have been designed with specific characteristics:

- Strictly connection-oriented connectivity, which may be long-lived and may be provisioned manually, for example, by network management systems or direct node configuration using a command line interface.
- A high level of availability.
- Quality of service.
- Extensive Operations, Administration, and Maintenance (OAM) capabilities.

Carriers wish to evolve such transport networks to take advantage of the flexibility and cost benefits of packet switching technology and to support packet-based services more efficiently. While MPLS is a maturing packet technology that already plays an important role in transport networks and services, not all MPLS capabilities and mechanisms are needed in, or consistent with, the transport network operational model. There are also transport technology characteristics that are not currently reflected in MPLS.

There are thus two objectives for MPLS-TP:

1. To enable MPLS to be deployed in a transport network and operated in a similar manner to existing transport technologies.

2. To enable MPLS to support packet transport services with a similar degree of predictability to that found in existing transport networks.

In order to achieve these objectives, there is a need to define a common set of MPLS protocol functions -- an MPLS Transport Profile -- for the use of MPLS in transport networks and applications. Some of the necessary functions are provided by existing MPLS specifications, while others require additions to the MPLS tool-set. Such additions should, wherever possible, be applicable to MPLS networks in general as well as those that conform strictly to the transport network model.

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.

1.2. Scope

This document describes an architectural framework for the application of MPLS to the construction of packet-switched transport networks. It specifies the common set of protocol functions that meet the requirements in [RFC5654], and that together constitute the MPLS Transport Profile (MPLS-TP) for point-to-point MPLS-TP transport paths. The remaining MPLS-TP functions, applicable specifically to point-to-multipoint transport paths, are outside the scope of this document.

1.3. Terminology

Term	Definition
AC	Attachment Circuit
ACH	Associated Channel Header
Adaptation	The mapping of client information into a format suitable for transport by the server layer
APS	Automatic Protection Switching
ATM	Asynchronous Transfer Mode
BFD	Bidirectional Forwarding Detection
CE	Customer Edge

CL-PS	Connectionless - Packet Switched
CM	Configuration Management
CO-CS	Connection Oriented - Circuit Switched
CO-PS	Connection Oriented - Packet Switched
DCN	Data Communication Network
EMF	Equipment Management Function
FCAPS	Fault, Configuration, Accounting, Performance, and Security
FM	Fault Management
G-ACh	Generic Associated Channel
GAL	G-ACh Label
LER	Label Edge Router
LSP	Label Switched Path
LSR	Label Switching Router
MAC	Media Access Control
MCC	Management Communication Channel
ME	Maintenance Entity
MEG	Maintenance Entity Group
MEP	Maintenance Entity Group End Point
MIP	Maintenance Entity Group Intermediate Point
MPLS	Multiprotocol Label Switching
MPLS-TP	MPLS Transport Profile
MPLS-TP P	MPLS-TP Provider LSR
MPLS-TP PE	MPLS-TP Provider Edge LSR
MS-PW	Multi-Segment Pseudowire
Native Service	The traffic belonging to the client of the MPLS-TP network
OAM	Operations, Administration, and Maintenance (see [OAM-DEF])
OSI	Open Systems Interconnection
OTN	Optical Transport Network
PDU	Protocol Data Unit
PM	Performance Monitoring
PSN	Packet Switching Network
PW	Pseudowire
SCC	Signaling Communication Channel
SDH	Synchronous Digital Hierarchy
S-PE	PW Switching Provider Edge
SPME	Sub-Path Maintenance Element
SS-PW	Single-Segment Pseudowire
T-PE	PW Terminating Provider Edge
TE LSP	Traffic Engineered Label Switched Path
VCCV	Virtual Circuit Connectivity Verification

1.3.1. Transport Network

A Transport Network provides transparent transmission of user traffic between attached client devices by establishing and maintaining point-to-point or point-to-multipoint connections between such devices. The architecture of networks supporting point-to-multipoint connections is outside the scope of this document. A Transport Network is independent of any higher-layer network that may exist between clients, except to the extent required to supply this transmission service. In addition to client traffic, a Transport Network may carry traffic to facilitate its own operation, such as that required to support connection control, network management, and Operations, Administration, and Maintenance (OAM) functions.

See also the definition of packet transport service in Section 3.1.

1.3.2. MPLS Transport Profile

The MPLS Transport Profile (MPLS-TP) is the subset of MPLS functions that meet the requirements in [RFC5654]. Note that MPLS is defined to include any present and future MPLS capability specified by the IETF, including those capabilities specifically added to support transport network requirements [RFC5654].

1.3.3. MPLS-TP Section

MPLS-TP sections are defined in [DATA-PLANE]. See also the definition of "section layer network" in Section 1.2.2 of [RFC5654].

1.3.4. MPLS-TP Label Switched Path

An MPLS-TP Label Switched Path (MPLS-TP LSP) is an LSP that uses a subset of the capabilities of an MPLS LSP in order to meet the requirements of an MPLS transport network as set out in [RFC5654]. The characteristics of an MPLS-TP LSP are primarily that it:

1. Uses a subset of the MPLS OAM tools defined in [OAM-FRAMEWORK].

2. Supports 1+1, 1:1, and 1:N protection functions.

3. Is traffic engineered.

4. May be established and maintained via the management plane, or using GMPLS protocols when a control plane is used.

5. Is either point-to-point or point-to-multipoint. Multipoint-to-point and multipoint-to-multipoint LSPs are not supported.

6. It is either unidirectional, associated bidirectional, or co-routed bidirectional (i.e., the forward and reverse components of a bidirectional LSP follow the same path, and the intermediate nodes are aware of their association). These are further defined in [DATA-PLANE].

Note that an MPLS LSP is defined to include any present and future MPLS capability, including those specifically added to support the transport network requirements.

See [DATA-PLANE] for further details on the types and data-plane properties of MPLS-TP LSPs.

The lowest server layer provided by MPLS-TP is an MPLS-TP LSP. The client layers of an MPLS-TP LSP may be network-layer protocols, MPLS LSPs, or PWs. The relationship of an MPLS-TP LSP to its client layers is described in detail in Section 3.4.

1.3.5. MPLS-TP Label Switching Router

An MPLS-TP Label Switching Router (LSR) is either an MPLS-TP Provider Edge (PE) router or an MPLS-TP Provider (P) router for a given LSP, as defined below. The terms MPLS-TP PE router and MPLS-TP P router describe logical functions; a specific node may undertake only one of these roles on a given LSP.

Note that the use of the term "router" in this context is historic and neither requires nor precludes the ability to perform IP forwarding.

1.3.5.1. Label Edge Router

An MPLS-TP Label Edge Router (LER) is an LSR that exists at the endpoints of an LSP and therefore pushes or pops the LSP label, i.e., does not perform a label swap on the particular LSP under consideration.

1.3.5.2. MPLS-TP Provider Edge Router

An MPLS-TP Provider Edge (PE) router is an MPLS-TP LSR that adapts client traffic and encapsulates it to be transported over an MPLS-TP LSP. Encapsulation may be as simple as pushing a label, or it may require the use of a pseudowire. An MPLS-TP PE exists at the interface between a pair of layer networks. For an MS-PW, an MPLS-TP PE may be either an S-PE or a T-PE, as defined in [RFC5659] (see below). A PE that pushes or pops an LSP label is an LER for that LSP.

The term Provider Edge refers to the node's role within a provider's network. A provider edge router resides at the edge of a given MPLS-TP network domain, in which case it has links to another MPLS-TP network domain or to a CE, except for the case of a pseudowire switching provider edge (S-PE) router, which is not restricted to the edge of an MPLS-TP network domain.

1.3.5.3. MPLS-TP Provider Router

An MPLS-TP Provider router is an MPLS-TP LSR that does not provide MPLS-TP PE functionality for a given LSP. An MPLS-TP P router switches LSPs that carry client traffic, but does not adapt client traffic and encapsulate it to be carried over an MPLS-TP LSP. The term Provider Router refers to the node's role within a provider's network. A provider router does not have links to other MPLS-TP network domains.

1.3.5.4. Pseudowire Switching Provider Edge Router (S-PE)

RFC 5659 [RFC5659] defines an S-PE as:

> A PE capable of switching the control and data planes of the preceding and succeeding PW segments in an MS-PW. The S-PE terminates the PSN tunnels of the preceding and succeeding segments of the MS-PW. It therefore includes a PW switching point for an MS-PW. A PW switching point is never the S-PE and the T-PE for the same MS-PW. A PW switching point runs necessary protocols to set up and manage PW segments with other PW switching points and terminating PEs. An S-PE can exist anywhere a PW must be processed or policy applied. It is therefore not limited to the edge of a provider network.
>
> Note that it was originally anticipated that S-PEs would only be deployed at the edge of a provider network where they would be used to switch the PWs of different service providers. However, as the design of MS-PW progressed, other applications for MS-PW were recognized. By this time S-PE had become the accepted term for the equipment, even though they were no longer universally deployed at the provider edge.

1.3.5.5. Pseudowire Terminating Provider Edge (T-PE) Router

RFC 5659 [RFC5659] defines a T-PE as:

> A PE where the customer-facing attachment circuits (ACs) are bound to a PW forwarder. A terminating PE is present in the first and last segments of an MS-PW. This incorporates the functionality of a PE as defined in RFC 3985.

1.3.6. Customer Edge (CE)

A Customer Edge (CE) is the client function that sources or sinks native service traffic to or from the MPLS-TP network. CEs on either side of the MPLS-TP network are peers and view the MPLS-TP network as a single link.

1.3.7. Transport LSP

A Transport LSP is an LSP between a pair of PEs that may transit zero or more MPLS-TP provider routers. When carrying PWs, the Transport LSP is equivalent to the PSN tunnel LSP in [RFC3985] terminology.

1.3.8. Service LSP

A service LSP is an LSP that carries a single client service.

1.3.9. Layer Network

A layer network is defined in [G.805] and described in [RFC5654]. 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 a 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.

1.3.10. Network Layer

This document uses the term Network Layer in the same sense as it is used in [RFC3031] and [RFC3032]. Network-layer protocols are synonymous with those belonging to Layer 3 of the Open System Interconnect (OSI) network model [X.200].

1.3.11. Service Interface

The packet transport service provided by MPLS-TP is provided at a service interface. Two types of service interfaces are defined:

- User-Network Interface (UNI) (see Section 3.4.3.1).
- Network-Network Interface (NNI) (see Section 3.4.3.2).

A UNI service interface may be a Layer 2 interface that carries only network layer clients. MPLS-TP LSPs are both necessary and sufficient to support this service interface as described in Section 3.4.3. Alternatively, it may be a Layer 2 interface that carries both network-layer and non-network-layer clients. To support this service interface, a PW is required to adapt the client traffic received over the service interface. This PW in turn is a client of the MPLS-TP server layer. This is described in Section 3.4.2.

An NNI service interface may be to an MPLS LSP or a PW. To support this case, an MPLS-TP PE participates in the service interface signaling.

1.3.12. Native Service

The native service is the client layer network service that is transported by the MPLS-TP network, whether a pseudowire or an LSP is used for the adaptation (see Section 3.4).

1.3.13. Additional Definitions and Terminology

Detailed definitions and additional terminology may be found in [RFC5654] and [ROSETTA-STONE].

2. MPLS Transport Profile Requirements

The requirements for MPLS-TP are specified in [RFC5654], [RFC5860], and [NM-REQ]. This section provides a brief reminder to guide the reader. It is not normative or intended as a substitute for these documents.

MPLS-TP must not modify the MPLS forwarding architecture and must be based on existing pseudowire and LSP constructs.

Point-to-point LSPs may be unidirectional or bidirectional, and it must be possible to construct congruent bidirectional LSPs.

MPLS-TP LSPs do not merge with other LSPs at an MPLS-TP LSR and it must be possible to detect if a merged LSP has been created.

It must be possible to forward packets solely based on switching the MPLS or PW label. It must also be possible to establish and maintain LSPs and/or pseudowires both in the absence or presence of a dynamic control plane. When static provisioning is used, there must be no dependency on dynamic routing or signaling.

OAM and protection mechanisms, and forwarding of data packets, must be able to operate without IP forwarding support.

It must be possible to monitor LSPs and pseudowires through the use of OAM in the absence of control-plane or routing functions. In this case, information gained from the OAM functions is used to initiate path recovery actions at either the PW or LSP layers.

3. MPLS Transport Profile Overview

3.1. Packet Transport Services

One objective of MPLS-TP is to enable MPLS networks to provide packet transport services with a similar degree of predictability to that found in existing transport networks. Such packet transport services exhibit a number of characteristics, defined in [RFC5654]:

- In an environment where an MPLS-TP layer network is supporting a client layer network, and the MPLS-TP layer network is supported by a server layer network then operation of the MPLS-TP layer network must be possible without any dependencies on either the server or client layer network.

- The service provided by the MPLS-TP network to a given client will not fall below the agreed level as a result of the traffic loading of other clients.

- The control and management planes of any client network layer that uses the service is isolated from the control and management planes of the MPLS-TP layer network, where the client network layer is considered to be the native service of the MPLS-TP network.

