

Network Working Group
Internet Draft
Category: Informational
Expires: March 2012

F. Jounay (Ed.)
France Telecom Orange

Y. Kamite
NTT Communications

G. Heron
Cisco

M. Bocci
Alcatel-Lucent

September 08, 2011

Requirements and Framework for Point-to-Multipoint Pseudowires over MPLS PSNs

[draft-ietf-pwe3-p2mp-pw-requirements-05.txt](#)

Status of this Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

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

This Internet-Draft will expire on January 08, 2012.

Abstract

This document presents a set of requirements and a framework for providing a Point-to-Multipoint Pseudowire (PW) over MPLS PSNs. The requirements identified in this document are related to architecture, signaling and maintenance aspects of Point-to-Multipoint PW operation. They are proposed as guidelines for the standardization of such mechanisms. Among other potential applications, Point-to-

Multipoint PWs can be used to optimize the support of multicast layer 2 services (Virtual Private LAN Service and Virtual Private Multicast Service) as defined in the Layer 2 Virtual Private Network Working Group.

Table of Contents

1. Introduction	3
1.1. Problem Statement	3
1.2. Scope of the document	3
2. Definition	4
2.1. Acronyms	4
2.2. Terminology	4
3. P2MP SS-PW Requirements	5
3.1. P2MP SS-PW Reference Model	5
3.2. P2MP SS-PW Underlying Layer	7
3.3. P2MP SS-PW Construction	8
3.4. P2MP SS-PW Signaling Requirements	8
3.4.1. PW Identifier	8
3.4.2. PW type mismatch	9
3.4.3. Interface Parameters sub-TLV	9
3.4.4. Leaf Grafting/Pruning	9
3.5. Failure Detection and Reporting	9
3.6. Protection and Restoration	10
3.7. Scalability	11
4. P2MP MS-PW Requirements	12
4.1. P2MP MS-PW Pseudowire Reference Model	12
4.2. P2MP SS-PW Underlying Layer	13
4.3. P2MP MS-PW Signaling Requirements	14
4.3.1. Dynamically Instantiated P2MP MS-PW	14
4.3.2. P2MP MS-PW Setup Mechanisms	14
4.3.3. PW type mismatch	14
4.3.4. Interface Parameters sub-TLV	15
4.3.5. Leaf Grafting/Pruning	15
4.3.6. Explicit Routing	15
4.4. Failure Detection and Reporting	15
4.5. Protection and Restoration	16
4.6. Scalability	16
5. Manageability considerations	16
6. Backward Compatibility	17
7. Security Considerations	17
8. IANA Considerations	17
9. Acknowledgments	17
10. References	18
10.1. Informative References	18

Authors' Addresses.....	19
Copyright and Licence Notice.....	20

1. Introduction

1.1. Problem Statement

As defined in the pseudowire architecture [[RFC3985](#)], a Pseudowire (PW) is a mechanism that emulates the essential attributes of a telecommunications service (such as a T1 leased line or Frame Relay) over an IP or MPLS PSN. It provides a single service which is perceived by its user as an unshared link or circuit of the chosen service. A Pseudowire is used to transport layer 1 or layer 2 traffic (e.g. Ethernet, TDM, ATM, and FR) over a layer 3 PSN. PWE3 operates "edge to edge" to provide the required connectivity between the two endpoints of the PW.

The Point-to-Multipoint (P2MP) topology described in [VPMS REQ] and required to provide P2MP L2VPN services can be achieved using one or more P2MP PWs. The use of PW encapsulation enables P2MP services transporting layer 1 or layer 2 data. This could be achieved using a set of point to point PWs, with traffic replication on the PE, but at the cost of bandwidth efficiency, as duplicate traffic would be carried multiple times on shared links.

This document defines the requirements for a Point-to-Multipoint PW (P2MP PW). A P2MP PW is a mechanism that emulates the essential attributes of a P2MP Telecommunications service such as a P2MP ATM VC over a PSN. The required functions of P2MP PWs include encapsulating service-specific PDUs arriving at an ingress Attachment Circuit (AC), and carrying them across a tunnel to one or more egress ACs, managing their timing and order, and any other operations required to emulate the behavior and characteristics of the service as faithfully as possible.

