Network Working Group
Internet Draft
Expiration Date: August 2005

Luca Martini (Editor)
Cisco Systems Inc.

Matthew Bocci (Editor)
Alcatel

Nabil Bitar (Editor)
Verizon

February 2005

Requirements for inter domain Pseudo-Wires

draft-martini-pwe3-mh-pw-requirements-01.txt

Status of this Memo

Abstract

This document describes the necessary requirements to allow a service provider to extend the reach of pseudo-wires across multiple domains. These domains can be autonomous systems under one provider administrative control, IGP areas in one autonomous system, different autonomous systems under the administrative control of two or more

service providers, or administratively established pseudo-wire domains.

Table of Contents

1. Specification of Requirements

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

2. Acknowledgments

The editors gratefully acknowledge the following contributors: Dimitri Papadimitriou (Alcatel), Peter Busschbach (Lucent), Sasha Vainshtein (Axerra), Richard Spencer (British Telecom), Simon Delord (France Telecom), Deborah Brungard (AT&T), Rahul Aggawal (Juniper), Du Ke (ZTE), Cagatay Buyukkoc (ZTE).

3. Introduction

3.1. Scope

This document specifies requirements for extending pseudo-wires across more than one packet switched network (PSN) domain and more than one PSN tunnel. These pseudo-wires are called multi-segment pseudo-wires (MS-PW). Requirements for single-segment pseudo wires (SS-PW) that extend edge to edge across only one PSN domain are specified in [PWE3-REQ].

This document specifies additional requirements that apply to MS-PWs. These requirements do not apply to PSNs that only support SS-PWs.

3.2. Architecture

The following three figures describe the reference models which are derived from [PWE3-ARCH] to support PW emulated services.