- Where a client network makes use of an MPLS-TP server that provides a packet transport service, the level of coordination required between the client and server layer networks is minimal (preferably no coordination will be required).

- The complete set of packets generated by a client MPLS(-TP) layer network using the packet transport service, which may contain packets that are not MPLS packets (e.g., IP or CLNS (Connectionless Network Service) packets used by the control/management plane of the client MPLS(-TP) layer network), are transported by the MPLS-TP server layer network.

- The packet transport service enables the MPLS-TP layer network addressing and other information (e.g., topology) to be hidden from any client layer networks using that service, and vice-versa.

These characteristics imply that a packet transport service does not support a connectionless packet-switched forwarding mode. However, this does not preclude it carrying client traffic associated with a connectionless service.

3.2. Scope of the MPLS Transport Profile

Figure 1 illustrates the scope of MPLS-TP. MPLS-TP solutions are primarily intended for packet transport applications. MPLS-TP is a strict subset of MPLS, and comprises only those functions that are necessary to meet the requirements of [RFC5654]. This includes MPLS functions that were defined prior to [RFC5654] but that meet the requirements of [RFC5654], together with additional functions defined to meet those requirements. Some MPLS functions defined before [RFC5654] such as Equal Cost Multi-Path (ECMP), LDP signaling when used in such a way that it creates multipoint-to-point LSPs, and IP forwarding in the data plane are explicitly excluded from MPLS-TP by that requirements specification.

Note that MPLS as a whole will continue to evolve to include additional functions that do not conform to the MPLS Transport Profile or its requirements, and thus fall outside the scope of MPLS-TP.

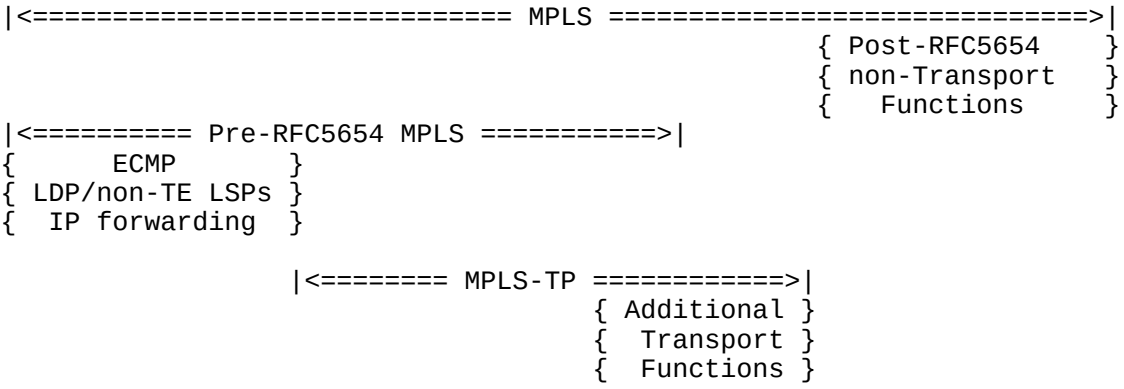

Figure 1: Scope of MPLS-TP

MPLS-TP can be used to construct packet networks and is therefore applicable in any packet network context. A subset of MPLS-TP is also applicable to ITU-T-defined packet transport networks, where the transport network operational model is deemed attractive.

3.3. Architecture

MPLS-TP comprises the following architectural elements:

- A standard MPLS data plane [RFC3031] as profiled in [DATA-PLANE].
- Sections, LSPs, and PWs that provide a packet transport service for a client network.
- Proactive and on-demand Operations, Administration, and Maintenance (OAM) functions to monitor and diagnose the MPLS-TP network, as outlined in [OAM-FRAMEWORK].
- Control planes for LSPs and PWs, as well as support for static provisioning and configuration, as outlined in [CP-FRAMEWORK].
- Path protection mechanisms to ensure that the packet transport service survives anticipated failures and degradations of the MPLS-TP network, as outlined in [SURVIVE-FWK].
- Control-plane-based restoration mechanisms, as outlined in [SURVIVE-FWK].
- Network management functions, as outlined in [NM-FRAMEWORK].

The MPLS-TP architecture for LSPs and PWs includes the following two sets of functions:

- MPLS-TP native service adaptation
- MPLS-TP forwarding

The adaptation functions interface the native service (i.e., the client layer network service) to MPLS-TP. This includes the case where the native service is an MPLS-TP LSP.

The forwarding functions comprise the mechanisms required for forwarding the encapsulated native service traffic over an MPLS-TP server layer network, for example, PW and LSP labels.

3.3.1. MPLS-TP Native Service Adaptation Functions

The MPLS-TP native service adaptation functions interface the client layer network service to MPLS-TP. For pseudowires, these adaptation functions are the payload encapsulation described in Section 4.4 of [RFC3985] and Section 6 of [RFC5659]. For network layer client services, the adaptation function uses the MPLS encapsulation format as defined in [RFC3032].

The purpose of this encapsulation is to abstract the data plane of the client layer network from the MPLS-TP data plane, thus contributing to the independent operation of the MPLS-TP network.

MPLS-TP is itself a client of an underlying server layer. MPLS-TP is thus also bounded by a set of adaptation functions to this server layer network, which may itself be MPLS-TP. These adaptation functions provide encapsulation of the MPLS-TP frames and for the transparent transport of those frames over the server layer network. The MPLS-TP client inherits its Quality of Service (QoS) from the MPLS-TP network, which in turn inherits its QoS from the server layer. The server layer therefore needs to provide the necessary QoS to ensure that the MPLS-TP client QoS commitments can be satisfied.

3.3.2. MPLS-TP Forwarding Functions

The forwarding functions comprise the mechanisms required for forwarding the encapsulated native service traffic over an MPLS-TP server layer network, for example, PW and LSP labels.

MPLS-TP LSPs use the MPLS label switching operations and Time-to-Live (TTL) processing procedures defined in [RFC3031], [RFC3032], and [RFC3443], as profiled in [DATA-PLANE]. These operations are highly optimized for performance and are not modified by the MPLS-TP profile.

In addition, MPLS-TP PWs use the SS-PW and optionally the MS-PW forwarding operations defined in [RFC3985] and [RFC5659].

Per-platform label space is used for PWs. Either per-platform, per-interface, or other context-specific label space [RFC5331] may be used for LSPs.

MPLS-TP forwarding is based on the label that identifies the transport path (LSP or PW). The label value specifies the processing operation to be performed by the next hop at that level of encapsulation. A swap of this label is an atomic operation in which the contents of the packet after the swapped label are opaque to the forwarder. The only event that interrupts a swap operation is TTL expiry. This is a fundamental architectural construct of MPLS to be taken into account when designing protocol extensions (such as those for OAM) that require packets to be sent to an intermediate LSR.

Further processing to determine the context of a packet occurs when a swap operation is interrupted in this manner, or a pop operation exposes a specific reserved label at the top of the stack, or the

packet is received with the GAL (Section 3.6) at the top of stack. Otherwise, the packet is forwarded according to the procedures in [RFC3032].

MPLS-TP supports Quality of Service capabilities via the MPLS Differentiated Services (Diffserv) architecture [RFC3270]. Both E-LSP and L-LSP MPLS Diffserv modes are supported.

Further details of MPLS-TP forwarding can be found in [DATA-PLANE].

3.4. MPLS-TP Native Service Adaptation

This document describes the architecture for two native service adaptation mechanisms, which provide encapsulation and demultiplexing for native service traffic traversing an MPLS-TP network:

- A PW
- An MPLS LSP

MPLS-TP uses IETF-defined pseudowires to emulate certain services, for example, Ethernet, Frame Relay, or PPP / High-Level Data Link Control (HDLC). A list of PW types is maintained by IANA in the "MPLS Pseudowire Type" registry. When the native service adaptation is via a PW, the mechanisms described in Section 3.4.4 are used.

An MPLS LSP can also provide the adaptation, in which case any native service traffic type supported by [RFC3031] and [RFC3032] is allowed. Examples of such traffic types include IP packets and MPLS-labeled packets. Note that the latter case includes TE-LSPs [RFC3209] and LSP-based applications such as PWs, Layer 2 VPNs [RFC4664], and Layer 3 VPNs [RFC4364]. When the native service adaptation is via an MPLS label, the mechanisms described in Section 3.4.5 are used.

3.4.1. MPLS-TP Client/Server Layer Relationship

The relationship between the client layer network and the MPLS-TP server layer network is defined by the MPLS-TP network boundary and the label context. It is not explicitly indicated in the packet. In terms of the MPLS label stack, when the native service traffic type is itself MPLS-labeled, then the S bits of all the labels in the MPLS-TP label stack carrying that client traffic are zero; otherwise, the bottom label of the MPLS-TP label stack has the S bit set to 1. In other words, there can be only one S bit set in a label stack.

The data-plane behavior of MPLS-TP is the same as the best current practice for MPLS. This includes the setting of the S bit. In each case, the S bit is set to indicate the bottom (i.e., innermost) label

in the label stack that is contiguous between the MPLS-TP LSP and its payload, and only one label stack entry (LSE) contains the S bit (Bottom of Stack bit) set to 1. Note that this best current practice differs slightly from [RFC3032], which uses the S bit to identify when MPLS label processing stops and network layer processing starts.

The relationship of MPLS-TP to its clients is illustrated in Figure 2. Note that the label stacks shown in the figure are divided between those inside the MPLS-TP network and those within the client network when the client network is MPLS(-TP). They illustrate the smallest number of labels possible. These label stacks could also include more labels.

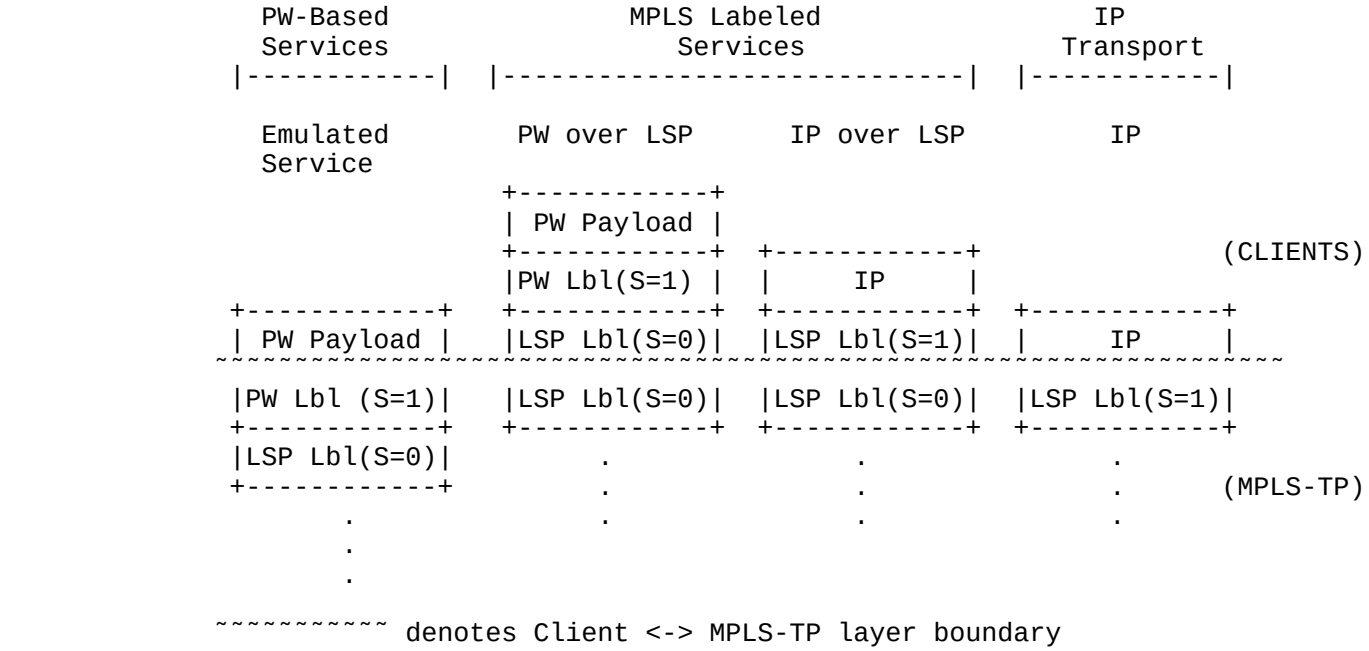

Figure 2: MPLS-TP - Client Relationship

3.4.2. MPLS-TP Transport Layers

An MPLS-TP network consists logically of two layers: the Transport Service layer and the Transport Path layer.

The Transport Service layer provides the interface between Customer Edge (CE) nodes and the MPLS-TP network. Each packet transmitted by a CE node for transport over the MPLS-TP network is associated at the receiving MPLS-TP Provider Edge (PE) node with a single logical point-to-point connection at the Transport Service layer between this

(ingress) PE and the corresponding (egress) PE to which the peer CE is attached. Such a connection is called an MPLS-TP Transport Service Instance, and the set of client packets belonging to the native service associated with such an instance on a particular CE-PE link is called a client flow.

The Transport Path layer provides aggregation of Transport Service Instances over MPLS-TP transport paths (LSPs), as well as aggregation of transport paths (via LSP hierarchy).