P2MP PWs therefore extend the PWE3 architecture [[RFC3985](#)] to offer a P2MP Telecommunications service.

This document also defines the associated requirements related to the P2MP PW operation (e.g. setup and maintenance, protection and scalability).

1.2. Scope of the document

The document describes the P2MP PW Reference Model architectures and outlines specific signaling requirements for the set up and maintenance of a P2MP PW. The requirements are divided into two parts, i.e. those applicable in a Single-Segment PW architecture and those applicable in a Multi-Segment PW architecture. For other aspects of P2MP PW implementation, such as packet processing ([section 4](#)) and Faithfulness of Emulated Services ([section 7](#)), the document

refers to [[RFC3916](#)].

Jounay et al.

Expires March 2012

[Page 3]

Some P2MP PW requirements are derived from the signaling requirements for P2MP Traffic-Engineered MPLS Label Switched Paths [[RFC4461](#)].

[2. Definition](#)

[2.1. Acronyms](#)

P2P: Point-to-Point

P2MP: Point-to-Multipoint

PW: Pseudowire

PSN: Packet Switched Network

SS-PW: Single-Segment Pseudowire

MS-PW: Multi-Segment Pseudowire

[2.2. Terminology](#)

This document uses terminology described in [[RFC5254](#)] and [[RFC5659](#)].

It also introduces additional terms needed in the context of P2MP PW.

P2MP PW, (also referred as PW Tree)

Point-to-Multipoint Pseudowire. A PW attached to a source CE used to distribute Layer 1 or Layer 2 traffic to a set of one or more receiver CEs. The P2MP PW is unidirectional and optionally bidirectional.

P2MP SS-PW

Point-to-Multipoint Single-Segment Pseudowire. A single segment P2MP PW set up between the PE attached to the source CE and the PEs attached to the receiver CEs. The P2MP SS-PW uses P2MP LSPs as PSN tunnels.

P2MP MS-PW

Point-to-Multipoint Multi-Segment Pseudowire. A multi-segment P2MP PW represents an End-to-End PW segmented by means of S-PEs which perform PW label switching. Each segment can use either a P2P LSP or a P2MP LSP as its PSN tunnel.

Root PE

P2MP PW Root Provider Edge. The PE attached to the traffic source CE

for the P2MP PW via an Attachment Circuit (AC). In a MS-PW architecture the term used is Root T-PE.

Leaf PE

P2MP PW Leaf Provider Edge. A PE attached to a set of one or more traffic receiver CEs, via ACs. The Leaf PE replicates traffic to the CEs based on its Forwarder function [[RFC3985](#)].

Branch S-PE

The Branch S-PE is only defined and required in the context of MS-PWs. The Branch S-PE has one upstream PW segment, which may be P2P or P2MP, and one or more downstream PW segments, which may also be P2P or P2MP.

P2MP PSN Tunnel

In the P2MP SS-PW topology, The PSN Tunnel is a general term indicating a virtual P2MP connection between the Root PE and the Leaf PEs. A P2MP tunnel may potentially carry multiple P2MP PWs inside (aggregation). This document uses terminology from the document describing the MPLS multicast architecture [[RFC5332](#)] for MPLS PSN.

[3.](#) P2MP SS-PW Requirements

[3.1.](#) P2MP SS-PW Reference Model

A P2MP SS-PW provides Point-to-Multipoint connectivity from a Root PE connected to a traffic source CE to one or more Leaf PEs connected to traffic receiver CEs.

Figure 1 describes the P2MP SS-PW reference model which is derived from [[RFC3985](#)] to support P2MP emulated services.

Specific operations that must be performed at the PE on the native data units are not described here since the required pre-processing (Forwarder (FWRD) and Native Service Processing (NSP)) defined in

[section 4.2 of \[RFC3985\]](#) are also applicable to P2MP PW.

P2MP PWs are generally unidirectional, but a Root PE may need to receive unidirectional P2P traffic from any Leaf PE. For that purpose the P2MP PW can support optional bidirectional connectivity between the Root PE and each Leaf PE

- Downstream: Point-to-Multipoint (Root PE to any Leaf PE)
- Upstream: Point-to-Point or Multipoint-to-Point (any Leaf PE to Root PE).

