Network Working Group
Internet Draft
Category: Informational
Expires: January 2011

F. Jounay (Ed.)
P. Niger
France Telecom Orange

Y. Kamite
NTT Communications

S. Delord
Testra

L. Wang
Telenor

G. Heron
BT

L. Martini
Cisco

R. Aggarwal
Juniper Networks

M. Bocci
M. Vigoureux
Alcatel-Lucent

L. Jin
ZTE

August 18, 2010

Requirements for Point-to-Multipoint Pseudowire

draft-ietf-pwe3-p2mp-pw-requirements-03.txt

Status of this Memo

Abstract

This document presents a set of requirements for providing a Point-to-Multipoint PWE3 (Pseudowire Emulation Edge to Edge) emulation. The requirements identified in this document are related to architecture, signaling and maintenance aspects of a Point-to-Multipoint PW operation. They are proposed as guidelines for the standardization of such mechanisms. Among other potential applications Point-to-Multipoint PWs SHOULD be used to optimize the support of multicast services (Virtual Private LAN Service and Virtual Private Multicast Service) as defined in the Layer 2 Virtual Private Network working group.

Conventions used in this document

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

Table of Contents

1. Introduction

1.1. Problem Statement

As defined in the PWE3 WG charter, a Pseudowire (PW) emulates a point-to-point bidirectional link over an IP/MPLS network, and 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 non IP traffic (e.g. Ethernet, TDM, ATM, and FR) in an IP/MPLS-based PSN (Packet Switched Network). PWE3 operates "edge to edge" to provide the required connectivity between the two endpoints of the PW.

The P2MP topology mentioned in [VPMS REQ] and required to provide P2MP L2VPN services can be achieved via a P2MP PW. The use of PW becomes necessary for some P2MP services requiring specific

encapsulation capabilities. This could be achieved using a set of point to point PWs, with traffic replication on the PE, but faces obvious bandwidth limitation issues, as traffic is carried multiple time on shared links.

This document defines the requirements for the use of a Point-to-Multipoint PW (P2MP PW). A Point-to-Multipoint (P2MP) Pseudowire (PW) is a mechanism that emulates the essential attributes of a P2MP Telecommunications service such as P2MP ATM over PSN. One of the applicabilities of a P2MP PW is to deliver a non-IP multicast service that carries multicast frames from a multicast source to one or more multicast receivers. 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 extend the PWE3 architecture [RFC3985] to offer a P2MP Telecommunications service.
This document aims at defining the associated requirements related to the P2MP PW operation (e.g. setup and maintenance, protection, 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 topology and those applicable in a Multi-Segment topology. For other aspects of P2MP PW implementation like packet processing (section 4) and Faithfulness of Emulated Services (section 7), the document refers to [RFC3916].

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

SS-PW: Single-Segment Pseudowire

MS-PW: Multi-Segment Pseudowire

2.2. Terminology

This document uses terminology described in [RFC5254], [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 used to distribute L1/L2 format traffic to a set of one or more receivers (or leaves). 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 and the PEs attached to the receivers. The P2MP SS-PW relies on P2MP LSP as PSN tunnel.

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 are in charge of switching the PW label. Each segment can rely on either P2P LSP or P2MP LSP as PSN tunnel.

Root PE

P2MP PW Root Provider Edge. Router attached to a Customer Equipment (traffic source) via an Attachment Circuit (AC). In a MS-PW architecture the term used is Root T-PE.

Leaf PE

P2MP PW Leaf Provider Edge. Router attached to a set of one or more Customer Equipments (traffic receivers or leaves). The P2MP PW is attached to an Attachment Circuit (AC). The Leaf PE is therefore in charge of replicating the traffic over 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-PW. The Branch S-PE has one upstream PW (P2P or P2MP) segment and one or several downstream PW (P2P or P2MP) segments.

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 a Point-to-Multipoint connectivity from a Root PE connected to a traffic source to at least two Leaf PEs connected to traffic receivers. The PW endpoints connect the PW to its Attachment Circuit (AC). As for a P2P PW, an AC can be a Frame Relay DLC, an ATM VP/VC, an Ethernet port, a VLAN, a HDLC link on a physical interface.