Awareness of the Transport Service layer need exist only at PE nodes. MPLS-TP Provider (P) nodes need have no awareness of this layer. Both PE and P nodes participate in the Transport Path layer. A PE terminates (i.e., is an LER with respect to) the transport paths it supports, and is responsible for multiplexing and demultiplexing of Transport Service Instance traffic over such transport paths.

3.4.3. MPLS-TP Transport Service Interfaces

An MPLS-TP PE node can provide two types of interface to the Transport Service layer. The MPLS-TP User-Network Interface (UNI) provides the interface between a CE and the MPLS-TP network. The MPLS-TP Network-Network Interface (NNI) provides the interface between two MPLS-TP PEs in different administrative domains.

When MPLS-TP is used to provide a transport service for, e.g., IP services that are a part of a Layer 3 VPN, then packets are transported in the same manner as specified in [RFC4364].

3.4.3.1. User-Network Interface

The MPLS-TP User-Network interface (UNI) is illustrated in Figure 3. The UNI for a particular client flow may or may not involve signaling between the CE and PE, and if signaling is used, it may or may not traverse the same attachment circuit that supports the client flow.

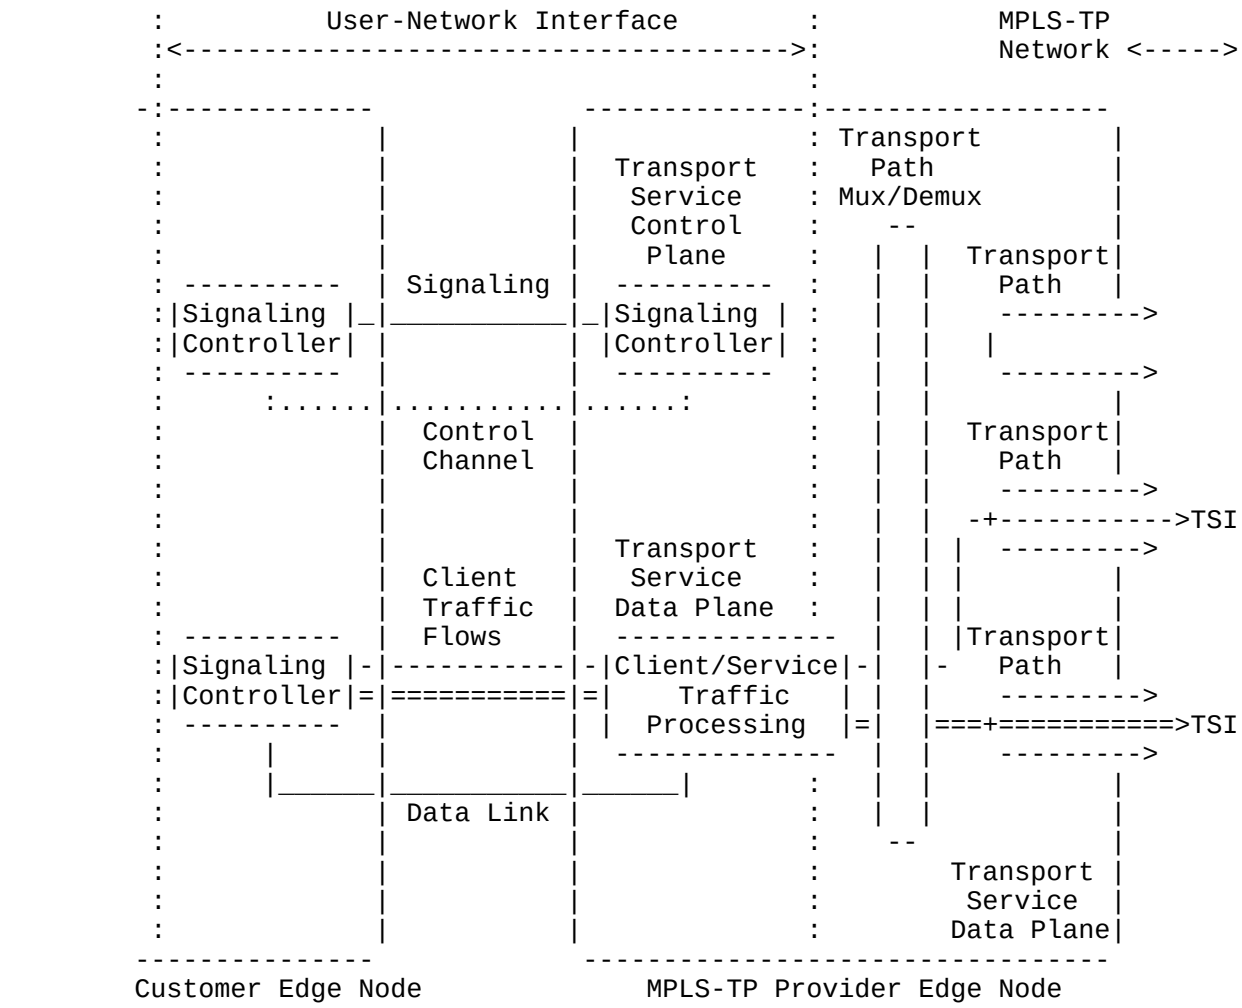

TSI = Transport Service Instance

Figure 3: MPLS-TP PE Containing a UNI

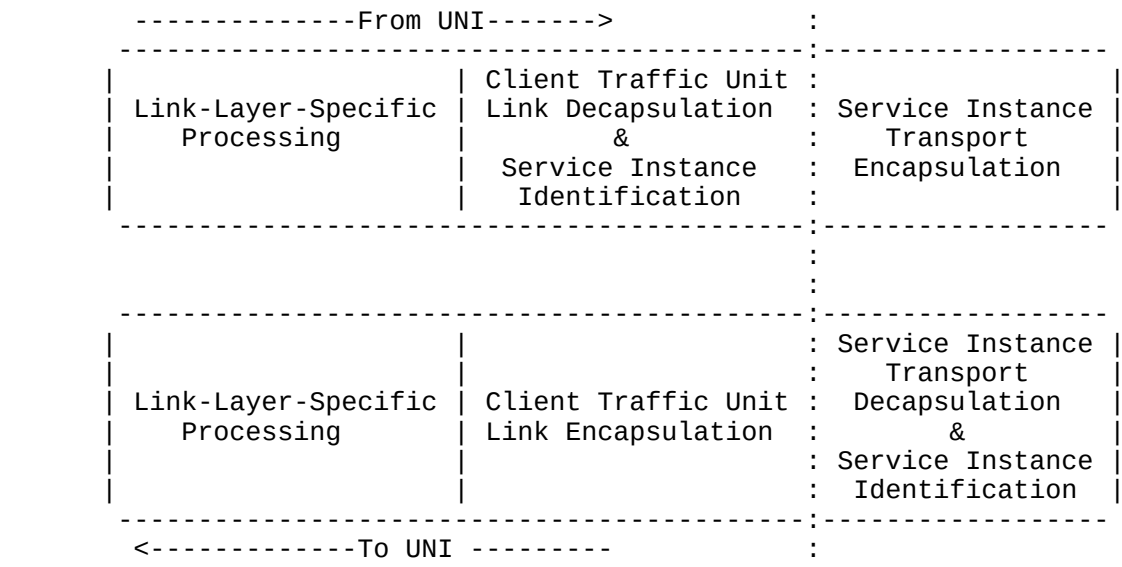

Figure 4: MPLS-TP UNI Client-Server Traffic Processing Stages

Figure 4 shows the logical processing steps involved in a PE both for traffic flowing from the CE to the MPLS-TP network (left to right), and from the network to the CE (right to left).

In the first case, when a packet from a client flow is received by the PE from the CE over the data-link, the following steps occur:

1. Link-layer-specific pre-processing, if any, is performed. An example of such pre-processing is the PREP function illustrated in Figure 3 of [RFC3985]. Such pre-processing is outside the scope of MPLS-TP.

2. The packet is extracted from the data-link frame, if necessary, and associated with a Transport Service Instance. At this point, UNI processing has completed.

3. A transport service encapsulation is associated with the packet, if necessary, for transport over the MPLS-TP network.

4. The packet is mapped to a transport path based on its associated Transport Service Instance, the transport path encapsulation is added, if necessary, and the packet is transmitted over the transport path.

In the second case, when a packet associated with a Transport Service Instance arrives over a transport path, the following steps occur:

1. The transport path encapsulation is disposed of.

2. The transport service encapsulation is disposed of and the Transport Service Instance and client flow identified.

3. At this point, UNI processing begins. A data-link encapsulation is associated with the packet for delivery to the CE based on the client flow.

4. Link-layer-specific postprocessing, if any, is performed. Such postprocessing is outside the scope of MPLS-TP.

3.4.3.2. Network-Network Interface

The MPLS-TP NNI is illustrated in Figure 5. The NNI for a particular Transport Service Instance may or may not involve signaling between the two PEs; and if signaling is used, it may or may not traverse the same data-link that supports the service instance.