Depending on the service using the P2MP PW, the Root PE may benefit from information sent by e.g. a Leaf PE using P2P connectivity at the expense of the amount of state and configuration overhead for the P2P return path. However, in most situations a Multipoint-to-point (MP2P) connectivity is expected to be sufficient. Hence it must be possible for the operator to configure the attributes (P2P or MP2P) of the return path.

[3.2.](#) P2MP SS-PW Underlying Layer

If Ingress PE replication is used, a P2MP PW may be supported over multiple P2MP PSN tunnels, or optionally P2P PSN tunnels, or a mix of both. These PSN tunnels must be able to serve more than one P2MP PW. The P2MP SS-PW underlying layer may be P2P, but this will be at the expense of bandwidth consumption.

Typically the P2MP SS-PW implies an underlying P2MP PSN tunnel. Figure 2 gives an example of P2MP SS-PW topology relying on a P2MP LSP. The PW tree is composed of one Root PE (i1) and several Leaf PEs (e1, e2, e3, e4).




Figure 2 Example of P2MP Underlying Layer for P2MP SS-PW

The P2MP Tunnels may also be of different technology (ex. MPLS over GRE, or P-to-MP MPLS LSP) or just use different setup protocols. (ex.

MLDP, and P2MP RSVP-TE).

Jounay et al.

Expires March 2012

[Page 7]

The P2MP LSP associated to the P2MP PW can be selected either by user configuration or by dynamically using a multiplexing/demultiplexing mechanism.

The P2MP PW multiplexing should be used based on the overlap rate between P2MP LSP and P2MP PW. As an example, an existing P2MP LSP may attach more leaves than the ones defined as Leaf PEs for a given P2MP PW. It may be attractive to reuse it to minimize new configuration, but using this P2MP LSP would imply non-Leaf PEs receive unwanted traffic, not destined to Leaf PE at the service layer. The operator should determine whether the P2MP PW can accept partially multiplexing with P2MP LSP, and a minimum congruency rate may be defined. The Root PE can determine whether P2MP PW can multiplex to a P2MP LSP according to the congruency rate. The congruency rate should take into account several items, such as:

- the amount of overlap between the number of Leaf PEs of P2MP PW and existing egress PE routers of a P2MP LSP. If there is a complete overlap, the congruency is perfect and the rate is 100%.
- at the expense of the additional traffic (e.g. other VPNs) supported over the P2MP LSP.

With this procedure a P2MP PW is nested within a P2MP LSP. This allows multiplexing several PWs over a common P2MP LSP. Prior to the P2MP PW signaling phase, the Root PE must determine which P2MP LSP will be used for this P2MP PW. The PSN Tunnel can be an existing PSN tunnel or the Root PE can create a new P2MP PSN tunnel.

3.3. P2MP SS-PW Construction

The following requirements apply to the establishment of P2MP SS-PWs:

- PE nodes must be configurable with the P2MP PW identifiers and ACs.
- A discovery mechanism should allow the Root PE to discover the Leaf PEs, or vice versa.
- Solutions should allow single-sided operation at the Root PE for the selection of some AC(s) at the Leaf PE(s) to be attached to the PW tree so that the Root PE controls the Leaf attachment.

The Root PE should support a method to be informed about whether a Leaf PE has successfully attached to the PW tree.

3.4. P2MP SS-PW Signaling Requirements

3.4.1. PW Identifier

Jounay et al.

Expires March 2012

[Page 8]

The P2MP PW must be uniquely identified. This unique P2MP PW identifier must be used for all signaling procedures related to this PW (PW setup, monitoring, etc).

3.4.2. PW type mismatch

The Root PE and Leaf PEs of a P2MP PW must be configured with the same PW type as defined in [[RFC4446](#)] for P2P PW. In case of a different type, a PE must abort attempts to establish the P2MP PW.

3.4.3. Interface Parameters sub-TLV

Some interface parameters [[RFC4446](#)] related to the AC capability have been defined according to the PW type and are signaled during the PW setup.