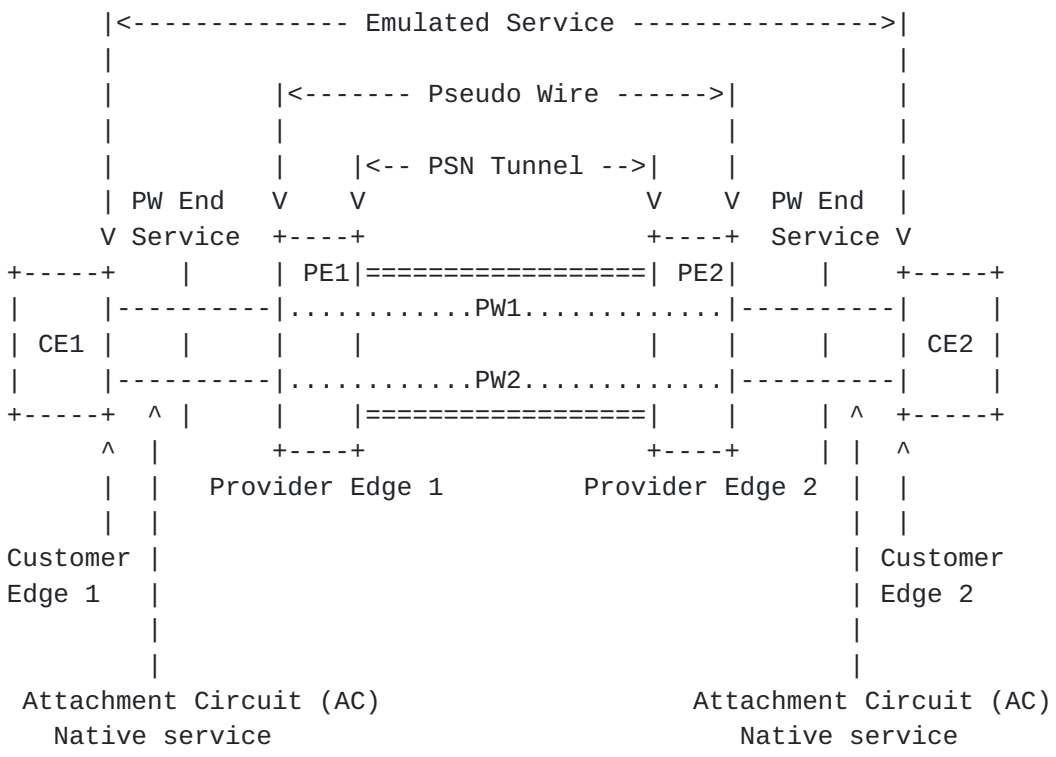

Figure 1: PWE3 Reference Configuration

Figure 1 shows the PWE3 reference architecture [PWE3-ARCH]. This architecture applies to the case where a PSN tunnel extends between two edges of a single PSN domain to transport a PW with endpoints at these edges.

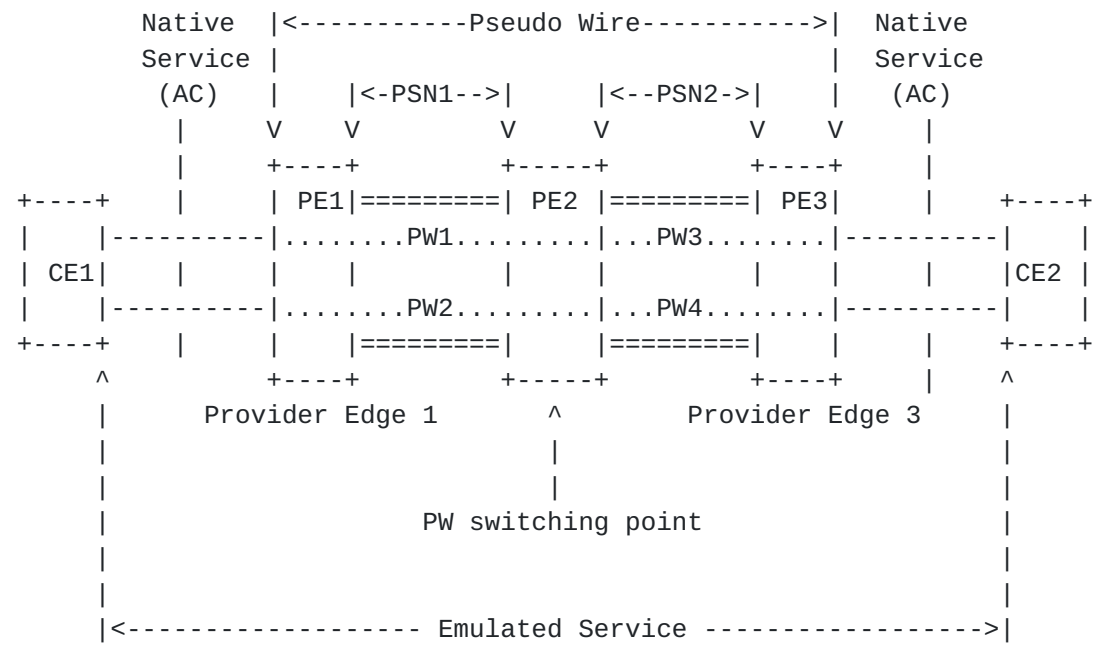

Figure 2: PW switching Reference Model

Figure 2 extends this architecture to show a multi-segment case. PE1 and PE3 provide PWE3 to CE1 and CE2. These PEs reside in different PSNs. A PSN tunnel extends from PE1 to PE2 across PSN1, and a second PSN tunnel extends from PE2 to PE3 across PSN2. PWs are used to connect the Attachment circuits (ACs) attached to PE1 to the corresponding ACs attached to PE3. Each PW on the tunnel across PSN1 is stitched to a PW in the tunnel across PSN2 at PE2 to complete the multi-segment PW (MS-PW) between PE1 and PE3. PE2 is therefore the PW switching point and will be referred to as the PW switching provider edge (S-PE). PW1 and PW3 are segments of the same MS-PW while PW2 and PW4 are segments of another pseudo-wire. PW segments of the same MS-PW (e.g., PW1 and PW3) MAY be of the same PW type or different type, and PSN tunnels (e.g., PSN1 and PSN2) can be the same or different technology. This document requires support for MS-PWs with segments of the same type. An S-PE switches an MS-PW from one segment to another based on the PW identifiers (e.g., PW label in case of MPLS PWs).