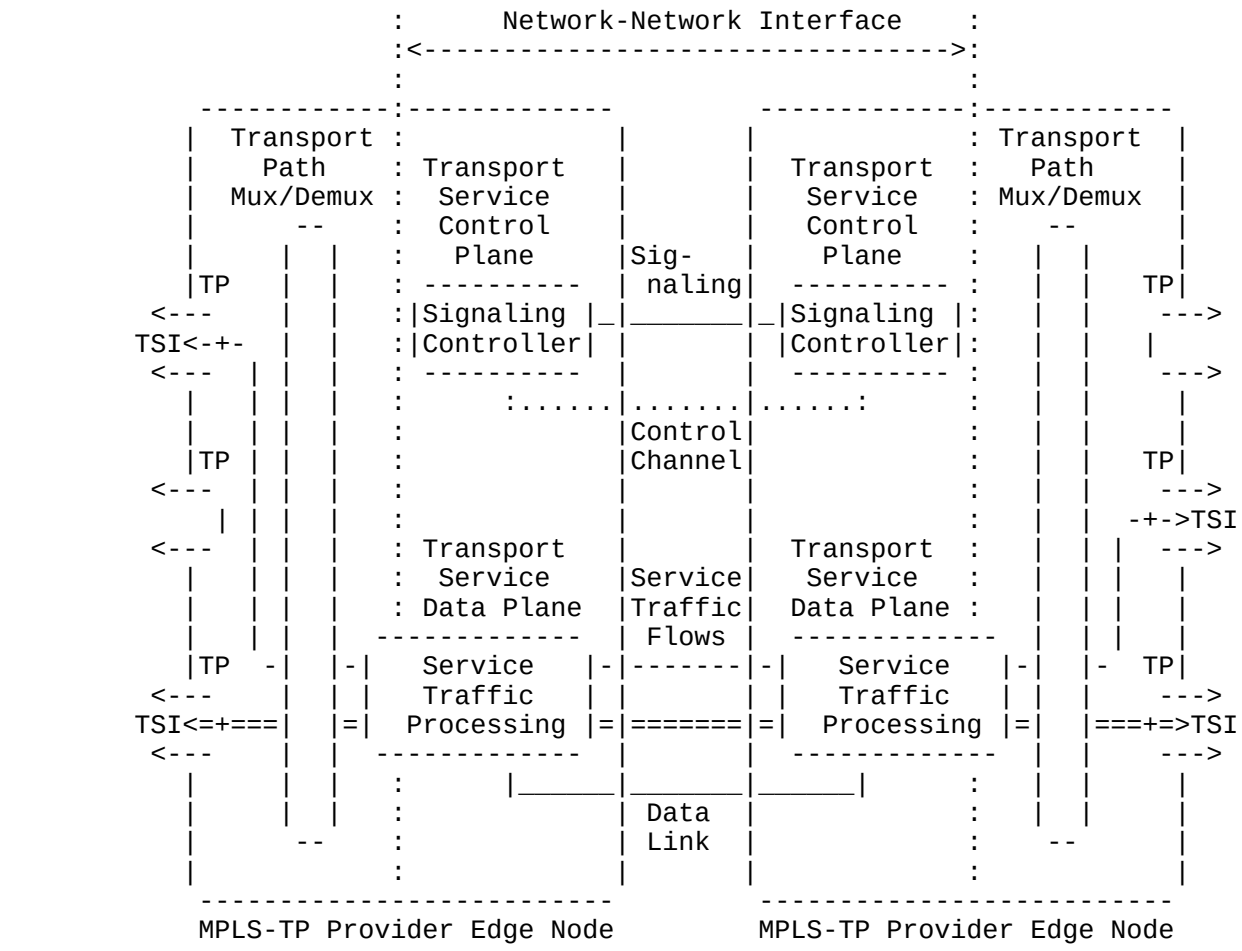

TP = Transport Path
TSI = Transport Service Instance

Figure 5: MPLS-TP PE Containing an NNI

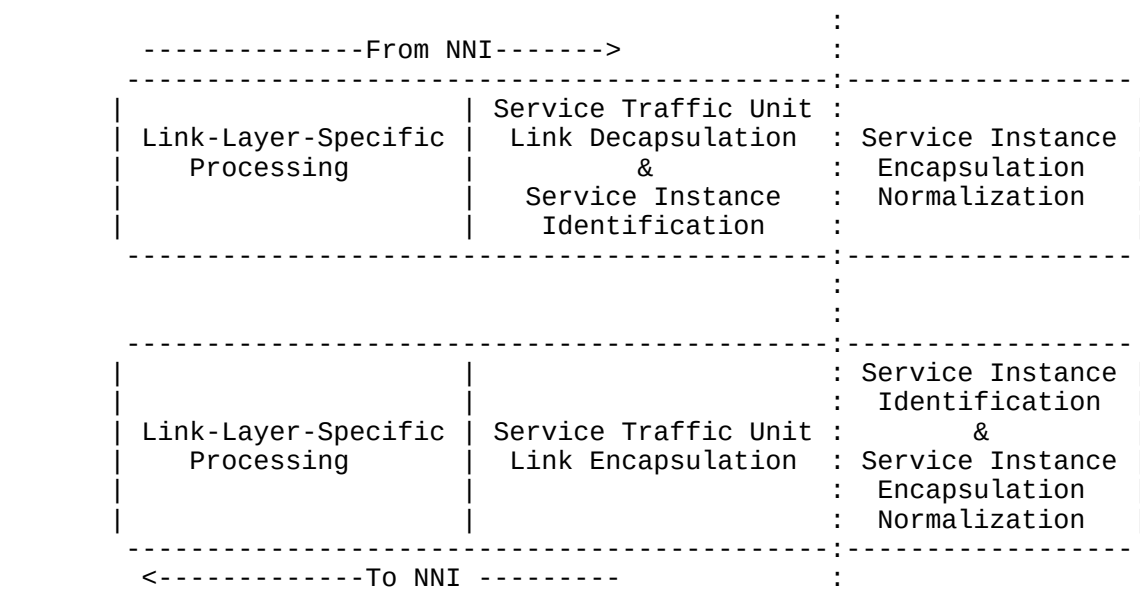

Figure 6: MPLS-TP NNI Service Traffic Processing Stages

Figure 6 shows the logical processing steps involved in a PE for traffic flowing both from the peer PE (left to right) and to the peer PE (right to left).

In the first case, when a packet from a Transport Service Instance is received by the PE from the peer PE over the data-link, the following steps occur:

1. Link-layer specific pre-processing, if any, is performed. Such pre-processing is outside the scope of MPLS-TP.

2. The packet is extracted from the data-link frame if necessary, and associated with a Transport Service Instance. At this point, NNI processing has completed.

3. The transport service encapsulation of the packet is normalized for transport over the MPLS-TP network. This step allows a different transport service encapsulation to be used over the NNI than that used in the internal MPLS-TP network. An example of such normalization is a swap of a label identifying the Transport Service Instance.

4. The packet is mapped to a transport path based on its associated Transport Service Instance, the transport path encapsulation is added, if necessary, and the packet is transmitted over the transport path.

In the second case, when a packet associated with a Transport Service Instance arrives over a transport path, the following steps occur:

1. The transport path encapsulation is disposed of.

2. The Transport Service Instance is identified from the transport service encapsulation, and this encapsulation is normalized for delivery over the NNI (see Step 3 above).

3. At this point, NNI processing begins. A data-link encapsulation is associated with the packet for delivery to the peer PE based on the normalized Transport Service Instance.

4. Link-layer-specific postprocessing, if any, is performed. Such postprocessing is outside the scope of MPLS-TP.

3.4.3.3. Example Interfaces

This section considers some special cases of UNI processing for particular transport service types. These are illustrative, and do not preclude other transport service types.

3.4.3.3.1. Layer 2 Transport Service

In this example the MPLS-TP network is providing a point-to-point Layer 2 transport service between attached CE nodes. This service is provided by a Transport Service Instance consisting of a PW established between the associated PE nodes. The client flows associated with this Transport Service Instance are the sets of all Layer 2 frames transmitted and received over the attachment circuits.

The processing steps in this case for a frame received from the CE are:

1. Link-layer specific pre-processing, if any, is performed, corresponding to the PREP function illustrated in Figure 3 of [RFC3985].

2. The frame is associated with a Transport Service Instance based on the attachment circuit over which it was received.

3. A transport service encapsulation, consisting of the PW control word and PW label, is associated with the frame.

4. The resulting packet is mapped to an LSP, the LSP label is pushed, and the packet is transmitted over the outbound interface associated with the LSP.

For PW packets received over the LSP, the steps are performed in the reverse order.

3.4.3.3.2. IP Transport Service

In this example, the MPLS-TP network is providing a point-to-point IP transport service between CE1, CE2, and CE3, as follows. One point-to-point Transport Service Instance delivers IPv4 packets between CE1 and CE2, and another instance delivers IPv6 packets between CE1 and CE3.

The processing steps in this case for an IP packet received from CE1 are:

1. No link-layer-specific processing is performed.

2. The IP packet is extracted from the link-layer frame and associated with a Service LSP based on the source MAC address (CE1) and the IP protocol version.

3. A transport service encapsulation, consisting of the Service LSP label, is associated with the packet.

4. The resulting packet is mapped to a tunnel LSP, the tunnel LSP label is pushed, and the packet is transmitted over the outbound interface associated with the LSP.

For packets received over a tunnel LSP carrying the Service LSP label, the steps are performed in the reverse order.

3.4.4. Pseudowire Adaptation

MPLS-TP uses pseudowires to provide a Virtual Private Wire Service (VPWS), a Virtual Private Local Area Network Service (VPLS), a Virtual Private Multicast Service (VPMS), and an Internet Protocol Local Area Network Service (IPLS). VPWS, VLPS, and IPLS are described in [RFC4664]. VPMS is described in [VPMS-REQS].

If the MPLS-TP network provides a layer 2 interface (that can carry both network-layer and non-network-layer traffic) as a service interface, then a PW is required to support the service interface. The PW is a client of the MPLS-TP LSP server layer. The architecture for an MPLS-TP network that provides such services is based on the MPLS [RFC3031] and pseudowire [RFC3985] architectures. Multi-segment

pseudowires may optionally be used to provide a packet transport service, and their use is consistent with the MPLS-TP architecture. The use of MS-PWs may be motivated by, for example, the requirements specified in [RFC5254]. If MS-PWs are used, then the MS-PW architecture [RFC5659] also applies.

Figure 7 shows the architecture for an MPLS-TP network using single-segment PWs. Note that, in this document, the client layer is equivalent to the emulated service described in [RFC3985], while the Transport LSP is equivalent to the Packet Switched Network (PSN) tunnel of [RFC3985].

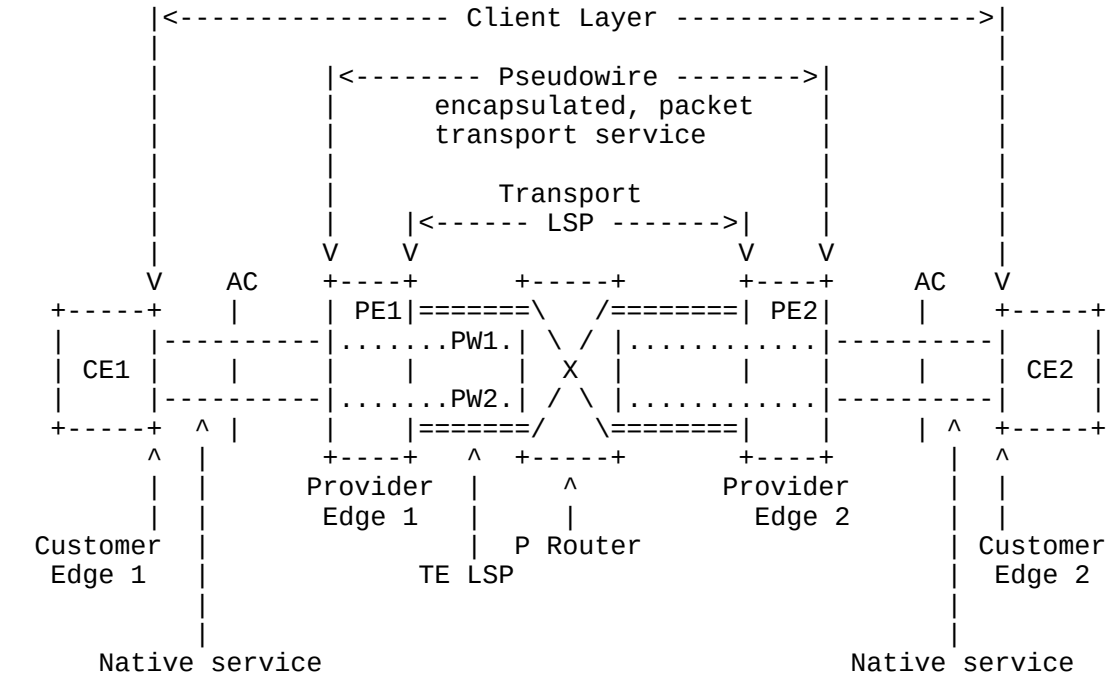

Figure 7: MPLS-TP Architecture (Single Segment PW)

Figure 8 shows the architecture for an MPLS-TP network when multi-segment pseudowires are used. Note that as in the SS-PW case, P-routers may also exist.

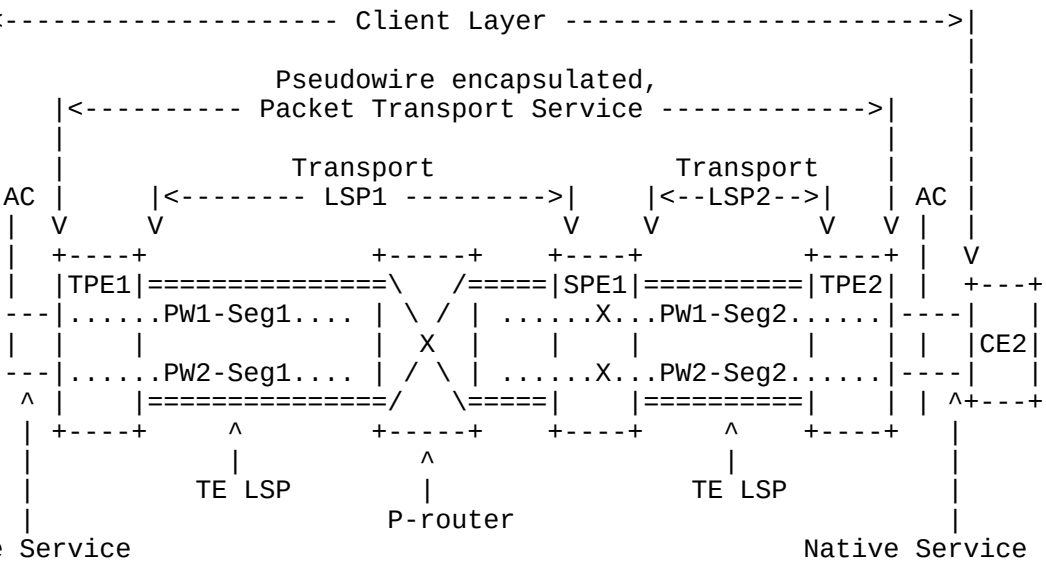

PW1-segment1 and PW1-segment2 are segments of the same MS-PW,
while PW2-segment1 and PW2-segment2 are segments of another MS-PW.

Figure 8: MPLS-TP Architecture (Multi-Segment PW)

The corresponding MPLS-TP protocol stacks including PWs are shown in Figure 9. In this figure, the Transport Service layer [RFC5654] is identified by the PW demultiplexer (Demux) label, and the Transport Path layer [RFC5654] is identified by the LSP Demux Label.

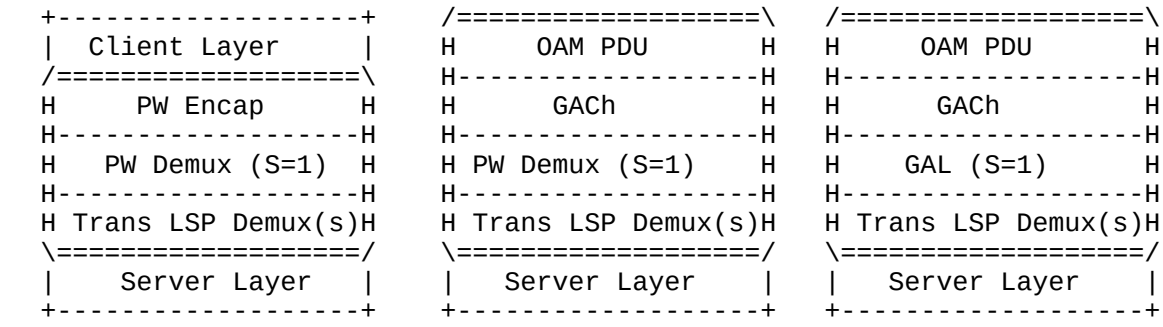

Note: H(ighlighted) indicates the part of the protocol stack considered in this document.

Figure 9: MPLS-TP Label Stack Using Pseudowires

PWs and their associated labels may be configured or signaled. See Section 3.11 for additional details related to configured service types. See Section 3.9 for additional details related to signaled service types.