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

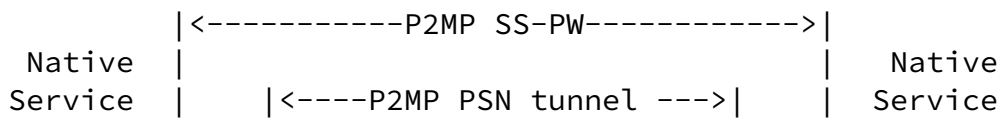

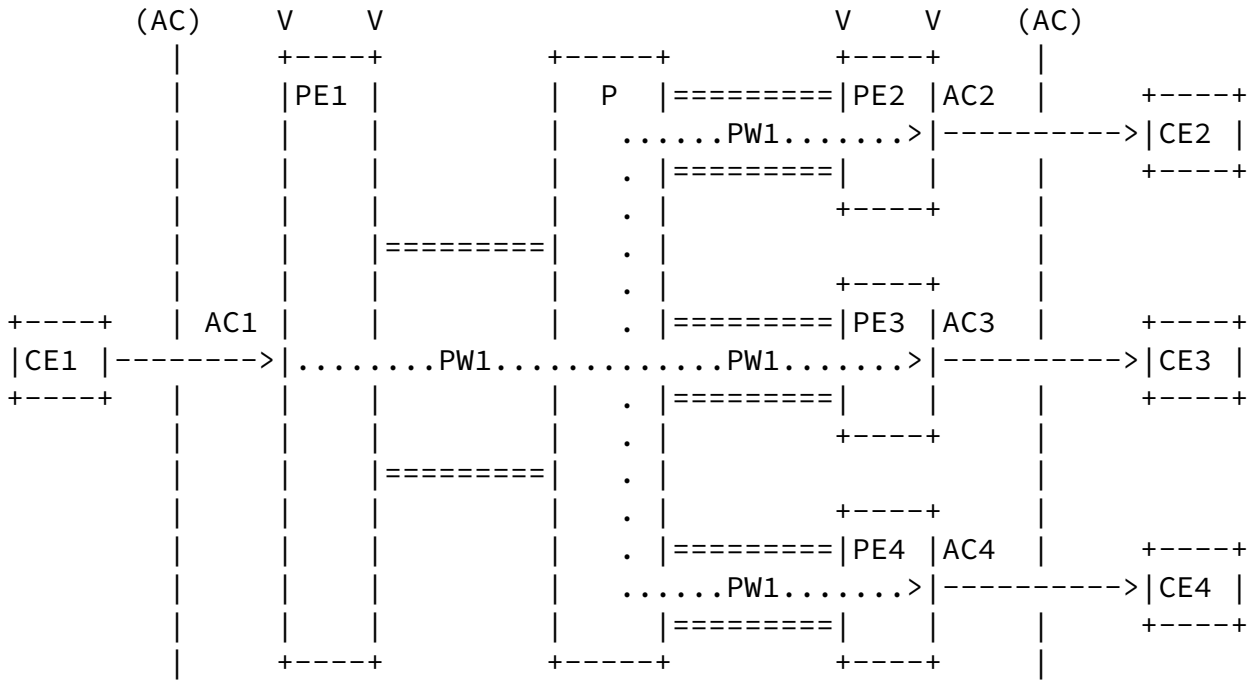

Figure 1 P2MP SS-PW Reference Model

This architecture applies to the case where a P2MP PSN tunnel extends between edge nodes of a single PSN domain to transport a unidirectional P2MP PW with endpoints at these edge nodes.
In this model a single copy of each PW packet is sent over the P2MP PSN tunnel and is received by all Leaf PEs due to the P2MP nature of the PSN tunnel. P2MP PW MUST be traffic optimised, only one copy of P2MP PW packet on one single link. P Router is joining P2MP PSN tunnel operation but is not participating in the signaling of P2MP PW. P2MP PW operation is associated with PE1, PE2, PE3 and PE4.

Specifics operations that must be perfomed 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 the section 4.2 [RFC3985] are also applicable to P2MP PW.

In nature the P2MP PW is unidirectional, but it may be required for a Root PE to receive unidirectional P2P traffic from any Leaf PE. For that purpose the P2MP PW MAY 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 (any Leaf PE to Root PE)

3.2. P2MP SS-PW Underlying Layer

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).

The P2MP PSN MAY be signaled with P2MP RSVP-TE [RFC4875] or MLDP [MLDP].