Where applicable, a solution is required to ascertain whether the AC at the Leaf PE is capable of supporting traffic coming from the AC at the Root PE.

In case of a mismatch, the passive PE (Root or Leaf PE, depending on the signaling process) must support mechanisms to reject attempts to establish the P2MP SS-PW.

3.4.4. Leaf Grafting/Pruning

Once the PW tree is established, the solution must allow the addition or removal of a Leaf PE, or a subset of leaves to/from the existing tree, without any impact on the PW tree (data and control planes) for the remaining Leaf PEs.

The addition or removal of a Leaf PE must also allow the P2MP PSN tunnel to be updated accordingly. This may cause the P2MP PSN tunnel to add or remove the corresponding Leaf PE.

3.5. Failure Detection and Reporting

Since the underlying layer has an End-to-End P2MP topology between the Root PE and the Leaf PEs, the failure reporting and processing procedures are implemented only on the edge nodes.

Failure events may cause one or more Leaf PEs to become detached from the PW tree. These events must be reported to the Root PE, using appropriate out-of-band or inband OAM messages.

It must be possible for the operator to choose the out-of-band or inband OAM tools or both to monitor the Leaf PE status.

The solution should allow the Root PE to be informed of Leaf PEs failure for management purposes.

Based on these failure notifications, solutions must allow the Root PE to update the remaining leaves of the PW tree.


- A solution must support in-band OAM mechanism to detect failures: unidirectional point-to-multipoint traffic failure. This should be realized by enhancing existing unicast PW methods, such as VCCV for seamless and familiar operation defined in [[RFC5085](#)] and [[RFC6073](#)].
- In case of failure, it should correctly report which Leaf PEs are affected. This should be realized by enhancing existing PW methods, such as LDP Status Notification. The notification message should include the type of fault (P2MP PW, AC or PSN tunnel).
- A Leaf PE may be notified of the status of the Root PE's AC.
- A solution must support OAM message mapping [[RFC6310](#)] at the Root PE and Leaf PE if a failure is detected on the source CE AC.

[3.6.](#) Protection and Restoration


It is assumed that if recovery procedures are required, the P2MP PSN tunnel will support standard MPLS-based recovery techniques (typically based on RSVP-TE). In that case a mechanism should be implemented to avoid race conditions between recovery at the PSN level and recovery at the PW level.

An alternative protection scheme may rely on the PW layer.

Leaf PEs may be protected via a P2MP PW redundancy mechanism. In the example depicted below, a standby P2MP PW is used to protect the active P2MP. In that protection scheme the AC at the Root PE must serve both P2MP PWs. In this scenario, the condition when to do the switchover should be implemented, e.g. one or all Leaf failure of active P2MP PW will course P2MP PW switchover.



The Root PE may be protected via a P2MP PW redundancy mechanism. In the example depicted below, a standby P2MP PW is used to protect the active P2MP. A single AC at the Leaf PE must be used to attach the CE to the primary and the standby P2MP PW. The Leaf PE must support protection mechanisms in order to select the active P2MP PW.



3.7. Scalability

The solution should scale at least linearly with the number of Leaf PEs.

Increasing the number of P2MP PWs between a Root PE and a given set of Leaf PEs should NOT cause the P router to increase the number of entries in its forwarding table by the same or greater proportion. Multiplexing P2MP PWs to P2MP PSN Tunnels achieves this.

4. P2MP MS-PW Requirements

4.1. P2MP MS-PW Pseudowire Reference Model

Figure 3 describes the P2MP MS-PW reference model which is derived from [RFC5659] to support P2MP emulated services.