3.4.5. Network Layer Adaptation

MPLS-TP LSPs can be used to transport network-layer clients. This document uses the term Network Layer in the same sense as it is used in [RFC3031] and [RFC3032]. The network-layer protocols supported by [RFC3031] and [RFC3032] can be transported between service interfaces. Support for network-layer clients follows the MPLS architecture for support of network-layer protocols as specified in [RFC3031] and [RFC3032].

With network-layer adaptation, the MPLS-TP domain provides either a unidirectional or bidirectional point-to-point connection between two PEs in order to deliver a packet transport service to attached customer edge (CE) nodes. For example, a CE may be an IP, MPLS, or MPLS-TP node. As shown in Figure 10, there is an attachment circuit between the CE node on the left and its corresponding provider edge (PE) node (which provides the service interface), a bidirectional LSP across the MPLS-TP network to the corresponding PE node on the right, and an attachment circuit between that PE node and the corresponding CE node for this service.

The attachment circuits may be heterogeneous (e.g., any combination of SDH, PPP, Frame Relay, etc.) and network-layer protocol payloads arrive at the service interface encapsulated in the Layer 1 / Layer 2

encoding defined for that access link type. It should be noted that the set of network-layer protocols includes MPLS, and hence MPLS-encoded packets with an MPLS label stack (the client MPLS stack) may appear at the service interface.

The following figures illustrate the reference models for network-layer adaptation. The details of these figures are described further in the following paragraphs.

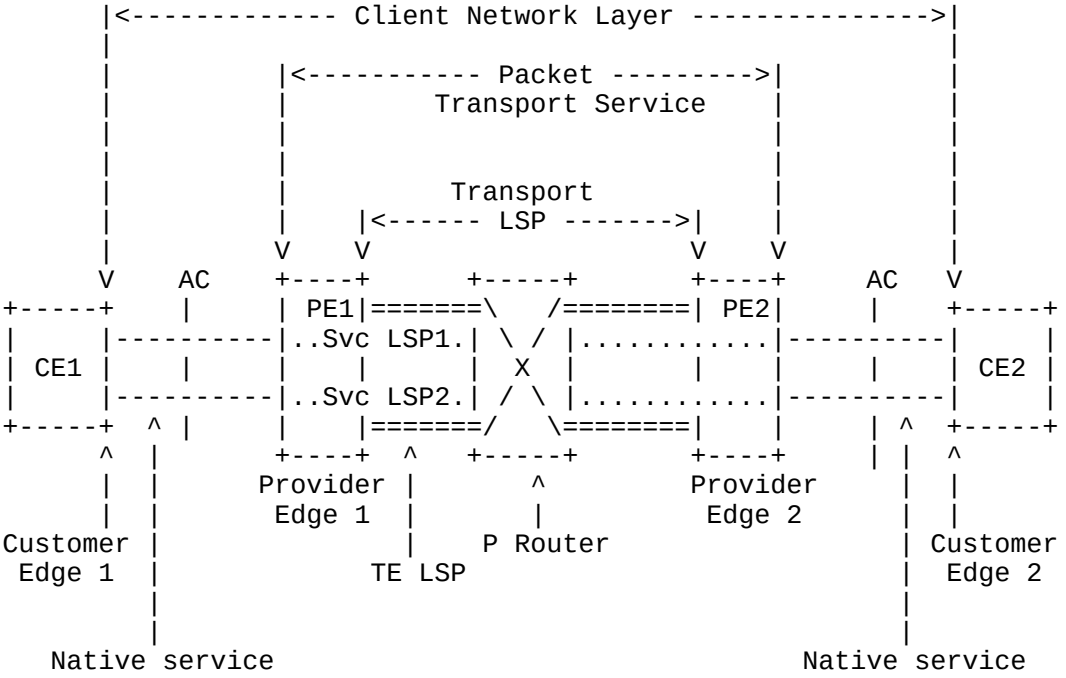

Figure 10: MPLS-TP Architecture for Network-Layer Clients

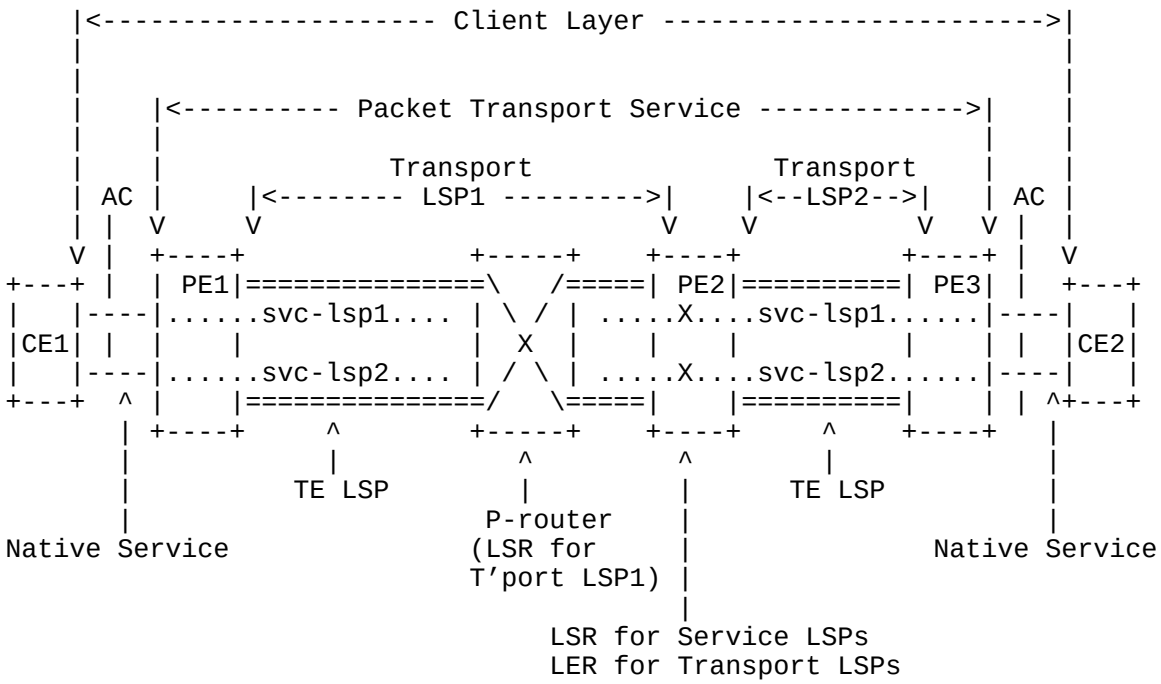

Figure 11: MPLS-TP Architecture for Network Layer Adaptation, Showing Service LSP Switching

Client packets are received at the ingress service interface. The PE pushes one or more labels onto the client packets that are then label switched over the transport network. Correspondingly, the egress PE pops any labels added by the MPLS-TP networks and transmits the packet for delivery to the attached CE via the egress service interface.

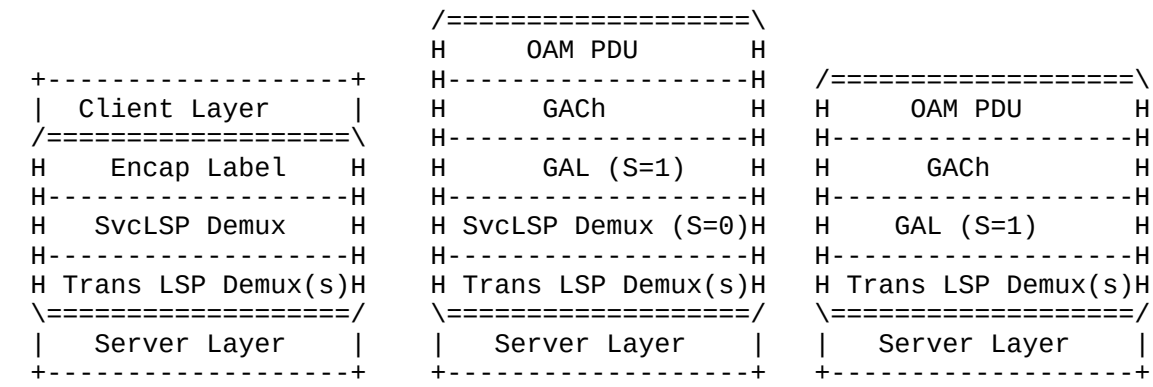

Note: H(ighlighted) indicates the part of the protocol stack considered in this document.

Figure 12: MPLS-TP Label Stack for IP and LSP Clients

In the figures above, the Transport Service layer [RFC5654] is identified by the Service LSP (SvcLSP) demultiplexer (Demux) label, and the Transport Path layer [RFC5654] is identified by the Transport (Trans) LSP Demux Label. Note that the functions of the Encapsulation Label (Encap Label) and the Service Label (SvcLSP Demux) shown above may alternatively be represented by a single label stack entry. Note that the S bit is always zero when the client layer is MPLS-labeled. It may be necessary to swap a service LSP label at an intermediate node. This is shown in Figure 11.

Within the MPLS-TP transport network, the network-layer protocols are carried over the MPLS-TP network using a logically separate MPLS label stack (the server stack). The server stack is entirely under the control of the nodes within the MPLS-TP transport network and it is not visible outside that network. Figure 12 shows how a client network protocol stack (which may be an MPLS label stack and payload) is carried over a network layer client service over an MPLS-TP transport network.

A label may be used to identify the network-layer protocol payload type. Therefore, when multiple protocol payload types are to be carried over a single service LSP, a unique label stack entry needs to be present for each payload type. Such labels are referred to as "Encapsulation Labels", one of which is shown in Figure 12. An Encapsulation Label may be either configured or signaled.

Both an Encapsulation Label and a Service Label should be present in the label stack when a particular packet transport service is supporting more than one network-layer protocol payload type. For example, if both IP and MPLS are to be carried, then two Encapsulation Labels are mapped on to a common Service Label.

Note: The Encapsulation Label may be omitted when the service LSP is supporting only one network-layer protocol payload type. For example, if only MPLS labeled packets are carried over a service, then the Service Label (stack entry) provides both the payload type indication and service identification. The Encapsulation Label cannot have any of the reserved label values [RFC3032].

Service labels are typically carried over an MPLS-TP Transport LSP edge-to-edge (or transport path layer). An MPLS-TP Transport LSP is represented as an LSP Transport Demux label, as shown in Figure 12. Transport LSP is commonly used when more than one service exists between two PEs.

Note that, if only one service exists between two PEs, the functions of the Transport LSP label and the Service LSP Label may be combined into a single label stack entry. For example, if only one service is carried between two PEs, then a single label could be used to provide both the service indication and the MPLS-TP Transport LSP. Alternatively, if multiple services exist between a pair of PEs, then a per-client Service Label would be mapped on to a common MPLS-TP Transport LSP.

As noted above, the Layer 2 and Layer 1 protocols used to carry the network-layer protocol over the attachment circuits are not transported across the MPLS-TP network. This enables the use of different Layer 2 and Layer 1 protocols on the two attachment circuits.

At each service interface, Layer 2 addressing needs to be used to ensure the proper delivery of a network-layer packet to the adjacent node. This is typically only an issue for LAN media technologies (e.g., Ethernet) that have Media Access Control (MAC) addresses. In cases where a MAC address is needed, the sending node sets the destination MAC address to an address that ensures delivery to the adjacent node. That is, the CE sets the destination MAC address to an address that ensures delivery to the PE, and the PE sets the destination MAC address to an address that ensures delivery to the CE. The specific address used is technology type specific and is not specified in this document. In some technologies, the MAC address will need to be configured.

Note that when two CEs, which peer with each other, operate over a network layer transport service and run a routing protocol such as IS-IS or OSPF, some care should be taken to configure the routing protocols to use point-to-point adjacencies. The specifics of such configuration is outside the scope of this document. See [RFC5309] for additional details.

The CE-to-CE service types and corresponding labels may be configured or signaled.

3.5. Identifiers

Identifiers are used to uniquely distinguish entities in an MPLS-TP network. These include operators, nodes, LSPs, pseudowires, and their associated maintenance entities. MPLS-TP defined two types of sets of identifiers: those that are compatible with IP, and those that are compatible with ITU-T transport-based operations. The definition of these sets of identifiers is outside the scope of this document and is provided by [IDENTIFIERS].

3.6. Generic Associated Channel (G-ACh)

For correct operation of OAM mechanisms, it is important that OAM packets fate-share with the data packets. In addition, in MPLS-TP it is necessary to discriminate between user data payloads and other types of payload. For example, a packet may be associated with a Signaling Communication Channel (SCC) or a channel used for a protocol to coordinate path protection state. This is achieved by carrying such packets in either:

- A generic control channel associated to the LSP, PW, or section, with no IP encapsulation, e.g., in a similar manner to Bidirectional Forwarding Detection for Virtual Circuit Connectivity Verification (VCCV-BFD) with PW ACH encapsulation [RFC5885]).

- An IP encapsulation where IP capabilities are present, e.g., PW ACH encapsulation with IP headers for VCCV-BFD [RFC5885] or IP encapsulation for MPLS BFD [RFC5884].