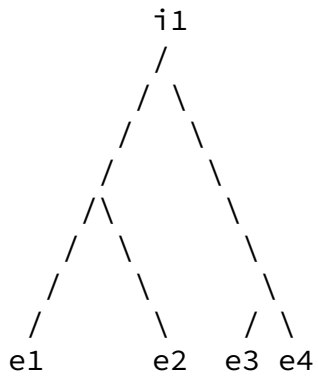

Figure 2 Example of P2MP Underlying Layer for P2MP SS-PW

The P2MP PW MAY be supported over multiple P2MP PSN tunnel. These P2MP PSN tunnels MUST be able to serve more than one P2MP 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).

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

The P2MP PW multiplexing will be based on the overlap rate between P2MP LSP and P2MP PW. The operator should determine whether the P2MP PW can accept partially multiplexing with P2MP LSP, and a minimum congruency rate may be defined. The congruency rate reflects the amount of overlap in the Leaf PE of P2MP PW that is multiplexed to a P2MP LSP. If there is a complete overlap, the congruency is perfect and the rate is 100%. The Root PE can determine whether P2MP PW can multiplex to a P2MP LSP according to the congruency rate. It is also possible to extend P2MP LSP to do P2MP PW multiplexing, but this will reduce the current congruency rate that the P2MP PW is currently taken. The multiplexing should ensure that the P2MP PW congruency that is currently taken under P2MP LSP should be larger than minimum

congruency that is configured.

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

As initial step PE nodes have to be configured with P2MP PW identifier and ACs.
Then discovery mechanism SHOULD allow PE to discover remote PEs configuration.
Eventually the solution 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.
Note that the Root PE single sided operation is a management requirement and does not presume any signaling requirement.

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

3.4. P2MP SS-PW Signaling Requirements

3.4.1. PW Identifier

The P2MP PW MUST be uniquely identified. This unique P2MP PW identifier MUST be used for all the signaling procedure related to this PW (PW setup, monitoring).

3.4.2. PW type mismatch

As for P2P PW, the ACs configured at Root PE and Leaf PEs of a P2MP PW MUST be of the same PW type [RFC4446]. In case of a different type, the passive PE (Root or Leaf PE, depending on the signaling process) MUST support mechanisms to reject 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.
When applicable, this mechanism used for the P2P PW setup MUST be enhanced for P2MP PW setup so as to ascertain that AC at the Leaf PE is capable to support traffic coming from AC at the Root PE.

In case of mismatch, the passive PE (Ingress 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 setup, 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 to the P2MP PSN tunnel to be updated accordingly. This MAY cause 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-band or inband OAM messages.
The solution SHOULD allow the Root PE to be informed of Leaf PEs failure for management purposes.

Based on these failure notifications the solution 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.

- 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 Notification for seamless and familiar operation. The notification message SHOULD include the type of fault (P2MP PW, AC or PSN tunnel).

- Respectively a Leaf PE also MAY receive the status of the Root PE's AC status.

- A solution MUST support OAM message mapping at the Root PE if failure is detected on the AC. The Leaf PE MUST report accordingly at the service layer this OAM message on its associated 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.

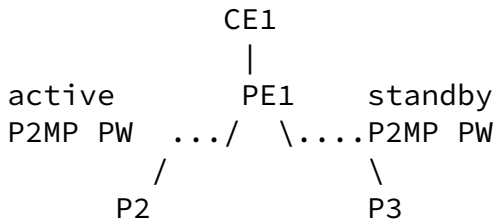

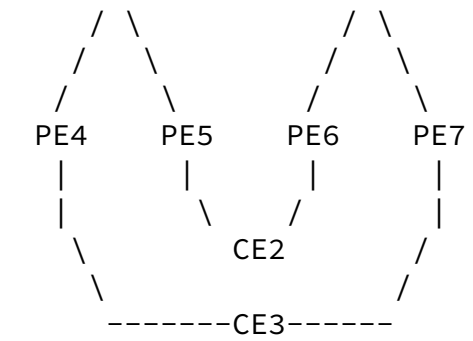

Root PE MAY be protected via a P2MP PW redundancy mechanism. In the example depicted below, a standby P2MP PW is used to protect the active P2MP. A single AC at the Leaf PE MUST be used to attach the CE to the primary and the standby P2MP PW. The Leaf PE MUST support protection mechanism in order to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