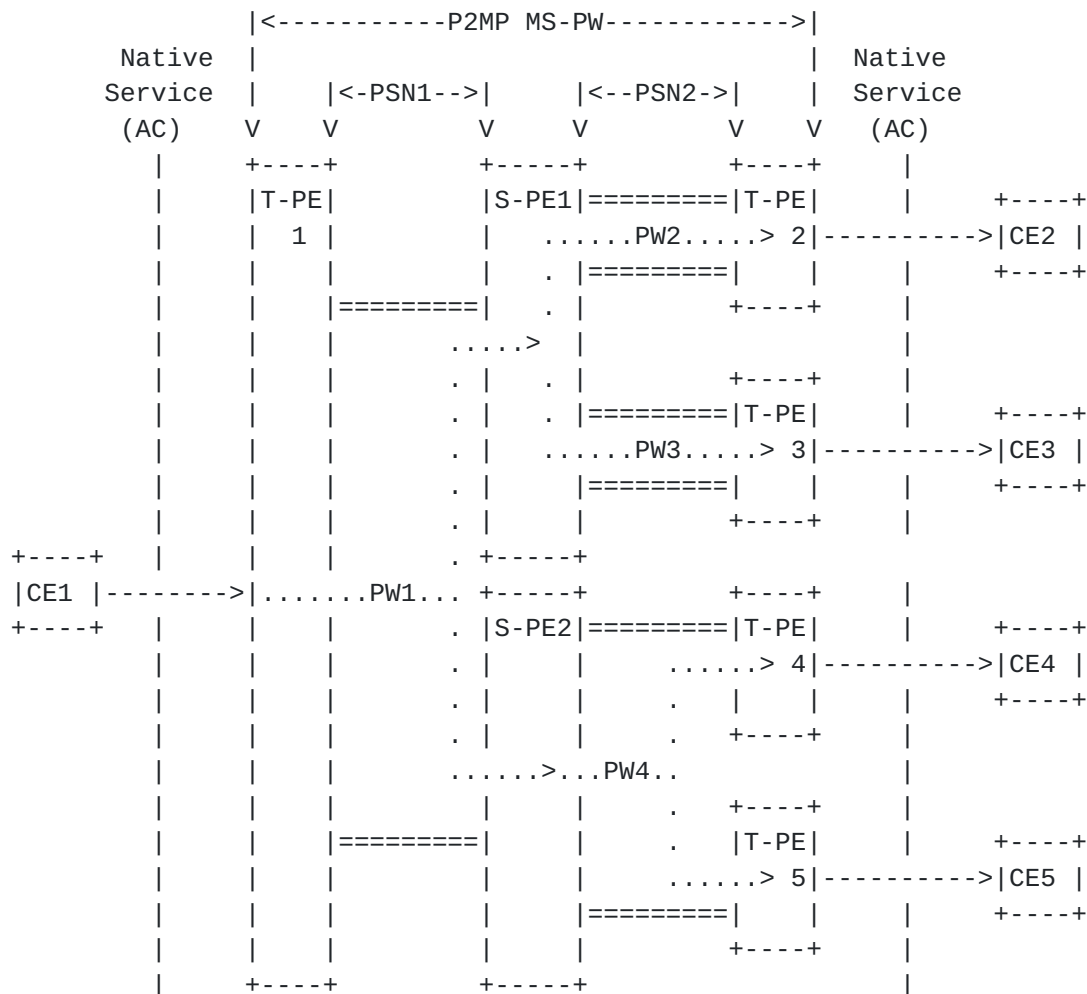


Figure 3 P2MP MS-PW Reference Model

Figure 3 extends the P2MP SS-PW architecture of Figure 1 to a multi-segment configuration. In a P2P MS-PW configuration as described in [RFC5659] the S-PE is responsible for switching a MS-PW from one ingress segment to only one egress segment, based on the PW identifier. Here in a P2MP MS-PW configuration the S-PE is responsible for switching a MS-PW from one ingress segment to one or more egress segments.

Referring to Figure 3, T-PE1 is the Root T-PE and T-PE2, T-PE3, T-PE4 and T-PE5 are the Leaf T-PEs. In the reference model, the Leaf T-PEs

are assumed to be located in the same PSN (PSN2), but it could be envisioned that each egress PW is located in a different PSN (PSN2, PSN3, PSN4). S-PEs play the role of Branch S-PEs since S-PE1 and S-

PE2 are in charge respectively of switching the ingress P2MP PW1 segment to the egress P2P PW2, P2P PW3 and P2MP PW4 segments.

A P2MP MS-PW may transit through more than one S-PE along its path.

As depicted in Figure 3 a PW segment belonging to a P2MP MS-PW can be supported over a P2MP PSN tunnel or a P2P PSN tunnel.