MPLS-TP makes use of such a generic associated channel (G-ACh) to support Fault, Configuration, Accounting, Performance, and Security (FCAPS) functions by carrying packets related to OAM, a protocol used to coordinate path protection state, SCC, MCC or other packet types in-band over LSPs, PWs, or sections. The G-ACh is defined in [RFC5586] and is similar to the Pseudowire Associated Channel [RFC4385], which is used to carry OAM packets over pseudowires. The G-ACh is indicated by an Associated Channel Header (ACH), similar to

the Pseudowire VCCV control word; this header is present for all sections, LSPs, and PWs that make use of FCAPS functions supported by the G-ACh.

As specified in [RFC5586], the G-ACh must only be used for channels that are an adjunct to the data service. Examples of these are OAM, a protocol used to coordinate path protection state, MCC, and SCC, but the use is not restricted to these services. The G-ACh must not be used to carry additional data for use in the forwarding path, i.e., it must not be used as an alternative to a PW control word, or to define a PW type.

At the server layer, bandwidth and QoS commitments apply to the gross traffic on the LSP, PW, or section. Since the G-ACh traffic is indistinguishable from the user data traffic, protocols using the G-ACh need to take into consideration the impact they have on the user data with which they are sharing resources. Conversely, capacity needs to be made available for important G-ACh uses such as protection and OAM. In addition, the security and congestion considerations described in [RFC5586] apply to protocols using the G-ACh.

Figure 13 shows the reference model depicting how the control channel is associated with the pseudowire protocol stack. This is based on the reference model for VCCV shown in Figure 2 of [RFC5085].

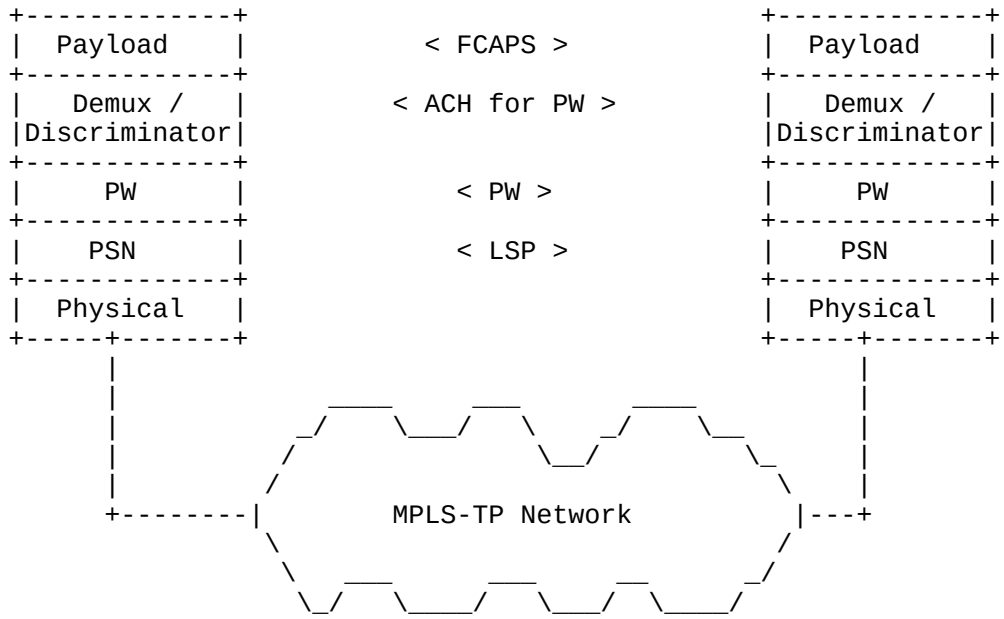

Figure 13: PWE3 Protocol Stack Reference Model Showing the G-ACh

PW-associated channel messages are encapsulated using the PWE3 encapsulation, so that they are handled and processed in the same manner (or in some cases, an analogous manner) as the PW PDUs for which they provide a control channel.

Figure 14 shows the reference model depicting how the control channel is associated with the LSP protocol stack.

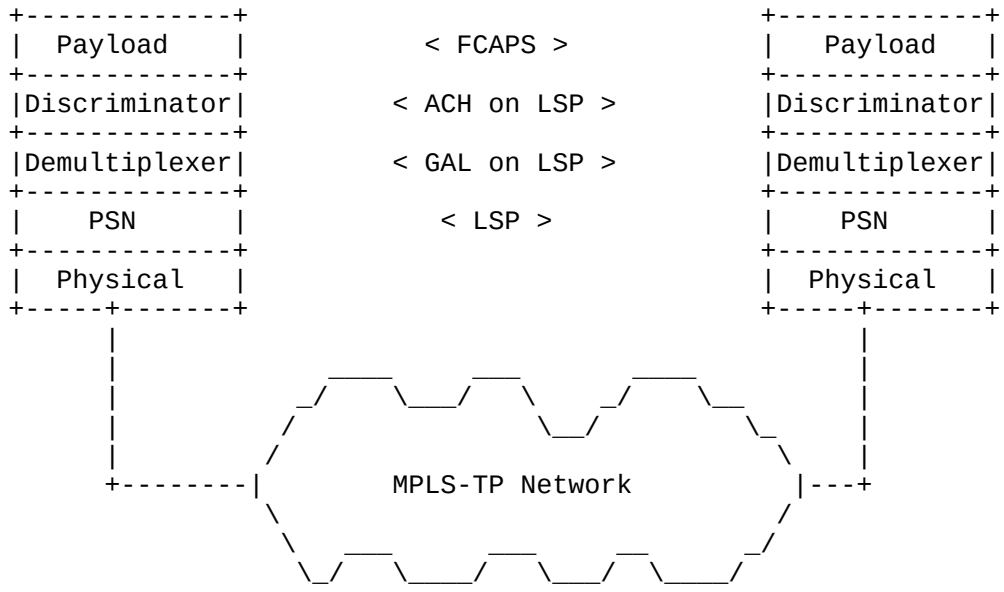

Figure 14: MPLS Protocol Stack Reference Model Showing the LSP Associated Control Channel

3.7. Operations, Administration, and Maintenance (OAM)

The MPLS-TP OAM architecture supports a wide range of OAM functions to check continuity, to verify connectivity, to monitor path performance, and to generate, filter, and manage local and remote defect alarms. These functions are applicable to any layer defined within MPLS-TP, i.e., to MPLS-TP sections, LSPs, and PWs.

The MPLS-TP OAM tool-set is able to operate without relying on a dynamic control plane or IP functionality in the data path. In the case of an MPLS-TP deployment in a network in which IP functionality is available, all existing IP/MPLS OAM functions (e.g., LSP Ping, BFD, and VCCV) may be used. Since MPLS-TP can operate in environments where IP is not used in the forwarding plane, the default mechanism for OAM demultiplexing in MPLS-TP LSPs and PWs is the Generic Associated Channel (Section 3.6). Forwarding based on IP addresses for OAM or user data packets is not required for MPLS-TP.

[RFC4379] and BFD for MPLS LSPs [RFC5884] have defined alert mechanisms that enable an MPLS LSR to identify and process MPLS OAM packets when the OAM packets are encapsulated in an IP header. These alert mechanisms are based on TTL expiration and/or use an IP destination address in the range 127/8 for IPv4 and that same range embedded as IPv4-mapped IPv6 addresses for IPv6 [RFC4379]. When the

OAM packets are encapsulated in an IP header, these mechanisms are the default mechanisms for MPLS networks (in general) for identifying MPLS OAM packets, although the mechanisms defined in [RFC5586] can also be used. MPLS-TP is able to operate in environments where IP forwarding is not supported, and thus the G-ACh/GAL is the default mechanism to demultiplex OAM packets in MPLS-TP in these environments.

MPLS-TP supports a comprehensive set of OAM capabilities for packet transport applications, with equivalent capabilities to those provided in SONET/SDH.

MPLS-TP requires [RFC5860] that a set of OAM capabilities is available to perform fault management (e.g., fault detection and localization) and performance monitoring (e.g., packet delay and loss measurement) of the LSP, PW, or section. The framework for OAM in MPLS-TP is specified in [OAM-FRAMEWORK].

MPLS-TP OAM packets share the same fate as their corresponding data packets, and are identified through the Generic Associated Channel mechanism [RFC5586]. This uses a combination of an Associated Channel Header (ACH) and a G-ACh Label (GAL) to create a control channel associated to an LSP, section, or PW.

OAM and monitoring in MPLS-TP is based on the concept of maintenance entities, as described in [OAM-FRAMEWORK]. A Maintenance Entity (ME) can be viewed as the association of two Maintenance Entity Group End Points (MEPs). A Maintenance Entity Group (MEG) is a collection of one or more MEs that belongs to the same transport path and that are maintained and monitored as a group. The MEPs that form an ME limit the OAM responsibilities of an OAM flow to within the domain of a transport path or segment, in the specific layer network that is being monitored and managed.

A MEG may also include a set of Maintenance Entity Group Intermediate Points (MIPs).

A G-ACh packet may be directed to an individual MIP along the path of an LSP or MS-PW by setting the appropriate TTL in the label stack entry for the G-ACh packet, as per the traceroute mode of LSP Ping [RFC4379] and the vccv-trace mode of [SEGMENTED-PW]. Note that this works when the location of MIPs along the LSP or PW path is known by the MEP. There may be circumstances where this is not the case, e.g., following restoration using a facility bypass LSP. In these cases, tools to trace the path of the LSP may be used to determine the appropriate setting for the TTL to reach a specific MIP.

Within an LSR or PE, MEPs and MIPs can only be placed where MPLS layer processing is performed on a packet. The MPLS architecture mandates that MPLS layer processing occurs at least once on an LSR.

Any node on an LSP can send an OAM packet on that LSP. Likewise, any node on a PW can send OAM packets on a PW, including S-PEs.

An OAM packet can only be received to be processed at an LSP endpoint, a PW endpoint (T-PE), or on the expiry of the TTL in the LSP or PW label stack entry.

3.8. Return Path

Management, control, and OAM protocol functions may require response packets to be delivered from the receiver back to the originator of a message exchange. This section provides a summary of the return path options in MPLS-TP networks. Although this section describes the case of an MPLS-TP LSP, it is also applicable to a PW.

In this description, U and D are LSRs that terminate MPLS-TP LSPs (i.e., LERs), and Y is an intermediate LSR along the LSP. Note that U is the upstream LER, and D is the downstream LER with respect to a particular direction of an LSP. This reference model is shown in Figure 15.

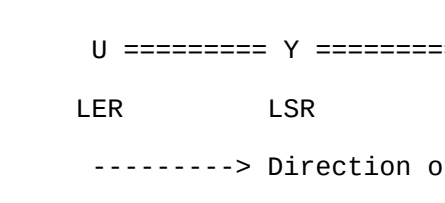

Figure 15: Return Path Reference Model

The following cases are described for the various types of LSPs:

Case 1 Return path packet transmission from D to U

Case 2 Return path packet transmission from Y to U

Case 3 Return path packet transmission from D to Y

Note that a return path may not always exist (or may exist but be disabled), and that packet transmission in one or more of the above cases may not be possible. In general, the existence and nature of return paths for MPLS-TP LSPs is determined by operational provisioning.

3.8.1. Return Path Types

There are two types of return path that may be used for the delivery of traffic from a downstream node D to an upstream node U. Either:

a. The LSP between U and D is bidirectional, and therefore D has a path via the MPLS-TP LSP to return traffic back to U, or

b. D has some other unspecified means of directing traffic back to U.

The first option is referred to as an "in-band" return path, the second as an "out-of-band" return path.

There are various possibilities for "out-of-band" return paths. Such a path may, for example, be based on ordinary IP routing. In this case, packets would be forwarded as usual to a destination IP address associated with U. In an MPLS-TP network that is also an IP/MPLS network, such a forwarding path may traverse the same physical links or logical transport paths used by MPLS-TP. An out-of-band return path may also be indirect, via a distinct Data Communication Network (DCN) (provided, for example, by the method specified in [RFC5718]); or it may be via one or more other MPLS-TP LSPs.

3.8.2. Point-to-Point Unidirectional LSPs

Case 1 If an in-band return path is required to deliver traffic from D back to U, it is recommended for reasons of operational simplicity that point-to-point unidirectional LSPs be provisioned as associated bidirectional LSPs (which may also be co-routed) whenever return traffic from D to U is required. Note that the two directions of such an LSP may have differing bandwidth allocations and QoS characteristics. The discussion below for such LSPs applies.

As an alternative, an out-of-band return path may be used.

Case 2 In this case, only the out-of-band return path option is available. However, an additional out-of-band possibility is worthy of note here: if D is known to have a return path to U, then Y can arrange to deliver return traffic to U by first sending it to D along the original LSP. The mechanism by which D recognizes the need for and performs this forwarding operation is protocol specific.

Case 3 In this case, only the out-of-band return path option is available. However, if D has a return path to U, then (in a manner analogous to the previous case) D can arrange to

deliver return traffic to Y by first sending it to U along that return path. The mechanism by which U recognizes the need for and performs this forwarding operation is protocol specific.