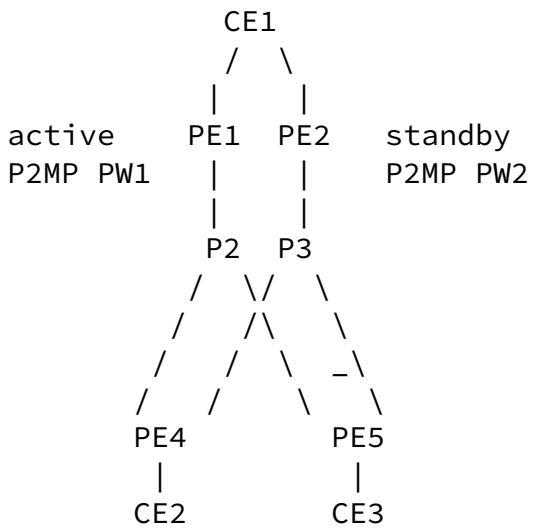

3.7. Scalability

The solution SHOULD scale at least as well as linearly with an increase in the number of Leaf PEs.

An increase in the number of P2MP PW SHOULD NOT cause the P router to increase its forwarding table linearly.

The P2MP PW multiplexed/demultiplexed to P2MP PSN Tunnel can improve the scalability.

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.

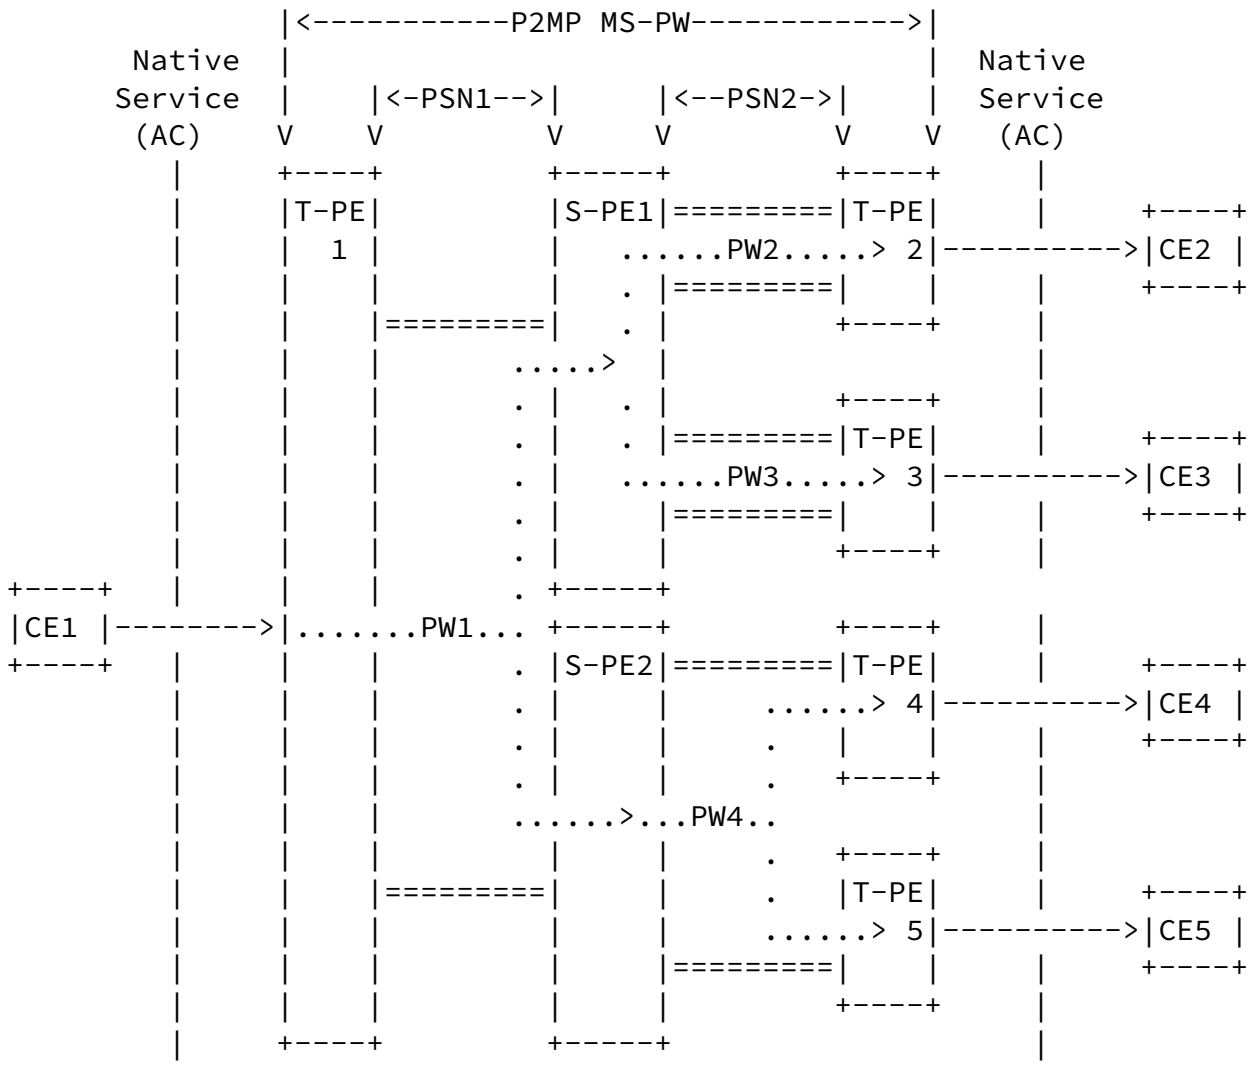

Figure 3 P2MP MS-PW Reference Model

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

[RFC5254] the S-PE is responsible to switch a MS-PW from one input segment to only one output segment, based on the PW identifier. Here in a P2MP MS-PW configuration the S-PE is responsible to switch a MS-PW from one input segment to one or several output 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 output PW is located in a different PSN (PSN2, PSN3, PSN4). The S-PE plays the role of Branch S-PE since S-PE1 and S-PE are in charge respectively of switching simultaneously the input P2MP PW1 segment to the output P2P PW2, P2P PW3 and P2MP PW4 segments.

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

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

4.2. P2MP SS-PW Underlying Layer

Figure 4 describes an example of P2MP MS-PW topology relying on a combination of both P2P and P2MP LSPs as PSN tunnels. PW segment over P2P LSP MAY address inter-provider requirement. 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 that case the traffic replication along the path of the PW tree is performed at the PW level. For instance 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 giving the fact that some PW segments MAY be supported over a P2MP LSP, the traffic replication along the path of these PW segments can be performed as well 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 section.