The Reference Model outlines the basic pieces of a P2MP MS-PW, however several levels of replication may be used when designing a P2MP MS-PW

- Ingress T-PE replication: traffic replicated to a set of P2P or P2MP PSN tunnels or to local receiver CEs
- P router replication: traffic replicated by means of P2MP PSN tunnel (P2MP LSP)
- S-PE replication: traffic replicated to a set of P2P or P2MP PSN tunnels
- Egress T-PE replication : traffic replicated to local receiver CEs

As described in [section 3.1](#), P2MP MS-PWs are generally unidirectional, but a Root T-PE may need to receive unidirectional P2P traffic from any Leaf PE. For that purpose the P2MP MS-PW may support bidirectional connectivity between the Root T-PE and each Leaf T-PE.

4.2. P2MP SS-PW Underlying Layer

Due to Ingress PE or S-PE replication, the P2MP PW segment may be supported over multiple concatenated P2MP PSN tunnels and optionally P2P PSN tunnels or a mix of both.

Figure 4 describes an example of a P2MP MS-PW architecture relying on a combination of both P2P and P2MP LSPs as PSN tunnels. PW segments over P2P LSPs may be used to address inter-provider requirements, for example. The PW tree is composed of one Root PE (i1) and several Leaf PEs (e1, e2, e3, e4). The Branch S-PEs are represented as b1, b2, b3, b4, b5. In this case the traffic replication along the path of the PW tree is performed at the PW level. For example, the Branch S-PE b5 must replicate incoming packets or data received from b2 and send them to Leaf T-PEs e3 and e4.

However since some PW segments may be supported over a P2MP LSP, the traffic replication along the path of these PW segments can be performed at the underlying LSP level.

Figure 4 describes the case where each segment is supported over a P2P LSP except for the b1-b3b4 P2MP segment which is conveyed over a P2MP LSP on this segment.




Figure 4 Example of P2P and P2MP underlying Layer for P2MP MS-PW

The mechanisms for establishing the PSN tunnel are outside the scope of this document, as long as they enable the essential attributes of the service to be emulated.

4.3. P2MP MS-PW Signaling Requirements

4.3.1. Dynamically Instantiated P2MP MS-PW

The PW tree could be statically configured at each T-PE and S-PE along its path. However it is recommended that a solution provides the ability to dynamically setup a MS-PW tree, by allowing the MS-PW segments to be dynamically discovered at S-PE.

During the PW tree setup, a Branch S-PE should be capable of informing the upstream PEs, including the Root T-PE that a set of Leaf T-PEs and associated leaves are not reachable.

4.3.2. P2MP MS-PW Setup Mechanisms

The requirements described in this section assume that dynamic setup of MS-PW segments allows the T-PEs and S-PEs to dynamically signal MS-PW segments and stitch these segments in order to build the MS-PW tree.

4.3.3. PW type mismatch

As described for P2MP SS-PW, the P2MP MS-PW requires ACs of the same PW type. Therefore the segments composing the P2MP MS-PW must be also

of the same PW type [[RFC4446](#)]. When P2MP MS-PW is statically configured, the S-PE must support switching PWs of the same PW type

as described in [[RFC5659](#)]. When MS-PW is dynamically configured by signaling, in case of a different type a PE must abort attempts to establish the P2MP MS-PW.

4.3.4. Interface Parameters sub-TLV

[Section 3.4.3](#) is also relevant to P2MP MS-PW. When applicable, the Leaf T-PE or the Root T-PE must signal its AC interface parameters to the Root T-PE or the Leaf T-PEs to make sure the AC at each Leaf T-PE is capable of supporting traffic coming from the AC at the Root T-PE. In the P2MP MS-PW case, S-PEs must propagate this information.

In case of a mismatch, the passive T-PE (Root or Leaf T-PE, depending on the signaling process) must support mechanisms to reject attempts to establish the P2MP MS-PW.

4.3.5. Leaf Grafting/Pruning

Once the PW tree is setup, the solution must allow the addition or removal of a Leaf T-PE, or a subset of leaves to/from the existing tree, without any impact on the PW tree (data and control planes) for the remaining Leaf T-PEs.