Note that although Figure 2 only shows a single S-PE, a PW may transit more one S-PE along its path. For instance, in the multi-provider case shown in Figure 3, there can be an S-PE at the border of one provider domain and another S-PE at the border of the other provider domain. A MS-PW that extends from the edge of one provider (PE1) to the edge of the other provider (PE4) is composed of three segments: (1) PW1, a segment in provider1 network, (2) PW2, a segment between the two borderrouters that are S-PEs, and (3) PWE3, a segment

in the provider2 network.

4. Terminology

[PWE3-REQ] provides terminology for PWE3. This document defines the following additional terms:

- Ultimate PE (U-PE). A PE where the customer-facing ACs (attachment circuits) are bound to a PW forwarder. An ultimate PE is present in the first and last segments of a MS-PW.

- Single-Segment PW (SS-PW). A PW setup directly between two U-PE devices. Each LSP in one direction of a SS-PW traverses one PSN tunnel that connects the two U-PEs.

- Multi-Segment PW (MS-PW). A static or dynamically configured set of two or more contiguous PW segments
  that behave and function as a single point-to-point PW. Each end of a MS-PW by definition MUST terminate on a U-PE.

- PW Switching Provider Edge (S-PE). A PE capable of switching the control and data planes of the preceding and succeeding PW segments in a MS-PW. It is therefore a PW switching point for a MS-PW. A PW Switching Point is never the S-PE and the U-PE for the same MS-PW. A PW switching point runs necessary protocols to setup and manage PW segments with other PW switching points and ultimate PEs.

- PW Segment. A part of a single-segment or multi-segment PW, which is set up between two PE devices, U-PEs and/or S-PEs.

5. Use Cases

PWE3 defines the signaling and encapsulation techniques for establishing SS-PWs between a pair of ultimate PEs and in the vast majority of cases this will be sufficient. MS-PWs may be useful in the following situations:

-i. Inter-Provider PWs: An Inter-Provider PW is a PW that extends from a U-PE in one provider domain to a U-PE in another provider domain.

-ii. It may not be possible, desirable or feasible to establish a direct PW control channel between the ultimate source and destination PEs to setup and maintain PWs. At minimum, a direct PW control channel establishment (e.g., targeted LDP session) requires knowledge of and reachability to the remote U-PE IP address. The local U-PE may not have access to this information due to operational or security constraints. Moreover, a SS-PW would require the existence of a PSN tunnel between the local U-PE and the remote U-PE. It may not be feasible or desirable to extend single, contiguous PSN tunnels between U-PEs in one domain and U-PEs in another domain for security and/or scalability reasons or for the fact that the two domains may be using different PSN technologies.

-iii. MS-PW setup, maintenance and forwarding procedures must satisfy requirements placed by the constraints of a multi-provider environment. An example is the inter-AS L2VPN scenario where the ultimate PEs reside in different provider networks (ASs) and it is the practice to MD5-key all control traffic exchanged between two networks. An MS-PW allows the providers to confine MD5 key administration to just the PW switching points connecting the two domains.

-iv. PSN Internetworking: PWE3 signaling protocols and PSN types may differ in different provider networks. The ultimate PEs may be connected to networks employing different PW signaling and /or PSN protocols. In this case it is not possible to use a SS-PW. A MS-PW with the appropriate interworking performed at the PW switching points can enable PW connectivity between the ultimate PEs in this scenario.