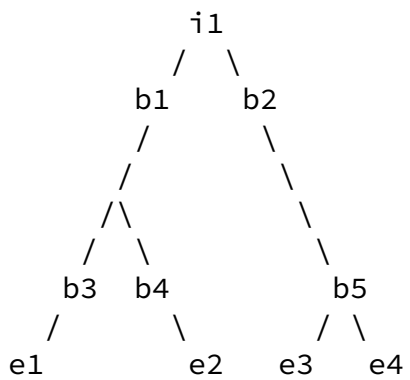

Figure 4 Example of P2P and P2MP underlying Layer for P2MP MS-PW

The P2MP PSN MAY be signaled with P2MP RSVP-TE [RFC4875] or MLDP [MLDP].

4.3. P2MP MS-PW Signaling Requirements

4.3.1. Dynamically Instantiated P2MP MS-PW

The PW tree could be statically configured at the T-PEs and each S-PE crossed. 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.

During the PW tree setup, a Branch S-PE SHOULD be capable to inform 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-PE 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]. The S-PE MAY only support switching PWs of the same PW type. In case of a different type, the passive PE (S-PE or T-PE) MUST support mechanisms to reject attempts to establish the P2MP MS-PW.

4.3.4. Interface Parameters sub-TLV

The section 3.4.3 is also relevant to P2MP MS-PW. When applicable,

the Leaf T-PE or the Root T-PE MUST signal respectively its AC' interface parameters to the Root T-PE or to the Leaf T-PE so as to make sure that AC at the Leaf T-PE is capable to support traffic coming from AC at the Root T-PE. In the P2MP MS-PW case, S-PEs MUST propagate correctly this information.
In case of 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 arbitrary P2MP MS-PW, according to pre-computed and configured S-PE paths as well as dynamically computed S-PE paths on the Root T-PE.

To support setup of explicitly routed MS-PW tree, the signaling solution SHOULD support some source-based control that can explicitly define particular S-PE nodes as Branch S-PEs for the PW tree.

The solution SHOULD let possible 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.

- 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.