3.8.3. Point-to-Point Associated Bidirectional LSPs

For Case 1, D has a natural in-band return path to U, the use of which is typically preferred for return traffic, although out-of-band return paths are also applicable.

For Cases 2 and 3, the considerations are the same as those for point-to-point unidirectional LSPs.

3.8.4. Point-to-Point Co-Routed Bidirectional LSPs

For all of Cases 1, 2, and 3, a natural in-band return path exists in the form of the LSP itself, and its use is preferred for return traffic. Out-of-band return paths, however, are also applicable, primarily as an alternative means of delivery in case the in-band return path has failed.

3.9. Control Plane

A distributed dynamic control plane may be used to enable dynamic service provisioning in an MPLS-TP network. Where the requirements specified in [RFC5654] can be met, the MPLS Transport Profile uses existing standard control-plane protocols for LSPs and PWs.

Note that a dynamic control plane is not required in an MPLS-TP network. See Section 3.11 for further details on statically configured and provisioned MPLS-TP services.

Figure 16 illustrates the relationship between the MPLS-TP control plane, the forwarding plane, the management plane, and OAM for point-to-point MPLS-TP LSPs or PWs.

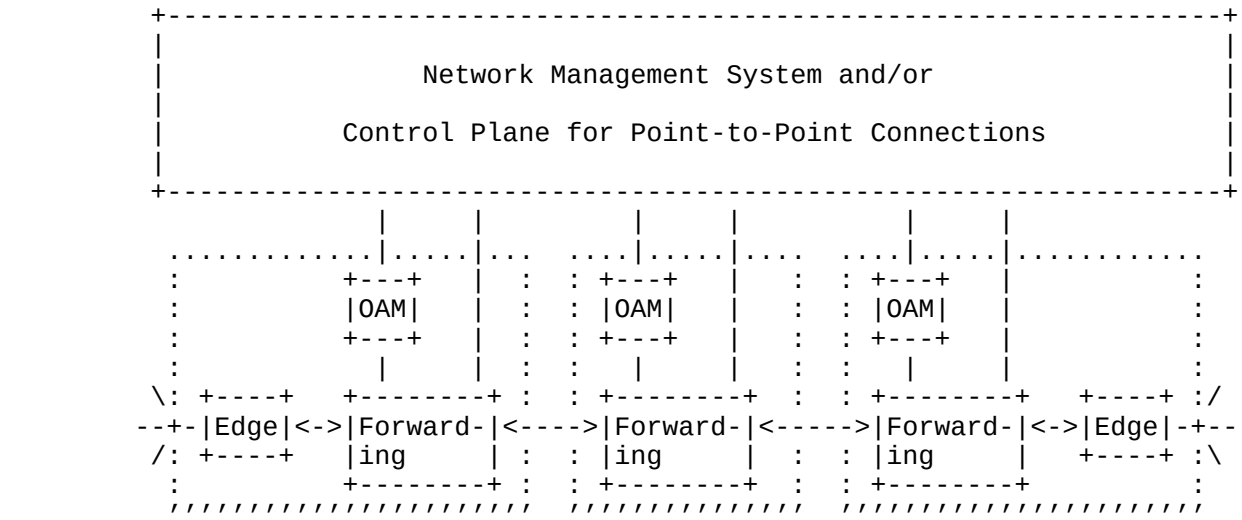

Note:
1) NMS may be centralized or distributed. Control plane is distributed.
2) 'Edge' functions refers to those functions present at the edge of a PSN domain, e.g., native service processing or classification.
3) The control plane may be transported over the server layer, an LSP, or a G-ACh.

Figure 16: MPLS-TP Control Plane Architecture Context

The MPLS-TP control plane is based on existing MPLS and PW control plane protocols, and is consistent with the Automatically Switched Optical Network (ASON) architecture [G.8080]. MPLS-TP uses:

- Generalized MPLS (GMPLS) signaling ([RFC3945], [RFC3471], [RFC3473]) for LSPs, and

- Targeted LDP (T-LDP) signaling ([RFC4447], [SEGMENTED-PW], [DYN-MS-PW]) for pseudowires.

MPLS-TP requires that any control-plane traffic be capable of being carried over an out-of-band signaling network or a signaling control channel such as the one described in [RFC5718]. Note that while T-LDP signaling is traditionally carried in-band in IP/MPLS networks, this does not preclude its operation over out-of-band channels. References to T-LDP in this document do not preclude the definition of alternative PW control protocols for use in MPLS-TP.

PW control (and maintenance) takes place separately from LSP tunnel signaling. The main coordination between LSP and PW control will occur within the nodes that terminate PWs. The control planes for PWs and LSPs may be used independently, and one may be employed without the other. This translates into the four possible scenarios: (1) no control plane is employed; (2) a control plane is used for both LSPs and PWs; (3) a control plane is used for LSPs, but not PWs; (4) a control plane is used for PWs, but not LSPs. The PW and LSP control planes, collectively, need to satisfy the MPLS-TP control plane requirements reviewed in the MPLS-TP Control Plane Framework [CP-FRAMEWORK]. When client services are provided directly via LSPs, all requirements must be satisfied by the LSP control plane. When client services are provided via PWs, the PW and LSP control planes operate in combination, and some functions may be satisfied via the PW control plane, while others are provided to PWs by the LSP control plane.

Note that if MPLS-TP is being used in a multi-layer network, a number of control protocol types and instances may be used. This is consistent with the MPLS architecture, which permits each label in the label stack to be allocated and signaled by its own control protocol.

The distributed MPLS-TP control plane may provide the following functions:

- Signaling
- Routing
- Traffic engineering and constraint-based path computation

In a multi-domain environment, the MPLS-TP control plane supports different types of interfaces at domain boundaries or within the domains. These include the User-Network Interface (UNI), Internal Network-Network Interface (I-NNI), and External Network-Network Interface (E-NNI). Note that different policies may be defined that control the information exchanged across these interface types.

The MPLS-TP control plane is capable of activating MPLS-TP OAM functions as described in the OAM section of this document Section 3.7, e.g., for fault detection and localization in the event of a failure in order to efficiently restore failed transport paths.

The MPLS-TP control plane supports all MPLS-TP data-plane connectivity patterns that are needed for establishing transport paths, including protected paths as described in Section 3.12.

Examples of the MPLS-TP data-plane connectivity patterns are LSPs utilizing the fast reroute backup methods as defined in [RFC4090] and ingress-to-egress 1+1 or 1:1 protected LSPs.

The MPLS-TP control plane provides functions to ensure its own survivability and to enable it to recover gracefully from failures and degradations. These include graceful restart and hot redundant configurations. Depending on how the control plane is transported, varying degrees of decoupling between the control plane and data plane may be achieved. In all cases, however, the control plane is logically decoupled from the data plane such that a control-plane failure does not imply a failure of the existing transport paths.

3.10. Inter-Domain Connectivity

A number of methods exist to support inter-domain operation of MPLS-TP, including the data-plane, OAM, and configuration aspects, for example:

- Inter-domain TE LSPs [RFC4726]
- Multi-segment Pseudowires [RFC5659]
- LSP stitching [RFC5150]
- Back-to-back attachment circuits [RFC5659]

An important consideration in selecting an inter-domain connectivity mechanism is the degree of layer network isolation and types of OAM required by the operator. The selection of which technique to use in a particular deployment scenario is outside the scope of this document.

3.11. Static Operation of LSPs and PWs

A PW or LSP may be statically configured without the support of a dynamic control plane. This may be either by direct configuration of the PEs/LSRs or via a network management system. Static operation is independent for a specific PW or LSP instance. Thus, it should be possible for a PW to be statically configured, while the LSP supporting it is set up by a dynamic control plane. When static configuration mechanisms are used, care must be taken to ensure that loops are not created. Note that the path of an LSP or PW may be dynamically computed, while the LSP or PW itself is established through static configuration.

3.12. Survivability

The survivability architecture for MPLS-TP is specified in [SURVIVE-FWK].

A wide variety of resiliency schemes have been developed to meet the various network and service survivability objectives. For example, as part of the MPLS/PW paradigms, MPLS provides methods for local repair using back-up LSP tunnels ([RFC4090]), while pseudowire redundancy [PW-REDUNDANCY] supports scenarios where the protection for the PW cannot be fully provided by the underlying LSP (i.e., where the backup PW terminates on a different target PE node than the working PW in dual-homing scenarios, or where protection of the S-PE is required). Additionally, GMPLS provides a well-known set of control-plane-driven protection and restoration mechanisms [RFC4872]. MPLS-TP provides additional protection mechanisms that are optimized for both linear topologies and ring topologies and that operate in the absence of a dynamic control plane. These are specified in [SURVIVE-FWK].

Different protection schemes apply to different deployment topologies and operational considerations. Such protection schemes may provide different levels of resiliency, for example:

- two concurrent traffic paths (1+1).
- one active and one standby path with guaranteed bandwidth on both paths (1:1).
- one active path and a standby path the resources of which are shared by one or more other active paths (shared protection).

The applicability of any given scheme to meet specific requirements is outside the scope of this document.

The characteristics of MPLS-TP resiliency mechanisms are as follows:

- Optimized for linear, ring, or meshed topologies.
- Use OAM mechanisms to detect and localize network faults or service degenerations.
- Include protection mechanisms to coordinate and trigger protection switching actions in the absence of a dynamic control plane.
- MPLS-TP recovery schemes are applicable to all levels in the MPLS-TP domain (i.e., section, LSP, and PW) providing segment and end-to-end recovery.

- o MPLS-TP recovery mechanisms support the coordination of protection switching at multiple levels to prevent race conditions occurring between a client and its server layer.

- o MPLS-TP recovery mechanisms can be data-plane, control-plane, or management-plane based.

- o MPLS-TP supports revertive and non-revertive behavior.

3.13. Sub-Path Maintenance

In order to monitor, protect, and manage a portion (i.e., segment or concatenated segment) of an LSP, a hierarchical LSP [RFC3031] can be instantiated. A hierarchical LSP instantiated for this purpose is called a Sub-Path Maintenance Element (SPME). Note that by definition an SPME does not carry user traffic as a direct client.

An SPME is defined between the edges of the portion of the LSP that needs to be monitored, protected or managed. The SPME forms an MPLS-TP Section [DATA-PLANE] that carries the original LSP over this portion of the network as a client. OAM messages can be initiated at the edge of the SPME and sent to the peer edge of the SPME or to a MIP along the SPME by setting the TTL value of the LSE at the corresponding hierarchical LSP level. A P router only pushes or pops a label if it is at the end of a SPME. In this mode, it is an LER for the SPME.

For example, in Figure 17, two SPMEs are configured to allow monitoring, protection, and management of the LSP concatenated segments. One SPME is defined between LER2 and LER3, and a second SPME is set up between LER4 and LER5. Each of these SPMEs may be monitored, protected, or managed independently.

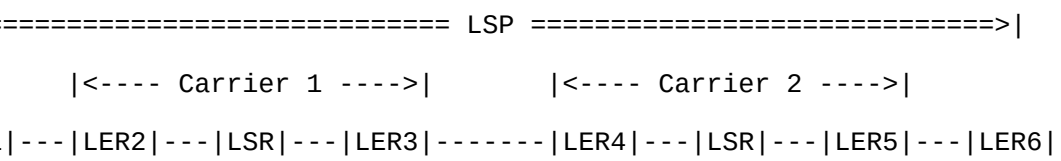

Note 1: LER2, LER3, LER4, and LER5 are with respect to the SPME, but LSRs with respect to the LSP.
Note 2: The LSP terminates in LERs outside of Carrier 1 and Carrier 2, for example, LER1 and LER6.

Figure 17: SPMEs in Inter-Carrier Network

The end-to-end traffic of the LSP, including data traffic and control traffic (OAM, Protection Switching Control, management and signaling messages) is tunneled within the hierarchical LSP by means of label stacking as defined in [RFC3031].

The mapping between an LSP and a SPME can be 1:1, in which case it is similar to the ITU-T Tandem Connection Element [G.805]. The mapping can also be 1:N to allow aggregated monitoring, protection, and management of a set of LSP segments or concatenated LSP segments. Figure 18 shows a SPME that is used to aggregate a set of concatenated LSP segments for the LSP from LERx to LERt and the LSP from LERa to LERd. Note that such a construct is useful, for example, when the LSPs traverse a common portion of the network and they have the same Traffic Class.

The QoS aspects of a SPME are network specific. [OAM-FRAMEWORK] provides further considerations on the QoS aspects of OAM.