-v. Traffic Engineered PSN Tunnels and bandwidth-managed PWs: There is a requirement to deploy PWs edge to edge in large service provider networks. Such networks typically encompass hundreds or thousands of aggregation devices at the edge, each of which would be a PE. Furthermore, there is a requirement That these PWs have explicit bandwidth guarantees. To satisfy these requirements, the PWs will be tunneled over PSN TE-tunnels with bandwidth constraints. A single segment pseudo-wire architecture would require that a full mesh of PSN TE-tunnels be provisioned to allow PWs to be established between all PEs. Inter provider PWs riding traffic engineered tunnels further add to the number of tunnels that would have to be supported by the PEs and the core network as the total number of PEs increases.

In this environment, there is a requirement either to

support a sparse mesh of PSN TE-tunnels and PW signaling adjacencies, or to partition the network into a number of smaller PWE3 domains. In either case, a PW would have to pass through more than one PSN tunnel hop along its path. An objective is to reduce the number of tunnels that must be supported, and thus the complexity and scalability problem that may arise. The following use case would benefit from this solution.

-vi. Pseudo-Wires in Access and Metro Networks: Service providers are looking to extend PWs to access and metro networks. The prime motive is cost reduction in capital and operation expenses. For instance, in metro networks, providers are looking into extending PWs to next generation SONET ADMs using MPLS mechanisms. The objective is to achieve statistical multiplexing of packets closer to the edge of the network, reducing the recurring costs of SONET circuits or maximizing the utilization of existing SONET infrastructure. In access and metro Ethernet networks, providers are looking to take advantage of MPLS mechanisms to setup point to point Ethernet virtual circuits with endpoints in the same or different access/metro networks.

Using the MS-PW solution, access and metro network elements need only maintain PW signaling adjacencies with the PEs to which they directly connect. They do not need PW signaling adjacencies with every other access and metro network device. PEs in the PSN backbone in turn maintain PW signaling adjacencies among each other. In addition, a PSN tunnel is setup between an access element and the PE to which it connects. Another PSN tunnel needs to be established between every PE pair in the PSN backbone. A MS-PW may be setup from one access network element to another another access element with three segments: (1) access-element - PSN PE, (2) PSN-PE to PSN-PE, and (3) PSN-PE to access element. In this MS-PW setup, access elements are U-PEs while PSN-PEs are S-PEs. It should be noted that the PSN backbone can be also segmented into PWE3 domains resulting in more segments per PW.

6. Multi-Segment Pseudo-Wire Requirements

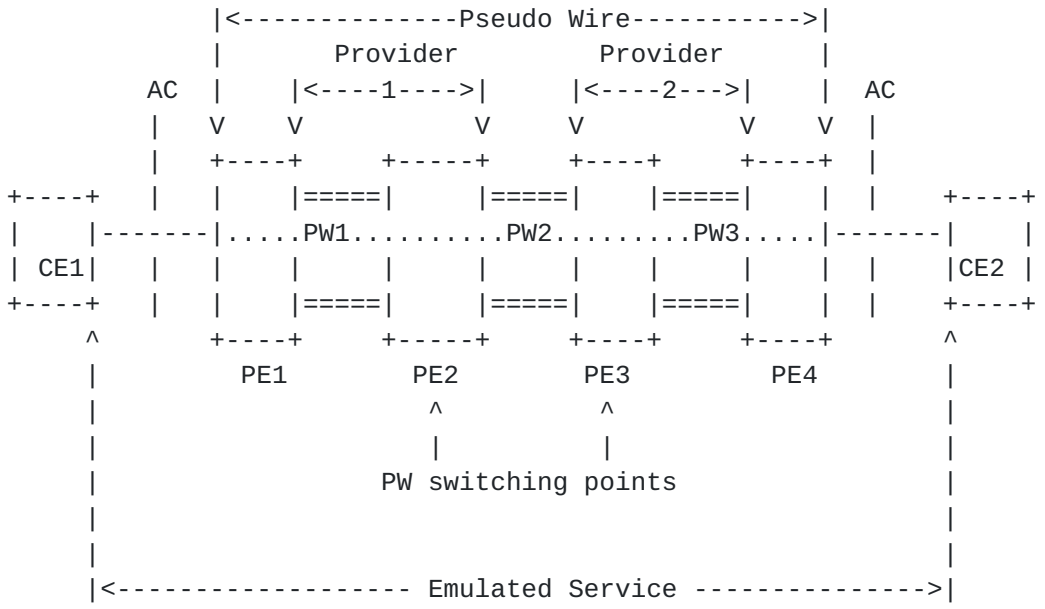

Figure 3: PW switching inter provider Reference Model

6.1. Architecture

The following requirements apply to the architecture for MS-PWs:

-i. S-PEs MAY only support switching PWs of the same PW type. In this case, the PW type is transparent to the S-PE in the forwarding plane, except for functions needed to provide for interworking between different PSN technologies.

-ii. If MS-PWs are tunneled, using a PSN tunnel overlay, across a PSN that only supports SS-PWs, then only the requirements of [PWE3-REQ] apply to that PSN. The fact that the overlay is carrying MS-PWs MUST be transparent to the routers in the PSN.

-iii. The PWs MUST remain transparent to the P-routers. A P-router is not an S-PE or an U-PE from the MS-PW architecture viewpoint. P-routers provide transparent PSN transport for PWs and MUST not have any knowledge of PW traversing them.

-iv. The MS-PWs MUST use the same encapsulation modes specified for SS-PWs.

-v. S-PEs MAY change the type or encapsulation mode of a PW in the NSP function. This enables the establishment of a MS-PW between two PEs with different attachment circuit encapsulation but with the same PW type.

-vi. A MS-PW SHOULD be able to pass across PSNs of all technologies specified by PWE3 for SS-PWs. When crossing from one PSN technology to another, an S-PE must provide the necessary PSN interworking functions in that case.

-vii. MS-PWs are composed of SS-PW, and SS-PW are bi-directional, therefore both directions of a PW segment MUST terminate on the same S-PE/U-PE.

6.2. Resiliency

Mechanisms to protect an MS-PW when the existing path of a MS-PW fails (including S-PE failure or any segment failure) MUST be provided. These mechanisms will depend on the methods of a MS-PW setup. The following are the resiliency requirements:

-i. The ability to configure an end-end backup PW path for a primary PW path MUST be supported. The backup and primary paths should have the ability to traverse separate S-PEs. The backup path MAY be signaled at configuration time or after failure detection.

-ii. The ability to configure a backup PW for a primary PW. The backup and primary PWs should have the ability to traverse separate S-PEs.

-iii. When a MS-PW segment is tunneled over PSN tunnels with fast reroute capability fast re-route events SHOULD be transparent to the MS-PW.

-iv. Automatic Mechanisms to perform a fast switchover from a primary PW to a backup PW upon failure detection SHOULD be provided to minimize packet loss.

-v. Configuration Mechanisms to perform a switchover from a primary PW to a backup PW upon failure detection SHOULD be provided.

-vi. A mechanism to automatically revert to a primary PW from a backup PW MAY be provided. When provided, it SHOULD be enabled/disabled by configuration.

-vii. Mechanisms for PW segment failure detection and notification to other segments of a MS-PW MUST be provided.

6.3. Quality of Service and Class of Service

Pseudo-wires are intended to support emulated services (e.g., TDM and ATM) which may have strict per-connection quality of service requirements. This may include either absolute or relative guarantees on packet loss, delay, and jitter. These guarantees are in part delivered by reserving sufficient network resources (e.g. BW), and by providing appropriate per-packet treatment (e.g. scheduling priority and drop precedence) throughout the network.

In SS-PWs, a traffic engineered PSN tunnel (i.e., MPLS-TE) may be used to ensure that sufficient resources are reserved in the P-routers to provide QoS to PWs on the tunnel. In this case, the ability to automatically manage the PSN tunnel resources (e.g. admission control of PWs onto the PSN tunnel) is a requirement at each PW segment. The ability to associate the appropriate QoS class with each PW PDU is a strict requirement at each router transited in the network.

For MS-PWs, each S-PE maps a PW to a PSN tunnel. That is, an S-PE decides what PW to bind to which PSN tunnel. Control over binding a PW to a specific PSN tunnel SHOULD be provided as a matter of policy configuration.

