DECLARATION

I, Alexa Morris, based on my personal knowledge and information, hereby declare as follows:

1. I am Managing Director of the IETF Administration LLC and have held that position since the LLC was formed in August 2018. Prior to that, starting on January 1, 2008, I was the Executive Director of the Internet Engineering Task Force, which was an activity of the Internet Society. Since the business of IETF did not change in any materially relevant manner with the formation of the LLC, I will collectively refer to both the activity and the LLC as IETF.

2. One of my responsibilities with IETF has been to act as the custodian of Internet-Drafts and records relating to Internet-Drafts. I am familiar with the record keeping practices relating to Internet-Drafts, including the creation and maintenance of such records.

3. I hereby declare that all statements made herein are of my own knowledge and information contained in the business records of IETF and are true, and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements may be punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

4. If depositions regarding the information in this declaration are required, the deposition should be taken by phone or videoconference or, if it must be in person, should be in California.

5. Since 1998, it has been the regular practice of the IETF to publish Internet-Drafts and make them available to the public on its website at www.ietf.org (the IETF website). The IETF maintains copies of Internet-Drafts in the ordinary course of its regularly conducted activities.

6. Any Internet-Draft published on the IETF website was reasonably accessible to the public and was disseminated or otherwise available to the extent that persons interested and ordinarily skilled in the subject matter or art exercising reasonable diligence could have located it. In particular, the Internet-Drafts were indexed and searchable on the IETF website.

7. Internet-Drafts are posted to an IETF online directory. When an Internet-Draft is published, an announcement of its publication that describes the Internet-Draft is disseminated. Typically, that dated announcement is made within 24 hours of the publication of the Internet-Draft. The announcement is kept in the IETF email archive and the date is affixed automatically.

8. The records of posting the Internet-Drafts in the IETF online repository are kept in the course of the IETF's regularly conducted activity and ordinary course of business. The records are made pursuant to established procedures and are relied upon by the IETF in the performance of its functions.

9. It is the regular practice of the IETF to make and keep the records in the online repository.

10. Exhibit 1 is a true and correct copy of an announcement of the publication of draft-pan-pwe3-protection-03.txt, titled "Pan Wire Protection." I have determined that the Internet-Draft was posted to the online repository in July 2006. Therefore, based on the normal practice of the IETF, that Internet-Draft was reasonably available to the public within 24 hours of posting. At that time, the Internet-Draft would have been disseminated or otherwise available to the extent that persons interested and ordinarily skilled in the subject matter or art, exercising reasonable diligence, could have located it.

Pursuant to Section 1746 of Title 28 of United States Code, I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct and that the foregoing is based upon personal knowledge and information and is believed to be true.

Date: January 3, 2024 By: Alexa Morris

Alexa Morris

4892-9810-3450

PWE3 Working Group
Internet Draft
P. Pan

(Hammerhead Systems)
Document: draft-pan-pwe3-protection-03.txt
M. Bocci

Mustapha Aissaoui

(Alcatel)

Florin Balus

Hamid Ould-Brahim

(Nortel)
Expires: December 2006
July 2006

Pseudo Wire Protection

Status of this Memo

Abstract

This document describes a mechanism that helps to protect and recover
user traffic when carried over pseudo-wires. The mechanism requires

some minor modification to the existing pseudo-wire setup procedure,
and is fully backward compatible.

The proposed mechanism allows the network operators to setup one or
multiple backup pseudo-wires to protect a working pseudo-wire. Upon
network failure, user traffic can be switched over to the next "best"
pseudo-wire base on preference levels.

This document first describes the motivation of the work base on the
discussions with a number of carriers. Then we define the protocol
extension itself.

Conventions used in this document

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

Table of Contents

1. Terminology

The reader is assumed to be familiar with the terminology in [LDP], [PW-CTRL] and [MHOP-PW]. The new terms are the following:

- Working Pseudo-wire: A pseudo-wire that carries user traffic, and may be protected by one or multiple associated backup pseudo-wires.

- Backup Pseudo-wire: A pseudo-wire that is used to re-route user traffic from a working pseudo-wire at head-end.

2. Introduction

Pseudo-wires have been deployed by a number of networks to carry customer layer-2 data traffic. Each Layer-2 data flow (or Attachment Circuit) is mapped to a pseudo-wire. Pseudo-wire setup, maintenance and packet encapsulation have been extensively described in a number of IETF PWE3 drafts [PWE3-CTRL, PWE3-TRANSPORT]. Recently, several carriers have requested that, when offered as a service, pseudo-wires need to possess the same protection and redundancy capabilities that have been deployed in transport networks.

In this draft, we extend the LDP pseudo-wire proposal [PWE3-CTRL] to support protection and restoration operation.

Why is such work necessary?

When it comes to traffic protection, the carriers need to ensure traffic protection on every network segment and in every layer of the network. Just because most of the pseudo-wire traffic will go through MPLS LSP's, we cannot therefore make the assumption that user traffic will be protected via MPLS Fast-Reroute [MPLS-FRR] or RSVP path protection. In the MHOP-PW scenario, even if Tunnel Protection schemes are present in individual PW domains, PW protection against S-PE failures is required.