- In case of failure, it SHOULD correctly report which Leaf T-PEs and Branch S-PEs are affected. This SHOULD be realized by enhancing existing unicast PW methods, such as LDP Notification for seamless and familiar operation. The notification message SHOULD include the type of fault (P2MP PW, AC or PSN tunnel).

- Respectively a Leaf T-PE also MAY receive the status of the Root PE's AC status.

- A solution MUST support OAM message mapping at the Root T-PE if failure is detected on the AC. The Leaf T-PE MUST report accordingly at the service layer this OAM message on its associated AC.

4.5. Protection and Restoration

The solution SHOULD provide mechanisms to recover as fast as possible following a failure event. The fast protection/recovery is typically dedicated to P2MP applications sensitive to traffic disruption.

Considering (i) a Root-initiated PW tree setup and (ii) that a local repair (PSN-tunnel or PW segment-based) is not feasible after a failure event and that (iii) the PE upstream to the failure receives by means of OAM mechanisms a message indicating that a subset of Leaf T-PEs are detached from the PW tree, the solution SHOULD allow the upstream PE to re-compute the path to those particular Leaf T-PEs. If the upstream PE failed to compute an alternative path, the procedure SHOULD be propagated upstream until the Root T-PE is reached.

It is also assumed that recovery procedures can be implemented at the underlying P2P or P2MP LSP layer, using standard MPLS-based recovery techniques. These procedures could be used to provide faster recovery time in case of link or node failure affecting this layer.

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

4.6. Scalability

In definition of solution for P2MP MS-PW a particular attention MUST be dedicated to scalability.

The solution MUST be designed to scale as well as linearly with an increase in the number of Leaf T-PEs, Branch S-PEs. The scalability issues MUST be addressed for the control plane (e.g. addressing of PW endpoints, number of signaling sessions, etc) and for data plane (e.g. duplication of PW segments, OAM mechanism, etc).

5. Manageability considerations

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

The solution SHOULD be completely backward compatible with the current PW standards. The solution SHOULD take into account the capability advertisement and negotiation procedures for the PEs implementing P2MP PW endpoints.

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 PW connection with non-compliant PEs. It MUST have a mechanism to report an error for non-compliant PEs. In this case, it SHOULD report which PE (S-PE and T-PEs) are not

compliant.

In some cases, upstream traffic is required from downstream CE to upstream CE. The P2MPPW solution SHOULD allow a return path (i.e. from the Leaf to the Root) that provides upstream connection. In particular, it is expected to be allowed that the same ACs are shared between downstream and upstream direction. For downstream, a CE receives from its connected AC traffic originated by the Root PE transported over a P2MP PW. For upstream, the CE MAY also send over the same AC traffic destined to the same remote PE.

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 that statements also apply to P2MP PW.

8. IANA Considerations

This draft does not define any new protocol element, and hence does not require any IANA action.

9. Acknowledgments

The authors thank the contributors 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. Normative References

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

[RFC3916] McPherson, D.,Pate, P., Xiao, X., "Requirements for Pseudo-Wire Emulation Edge-to-Edge", September 2004

[RFC3985] Bryant, S., Pate, P. "PWE3 Architecture", March 2005

[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

[RFC4875] Aggarwal, R., Papadimitriou, D., Yasukawa, S., "Extensions to RSVP-TE for Point-to-Multipoint TE LSPs", MAY 2007

[RFC4446] Martini, L. "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", April 2006

[RFC5254] Bitar, N., Bocci, M., and Martini, L., "Requirements for inter domain Pseudo-Wires", June 2008

[RFC5332] Rosen, E. et al., "MPLS Multicast Encapsulations", August 2008

[RFC5659] Bocci, M., and Bryant, S.,T., " An Architecture for Multi-Segment Pseudo Wire Emulation Edge-to-Edge", October 2009

10.2. Informative References

[MLDP] Minei, I., Wijnands, I., Thomas, B., "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", Internet Draft, draft-ietf-mpls-ldp-p2mp-10, July 2010

[VPMS REQ] Kamite, Y., Jounay, F. "Framework and Requirements for Virtual Private Multicast Service (VPMS)", Internet Draft, draft-ietf-l2vpn-vpms-frmwk-requirements-03, July 2010

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
BT
UK
EMail: giles.heron@gmail.com

Simon Delord
Telstra
242 Exhibition St
Melbourne VIC 3000
Australia
Email: simon.a.delord@team.telstra.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

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