When the U-PE attempts to signal a PW the following capability is required:

-i. Admission control to the PSN tunnel needs to be performed against available resources. PW admission control into a PSN tunnel MUST be configurable.

-ii. A per PW/QoS class setup priority should be provided.

-iii. In case the PSN tunnel lacks the resources necessary to accommodate the new PW, a PW setup failure message MUST be returned and the PW MUST fail to setup. Alternatively: In case the PSN tunnel lacks the resources necessary to accommodate the new PW, an attempt to signal an new PSN tunnel, or increase the capacity of the existing PSN tunnel MUST be made. If the expanded PSN tunnels fails to setup the PW MUST fail to setup.

-iv. PW traffic parameter representations MUST be the same for all types of PW.

-v. The PW signaling MUST enable separate traffic parameter values to be specified for the forward and reverse directions of the PW.

6.3.1. Traffic Engineering

The following requirements apply to the traffic engineering of MS-PWs:

-i. When setting up a MS-PW, S-PEs and U-PEs MUST be able to select a tunnel across the PSN in such a way as to satisfy the MS-PW QoS and bandwidth requirements. Restrictions may apply depending on the method used for setting up a PW segment and on PSN tunnel types.

-ii. Upon setting up a MS-PW for which QoS/bandwidth commitments are made over the PSN, S-PEs and U-PEs SHOULD be able to perform admission control for each PW segment over a PSN tunnel to ensure that PW bandwidth and QoS requirements can be satisfied.

6.4. MS-PW Setup Mechanisms.

The MS-PW setup mechanisms MUST accommodate Service Provider's practices, especially in relation to security and confidentiality and traffic engineering. Security and confidentiality are especially important when the MS-PW's are setup across ASs in different administrative domains.

There are four different SS-PW signaling protocols that are defined to signal PWs:

-i. Static configuration of the SS-PW (MPLS or L2TPv3).

-ii. LDP using PWid FEC 128

-iii. LDP using the generalized PW FEC 129

-iv. L2TPv3

Any combination of these signaling protocols MUST be supported.

Following are further requirements on MS-PW setup mechanisms:

-i. Static S-PE selection and PSN tunnel selection MUST be provided.

-ii. The MS-PW path MUST have the ability to be dynamically setup between the U-PEs by provisioning only the U-PEs.

-iii. Dynamic MS-PW pseudowire setup requires that a unique identifier be associated with a PW and be carried in the signaling message. That identifier must contain sufficient information to determine the path to the remote U-PE through intermediate S-PEs.

-iv. In a single-provider domain it is natural to have the U-PE identified by one of its IP addresses. This may also apply when a MS-PW is setup across multiple domains operated by the same provider. However, some service providers have security and confidentiality policies that prevent them from advertising reachability to routers in their networks to other providers (reachability to an ASBR is an exception). Thus, procedures MUST be provided to allow dynamic setup of MS-PWs under these conditions.

-v. Addressing of MS-PW end point at the U-PE MUST be independent of the IP address of the U-PEs themselves.

-vi. Solutions MUST strive to minimize the amount of configuration needed to setup an MS-PW.

-vii. MS-PW signaling procedures MUST define clear rules for triggering the setup of segments of a MS-PW.

-viii. The signaling procedures MUST be defined such that the setup of a MS-PW is considered successful if and only if all segments of the MS-PW are successfully setup.

-ix. Mechanisms MUST be developed to propagate setup explicit failure indications to the S-PEs and U-PEs associated with the failed MS-PW.

6.4.1. Routing

An objective of MS-PWs is to provide support for the following connectivity:

-i. MS-PW MUST be able to traverse multiple IGP domains.
-ii. MS-PW MUST be able to traverse multiple autonomous systems within the same administrative domain.
-iii. MS-PW MUST be able to traverse multiple autonomous systems belonging to different administrative domains.
-iv. MS-PW MUST be able to terminate, as well, in any of above mentioned domains.
-v. MS-PWs MUST be able to support any hybrid combination of the aforementioned connectivity scenarios.