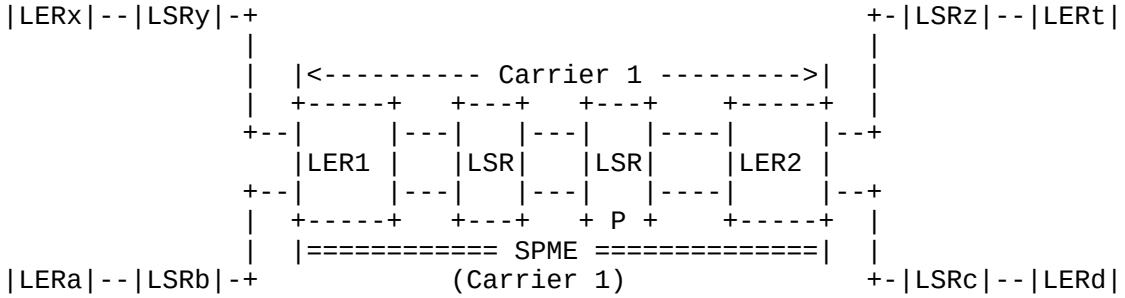

Figure 18: SPME for a Set of Concatenated LSP Segments

SPMEs can be provisioned either statically or using control-plane signaling procedures. The make-before-break procedures which are supported by MPLS allow the creation of a SPME on existing LSPs in-service without traffic disruption, as described in [SURVIVE-FWK]. A SPME can be defined corresponding to one or more end-to-end LSPs. New end-to-end LSPs that are tunneled within the SPME can be set up, which may require coordination across administrative boundaries. Traffic of the existing LSPs is switched over to the new end-to-end tunneled LSPs. The old end-to-end LSPs can then be torn down.

Hierarchical label stacking, in a similar manner to that described above, can be used to implement Sub-Path Maintenance Elements on pseudowires, as described in [OAM-FRAMEWORK].

3.14. Network Management

The network management architecture and requirements for MPLS-TP are specified in [NM-FRAMEWORK] and [NM-REQ]. These derive from the generic specifications described in ITU-T G.7710/Y.1701 [G.7710] for transport technologies. They also incorporate the OAM requirements for MPLS Networks [RFC4377] and MPLS-TP Networks [RFC5860] and expand on those requirements to cover the modifications necessary for fault, configuration, performance, and security in a transport network.

The Equipment Management Function (EMF) of an MPLS-TP Network Element (NE) (i.e., LSR, LER, PE, S-PE, or T-PE) provides the means through which a management system manages the NE. The Management Communication Channel (MCC), realized by the G-ACh, provides a logical operations channel between NEs for transferring management information. The Network Management System (NMS) can be used to provision and manage an end-to-end connection across a network. Maintenance operations are run on a connection (LSP or PW) in a manner that is independent of the provisioning mechanism. Segments may be created or managed by, for example, Netconf [RFC4741], SNMP [RFC3411], or CORBA (Common Object Request Broker Architecture) interfaces, but not all segments need to be created or managed using the same type of interface. Where an MPLS-TP NE is managed by an NMS, at least one of these standard management mechanisms is required for interoperability, but this document imposes no restriction on which of these standard management protocols is used. In MPLS-TP, the EMF needs to support statically provisioning LSPs for an LSR or LER, and PWs for a PE, as well as any associated MEPs and MIPs, as per Section 3.11.

Fault Management (FM) functions within the EMF of an MPLS-TP NE enable the supervision, detection, validation, isolation, correction, and alarm handling of abnormal conditions in the MPLS-TP network and its environment. FM needs to provide for the supervision of transmission (such as continuity, connectivity, etc.), software processing, hardware, and environment. Alarm handling includes alarm severity assignment, alarm suppression/aggregation/correlation, alarm reporting control, and alarm reporting.

Configuration Management (CM) provides functions to control, identify, collect data from, and provide data to MPLS-TP NEs. In addition to general configuration for hardware, software protection switching, alarm reporting control, and date/time setting, the EMF of the MPLS-TP NE also supports the configuration of maintenance entity identifiers (such as Maintenance Entity Group Endpoint (MEP) ID and MEG Intermediate Point (MIP) ID). The EMF also supports the configuration of OAM parameters as a part of connectivity management to meet specific operational requirements. These may specify whether

the operational mode is one-time on-demand or is periodic at a specified frequency.

The Performance Management (PM) functions within the EMF of an MPLS-TP NE support the evaluation and reporting of the behavior of the NEs and the network. One particular requirement for PM is to provide coherent and consistent interpretation of the network behavior in a hybrid network that uses multiple transport technologies. Packet loss measurement and delay measurements may be collected and used to detect performance degradation. This is reported via fault management to enable corrective actions to be taken (e.g., protection switching) and via performance monitoring for Service Level Agreement (SLA) verification and billing. Collection mechanisms for performance data should be capable of operating on-demand or proactively.

4. Security Considerations

The introduction of MPLS-TP into transport networks means that the security considerations applicable to both MPLS [RFC3031] and PWE3 [RFC3985] apply to those transport networks. When an MPLS function is included in the MPLS transport profile, the security considerations pertinent to that function apply to MPLS-TP. Furthermore, when general MPLS networks that utilize functionality outside of the strict MPLS Transport Profile are used to support packet transport services, the security considerations of that additional functionality also apply.

For pseudowires, the security considerations of [RFC3985] and [RFC5659] apply.

MPLS-TP nodes that implement the G-ACh create a Control Channel (CC) associated with a pseudowire, LSP, or section. This control channel can be signaled or statically configured. Over this control channel, control channel messages related to network maintenance functions such as OAM, signaling, or network management are sent. Therefore, three different areas are of concern from a security standpoint.

The first area of concern relates to control plane parameter and status message attacks, that is, attacks that concern the signaling of G-ACh capabilities. MPLS-TP Control Plane security is discussed in [RFC5920].

A second area of concern centers on data-plane attacks, that is, attacks on the G-ACh itself. MPLS-TP nodes that implement the G-ACh mechanisms are subject to additional data-plane denial-of-service attacks as follows:

An intruder could intercept or inject G-ACh packets effectively disrupting the protocols carried over the G-ACh.

An intruder could deliberately flood a peer MPLS-TP node with G-ACh messages to deny services to others.

A misconfigured or misbehaving device could inadvertently flood a peer MPLS-TP node with G-ACh messages that could result in denial of services. In particular, if a node has either implicitly or explicitly indicated that it cannot support one or all of the types of G-ACh protocol, but is sent those messages in sufficient quantity, it could result in a denial of service.

To protect against these potential (deliberate or unintentional) attacks, multiple mitigation techniques can be employed:

G-ACh message throttling mechanisms can be used, especially in distributed implementations that have a centralized control-plane processor with various line cards attached by some control-plane data path. In these architectures, G-ACh messages may be processed on the central processor after being forwarded there by the receiving line card. In this case, the path between the line card and the control processor may become saturated if appropriate G-ACh traffic throttling is not employed, which could lead to a complete denial of service to users of the particular line card. Such filtering is also useful for preventing the processing of unwanted G-ACh messages, such as those which are sent on unwanted (and perhaps unadvertised) control channel types.

A third and last area of concern relates to the processing of the actual contents of G-ACh messages. It is necessary that the definition of the protocols using these messages carried over a G-ACh include appropriate security measures.

Additional security considerations apply to each MPLS-TP solution. These are discussed further in [SEC-FRAMEWORK].

The security considerations in [RFC5920] apply.

5. IANA Considerations

IANA considerations resulting from specific elements of MPLS-TP functionality will be detailed in the documents specifying that functionality.

This document introduces no additional IANA considerations in itself.

6. Acknowledgements

The editors wish to thank the following for their contributions to this document:

- Rahul Aggarwal
- Dieter Beller
- Malcolm Betts
- Italo Busi
- John E Drake
- Hing-Kam Lam
- Marc Lasserre
- Vincenzo Sestito
- Nurit Sprecher
- Martin Vigoureux
- Yaacov Weingarten
- The participants of ITU-T SG15

7. References

7.1. Normative References

[G.7710] ITU-T, "Common equipment management function requirements", ITU-T Recommendation G.7710/Y.1701, July 2007.

[G.805] ITU-T, "Generic Functional Architecture of Transport Networks", ITU-T Recommendation G.805, November 1995.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.

[RFC3270] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, May 2002.

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

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

[RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.

[RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, February 2006.

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

[RFC4872] Lang, J., Rekhter, Y., and D. Papadimitriou, "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, May 2007.

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

[RFC5586] Bocci, M., Vigoureux, M., and S. Bryant, "MPLS Generic Associated Channel", RFC 5586, June 2009.

7.2. Informative References

[CP-FRAMEWORK] Andersson, L., Berger, L., Fang, L., Bitar, N., Takacs, A., Vigoureux, M., Bellagamba, E., and E. Gray, "MPLS-TP Control Plane Framework", Work in Progress, March 2010.

[DATA-PLANE] Frost, D., Bryant, S., and M. Bocci, "MPLS Transport Profile Data Plane Architecture", Work in Progress, July 2010.

[DYN-MS-PW] Martini, L., Bocci, M., Balus, F., Bitar, N., Shah, H., Aissaoui, M., Rusmisel, J., Serbest, Y., Malis, A., Metz, C., McDysan, D., Sugimoto, J., Duckett, M., Loomis, M., Doolan, P., Pan, P., Pate, P., Radoaca, V., Wada, Y., and Y. Seo, "Dynamic Placement of Multi Segment Pseudo Wires", Work in Progress, October 2009.

[G.8080] ITU-T, "Architecture for the automatically switched optical network (ASON)", ITU-T Recommendation G.8080/Y.1304, 2005.

[IDENTIFIERS] Bocci, M. and G. Swallow, "MPLS-TP Identifiers", Work in Progress, March 2010.

[NM-FRAMEWORK] Mansfield, S., Ed., Gray, E., Ed., and H. Lam, Ed., "MPLS-TP Network Management Framework", Work in Progress, February 2010.

[NM-REQ] Mansfield, S. and K. Lam, "MPLS TP Network Management Requirements", Work in Progress, October 2009.

[OAM-DEF] Andersson, L., Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "The OAM Acronym Soup", Work in Progress, June 2010.

[OAM-FRAMEWORK] Busi, I., Ed., Niven-Jenkins, B., Ed., and D. Allan, Ed., "MPLS-TP OAM Framework", Work in Progress, April 2010.

[PW-REDUNDANCY] Muley, P., "Pseudowire (PW) Redundancy", Work in Progress, May 2010.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[RFC3411] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, December 2002.

[RFC3443] Agarwal, P. and B. Akyol, "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", RFC 3443, January 2003.

[RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.

[RFC4377] Nadeau, T., Morrow, M., Swallow, G., Allan, D., and S. Matsushima, "Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks", RFC 4377, February 2006.

[RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.

[RFC4664] Andersson, L. and E. Rosen, "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, September 2006.

[RFC4726] Farrel, A., Vasseur, J., and A. Ayyangar, "A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering", RFC 4726, November 2006.

[RFC4741] Enns, R., "NETCONF Configuration Protocol", RFC 4741, December 2006.

[RFC5150] Ayyangar, A., Kompella, K., Vasseur, JP., and A. Farrel, "Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)", RFC 5150, February 2008.

[RFC5254] Bitar, N., Bocci, M., and L. Martini, "Requirements for Multi-Segment Pseudowire Emulation Edge-to-Edge (PWE3)", RFC 5254, October 2008.

[RFC5309] Shen, N. and A. Zinin, "Point-to-Point Operation over LAN in Link State Routing Protocols", RFC 5309, October 2008.

[RFC5331] Aggarwal, R., Rekhter, Y., and E. Rosen, "MPLS Upstream Label Assignment and Context-Specific Label Space", RFC 5331, August 2008.

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

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

[RFC5718] Beller, D. and A. Farrel, "An In-Band Data Communication Network For the MPLS Transport Profile", RFC 5718, January 2010.

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

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

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

[RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

[ROSETTA-STONE] Sprecher, N., "A Thesaurus for the Terminology used in Multiprotocol Label Switching Transport Profile (MPLS-TP) drafts/RFCs and ITU-T's Transport Network Recommendations.", Work in Progress, May 2010.

[SEC-FRAMEWORK] Fang, L. and B. Niven-Jenkins, "Security Framework for MPLS-TP", Work in Progress, March 2010.

[SEGMENTED-PW] Martini, L., Nadeau, T., Metz, C., Bocci, M., and M. Aissaoui, "Segmented Pseudowire", Work in Progress, June 2010.

[SURVIVE-FWK] Sprecher, N. and A. Farrel, "Multiprotocol Label Switching Transport Profile Survivability Framework", Work in Progress, June 2010.

[VPMS-REQS] Kamite, Y., JOUNAY, F., Niven-Jenkins, B., Brungard, D., and L. Jin, "Framework and Requirements for Virtual Private Multicast Service (VPMS)", Work in Progress, October 2009.

[X.200] ITU-T, "Information Technology - Open Systems Interconnection - Basic reference Model: The Basic Model", ITU-T Recommendation X.200, 1994.

Authors' Addresses

Matthew Bocci (editor)
Alcatel-Lucent
Voyager Place, Shoppenhangers Road
Maidenhead, Berks SL6 2PJ
United Kingdom

EMail: matthew.bocci@alcatel-lucent.com

Stewart Bryant (editor)
Cisco Systems
250 Longwater Ave
Reading RG2 6GB
United Kingdom

EMail: stbryant@cisco.com

Dan Frost (editor)
Cisco Systems

EMail: danfrost@cisco.com

Lieven Levrau
Alcatel-Lucent
7-9, Avenue Morane Sulnier
Velizy 78141
France

EMail: lieven.levrau@alcatel-lucent.com

Lou Berger
LabN Consulting, L.L.C.

EMail: lberger@labn.net