4.3.6. Explicit Routing

The P2MP MS-PW signaling solution must provide a means of establishing P2MP MS-PWs according to pre-computed and configured S-PE paths as well as dynamically computing S-PE paths at the Root T-PE.

To support the setup of an explicitly routed MS-PW tree, the signaling solution should support the ability for a Root PE to explicitly define particular S-PE nodes as Branch S-PEs for the PW tree.

The solution should enable Explicit Path Loose Hops. Therefore the P2MP MS-PW may be partially specified with only a subset of intermediate Branch S-PEs.

4.4. Failure Detection and Reporting

The solution should rely on specific OAM mechanisms to detect a node (T-PE and S-PE) or segment failure of a PW tree. The solution should also support the ability to inform the Root T-PE of the failure as well as to indicate the identity of affected Leaf T-PEs.

Based on these failure notifications the solution must allow the Root

T-PE to update the remaining Leaf T-PEs of the PW tree.

Jounay et al.

Expires March 2012

[Page 15]

- A solution must support in-band OAM mechanism to detect unidirectional point-to-multipoint traffic failure. This should be realized by enhancing existing unicast PW methods, such as VCCV for seamless and familiar operation.
- In case of a failure, it should report which Leaf T-PEs and Branch S-PEs are affected. This should be realized by enhancing existing unicast PW methods, such as LDP Status Notification. The notification message should include the type of fault (P2MP PW, AC or PSN tunnel).
- A Leaf T-PE may be notified of the status of the Root PE's AC.
- A solution must support OAM message mapping [[RFC6310](#)] at the Root T-PE and Leaf T-PE if a failure is detected on the source CE AC.

4.5. Protection and Restoration

The solution should provide mechanisms to recover the emulated service as fast as possible following a failure event.

In the case of Root-initiated PW tree setup, where a local repair (PSN-tunnel or PW segment-based) is not feasible after a failure event, and where the PE upstream to a failure is notified that a subset of Leaf T-PEs have become detached from the PW tree, solutions should allow the upstream PE to re-compute the path to those particular Leaf T-PEs. If the upstream PE fails to compute an alternative path, this procedure should be propagated upstream until the Root T-PE is reached.

Note that recovery procedures can be implemented at the underlying P2P or P2MP LSP layer, using standard MPLS-based recovery techniques.

A mechanism should be provided to avoid race conditions between recovery at the PSN level and recovery at the PW level.

4.6. Scalability

Solutions for P2MP MS-PW must take into account scalability considerations.

Solutions must scale linearly, or better, with an increase in the number of Leaf T-PEs and Branch S-PEs. Scalability issues must be addressed for the control plane (e.g. addressing of PW endpoints, number of signaling sessions, etc.) and the data plane (e.g. duplication of PW segments, OAM mechanism, etc.).

5. Manageability considerations

Jounay et al.

Expires March 2012

[Page 16]

The solution should provide a simple provisioning procedure to build a P2MP SS-PW or a P2MP MS-PW.

The solution must take into consideration the situation where the Root PE and Leaf PEs are not managed by a single NMS.

In that case it must be possible to manage the whole P2MP PW using a single NMS. Typically the P2MP PW could be managed from the Root PE.

6. Backward Compatibility

Solutions must be backward compatible with current PW standards. Solutions should utilize existing capability advertisement and negotiation procedures for the PEs implementing P2MP PW endpoints.

The implementation of OAM mechanisms also implies the advertisement of PE capabilities to support specific OAM features. The solution may allow advertising P2MP PW OAM capabilities.

A solution must NOT allow a P2PW to be established to PEs that do not support P2MP PW functionality. It must have a mechanism to report an error for incompatible PEs. In this case, it should report which PEs (S-PE and T-PEs) are not compatible.

In some cases, upstream traffic is required from downstream CEs to upstream CEs. The P2MP PW solution should allow a return path (i.e. from the Leaf to the Root) that provides upstream connectivity.

In particular, the same ACs may be shared between downstream and upstream directions. For downstream, a CE receives traffic originated by the Root PE over its AC. For upstream, the CE may also send traffic destined to the same Root PE over the same AC.