The routing function MUST support the various MS-PW setup methods and the various connectivity scenarios. Following are the consequent requirements:

-i. MUST support the setup of a statically configured segment of a MS-PW. (static S-PE selection)
-ii. The MS-PW MUST have the ability to automatically select the S-PEs along the MS-PW path. Some of the S-PEs MAY be statically selected.
-iii. The PW Routing function MUST support dynamic re-routing around failure points when segments are setup using the dynamic setup method.
-iv. The PW Routing function MUST support re-routing around failures that occur between the statically configured segment endpoints. This may be done by choosing another PSN tunnel between the two segment endpoints or setting up an alternative tunnel.
-v. Routing MUST support the operation of backup protection of primary paths.
-vi. Manual routing SHOULD be supported to allow diversity for 1:1 protected PWs.

7. Operations and Maintenance (OAM)

OAM mechanisms for the attachment circuits are defined in the specifications for PW emulated specific technologies (e.g., ITU-T I.610 for ATM). These mechanisms enable, among other things, defects in the network to be detected, localized and diagnosed. They also enable communication of PW defect states on the PW attachment circuit.

The interworking of OAM mechanisms for SS-PWs between ACs and PWs is

defined in [PWE3-OAM]. These enable defect states to be propagated across a PWE3 network following the failure and recovery from faults.

OAM mechanisms for MS-PWs MUST provide at least the same capabilities as those for SS-PWs.

In addition, it should be possible to support both segment and end-to-end OAM mechanisms for both defect notifications and connectivity verification in order to allow defects to be localized in a multi-segment network. That is, PW OAM segments can be U-PE to U-PE, U-PE to S-PE, or S-PE to S-PE.

It is desirable to use existing PW OAM mechanisms for MS-PWs or extensions to them if needed.

The following requirements apply to OAM for MS-PWs:

-i. MS-PW OAM SHOULD be supported end-to-end across the network.

-ii. OAM activation/deactivation SHOULD be tied to MS-PW setup/tear down.

-iii. The MS-PW SHOULD support a Forward Defect Indicator (FDI).

-iv. Single ended monitoring SHOULD be supported for both directions of the MS-PW.

-v. MS-PW OAM SHOULD support switch over between 1:1 protected LSPs end-to-end.

-vi. In-band monitoring SHOULD be provided both end-to-end across the MS-PW, and on a segment (i.e. SS-PW) basis.

-vii. At the S-PE, defect notifications on the upstream segment of PWs must be propagated to the downstream PW segment.

-viii. All PE routers along the MS-PW MUST agree on a common PW OAM mechanism to use to the MS-PW.

-ix. At the S-PE, defects on an PSN tunnel must be propagated to all PWs that utilize that particular PSN tunnel.

-x. The directionality of defect notifications must be maintained across the S-PE.

-xi. The S-PE should be able to behave as a segment endpoint for PW OAM mechanisms.

-xii. The S-PE MUST be able to pass U-PE to U-PE PW OAM messages transparently.

8. Security Considerations

Section to be completed later. Editor's note: This section needs extensive work. Security requirements are needed for inter-as , and inter -providers situations.

9. Full Copyright Statement

10. Intellectual Property Statement

11. IANA Considerations

This document has no IANA Actions.

12. Normative References

[PWE3-ARCH] "PWE3 Architecture" Bryant, et al., RFC3985.

[PWE3-OAM] "Pseudo Wire (PW) OAM Message Mapping" Nadeau et al., draft-ietf-pwe3-oam-msg-map-01.txt (work in progress), October 2004

13. Informative References

[ITUQ] ITU-T Recommendation Q.933, and Q.922 Specification for Frame Mode Basic call control, ITU Geneva 1995

14. Author Information

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

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

Nabil Bitar
Verizon
40 Sylvan Road
Waltham, MA 02145
e-mail: nabil.bitar@verizon.com