7. Security Considerations

The security requirements common to PW are raised in [Section 10 of \[RFC3916\]](#) and common to MS-PW in [section 7 of \[RFC5254\]](#). P2MP PW (SS or MS) is a variant of the initial P2P PW definition, and those sections also apply to P2MP PW.

8. IANA Considerations

This draft does not require any IANA action.

9. Acknowledgments

Jounay et al.

Expires March 2012

[Page 17]

The authors thank the authors of [[RFC4461](#)] since the structure and content of this document were, for some sections, largely inspired by [[RFC4461](#)].

Many thanks to JL Le Roux and A. Cauvin for the discussions, comments and support.

[10. References](#)

[10.1. Informative References](#)

- [RFC5332] Rosen, E. et al., "MPLS Multicast Encapsulations", August 2008
- [RFC4446] Martini, L. "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", April 2006
- [RFC5085] Nadeau, T., Pignataro, C. "Pseudowire Virtual Circuit Connectivity Verification (VCCV)", December 2007
- [RFC6073] Martini, L. et al. "Segmented Pseudowire", January 2011
- [RFC3985] Bryant, S., Pate, P. "PWE3 Architecture", March 2005
- [RFC3916] McPherson, D., Pate, P., Xiao, X., "Requirements for Pseudo-Wire Emulation Edge-to-Edge", September 2004
- [RFC4461] Aggarwal, R., Farrel, A., Jork, M., Kamite, Y., Kullberg, A., Le Roux, JL., Malis, A., Papadimitriou, D., Vasseur, JP., Yasukawa, S., "Signaling Requirements for P2MP TE MPLS LSPs", April 2006
- [RFC5254] Bitar, N., Bocci, M., and Martini, L., "Requirements for inter domain Pseudo-Wires", June 2008
- [RFC5659] Bocci, M., and Bryant, S., "An Architecture for Multi-Segment Pseudo Wire Emulation Edge-to-Edge", October 2009
- [RFC6310] Aissaoui, M., et al. "Pseudowire OAM Message Mapping", Internet Draft, July 2011
- [VPMS REQ] Kamite, Y., Jounay, F. "Framework and Requirements for Virtual Private Multicast Service (VPMS)", Internet Draft, [draft-ietf-l2vpn-vpms-frmwk-requirements-04](#), July 2011

Author's Addresses

Frederic Jounay
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
FRANCE
Email: frederic.jounay@orange-ftgroup.com

Philippe Niger
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
FRANCE
Email: philippe.niger@orange-ftgroup.com

Yuji Kamite
NTT Communications Corporation
Tokyo Opera City Tower
3-20-2 Nishi Shinjuku, Shinjuku-ku
Tokyo 163-1421
Japan
Email: y.kamite@ntt.com

Luca Martini
Cisco Systems, Inc.
9155 East Nichols Avenue, Suite 400
Englewood, CO, 80112
EMail: lmartini@cisco.com

Giles Heron
Cisco Systems, Inc.
9 New Square
Bedfont Lakes
Feltham
Middlesex
TW14 8HA
United Kingdom
EMail: giheron@cisco.com

Lei Wang
Telenor
Snaroyveien 30
Fornebu 1331
Norway
Email: lei.wang@telenor.com

Rahul Aggarwal

Juniper Networks
1194 North Mathilda Ave.
Sunnyvale, CA 94089
Email: rahul@juniper.net

Jounay et al.

Expires March 2012

[Page 19]

Simon Delord
Alcatel-Lucent
Building 3, 388 Ningqiao Road, Jinqiao, Pudong
Shanghai, 201206, P.R. China
Email: simon.delord@alcatel-lucent.com

Martin Vigoureux
Alcatel-Lucent France
Route de Villejust
91620 Nozay
FRANCE
Email: martin.vigoureux@alcatel-lucent.fr

Matthew Bocci
Alcatel-Lucent Telecom Ltd,
Voyager Place
Shoppenhangers Road
Maidenhead
Berks, UK
E-mail: matthew.bocci@alcatel-lucent.co.uk

Lizhong Jin
ZTE Corporation
889, Bibo Road,
Shanghai, 201203, China
Email: lizhong.jin@zte.com.cn

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.
This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may

not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.