There are a number of the deployment scenarios where pseudo-wire protection can be critically important:

2.1. Access Networks

Pseudo-wire has been in deployment for multi-service data access. One reason is that pseudo-wire enables data aggregation, which in turn improves bandwidth utilization. In a typical metro network access location (Hub or CO), the statistical multiplexing gain is approximately 3-4 [ATT-REPORT]. The earlier user flows get aggregated, the better bandwidth utilization will be gained by the carriers, especially at the access locations where bandwidth is still expensive.

More importantly, pseudo-wire provides a common data transport layer, where all layer-2 packets can be processed uniformly at provider edge. This enables the carriers to migrate from the traditional layer-2 (ATM or Frame Relay) circuits into high-speed Ethernet without service distraction. A common deployment scenario can be shown as the following:

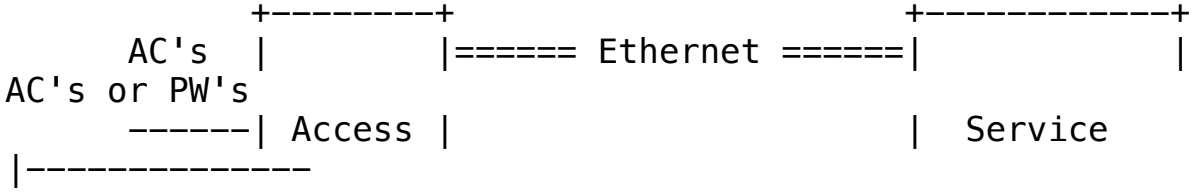

Figure-1: Pseudo-wire network access

Note that, given the size of access networks, the cost of access device and access link management are some of the key deployment considerations, such that the access devices may not be IP routers, and the extensive IP routing and MPLS signaling (such as RSVP-TE) may be not applied in this part of the network.

In this part of the network, one method may be to run pseudo-wires over the access links, and conduct traffic protection at per-pseudo-wire level.

2.2. Metro Networks

First of all, many of the MPLS-enabled metro networks today do not operate with RSVP-TE, which MPLS Fast-Reroute is based on. Secondly, many of the metro networks have already deployed pseudo-wires in one form or another (such as VPLS). Thus, pseudo-wire traffic protection becomes vital.

Another issue is that given the heterogeneous nature and subsequent

complexity in network topology, the metro networks may not be able to guarantee parallel MPLS tunnels between two edge nodes with the same bandwidth. In this case, pseudo-wire protection may be the only method for user traffic.

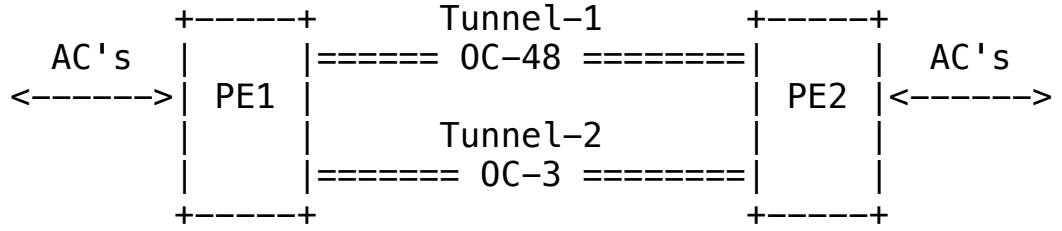

Figure-2: Bandwidth Mismatch

In Figure-2, there exist two parallel tunnels (LSP's) between two PE's with different link capacity. Whenever the bandwidth on a protecting link is smaller than that on the working link, we may run into trouble during protection and restoration.

In the example, let's assume that both tunnels are MPLS LSP's. Network operators have enabled MPLS fast-reroute to enable both LSP's protecting each other. From the PE's, a number of AC's are aggregated into the LSP's as pseudo-wires. Some AC's carry mission-critical data, while others transport best-effort data. If Tunnel-1 fails, all traffic on Tunnel-1 will be switched into Tunnel-2. However, since both tunnels have different bandwidth, mission-critical traffic could be dropped or delayed as a result of link congestion during switch-over.

This problem can be easily resolved if each pseudo-wire has its own preference, which allows the pseudo-wires to preempt each other when it becomes necessary. Also note that, since the pseudo-wires are always bi-directional, the preference assignment must be consistent on both ends of the pseudo-wires.

2.3. MS-PWs use cases

Multi-segment pseudo-wire [MS-ARCH, MHOP-PW, Segmented-PW] has gained much traction in carrier networks recently. It allows pseudo-wire traffic to transport over multiple PW Domains (Access/ Core Metro/WAN or different provider networks).

Within each network, the type of the PSN tunnels may be different. And there is no guarantee that the PSN tunnels within each network or over the inter-provider links will be protected. Also any failure of the S-PE nodes can not be addressed by the tunnel protection, and

therefore has to be addressed by a PW protection scheme. The head-end nodes (T-PE's) may use PW OAM (VCCV/PW Status TLV) to detect any network failure that may affect the pseudo-wires, and can reroute user traffic on a per-pseudo-wire basis.

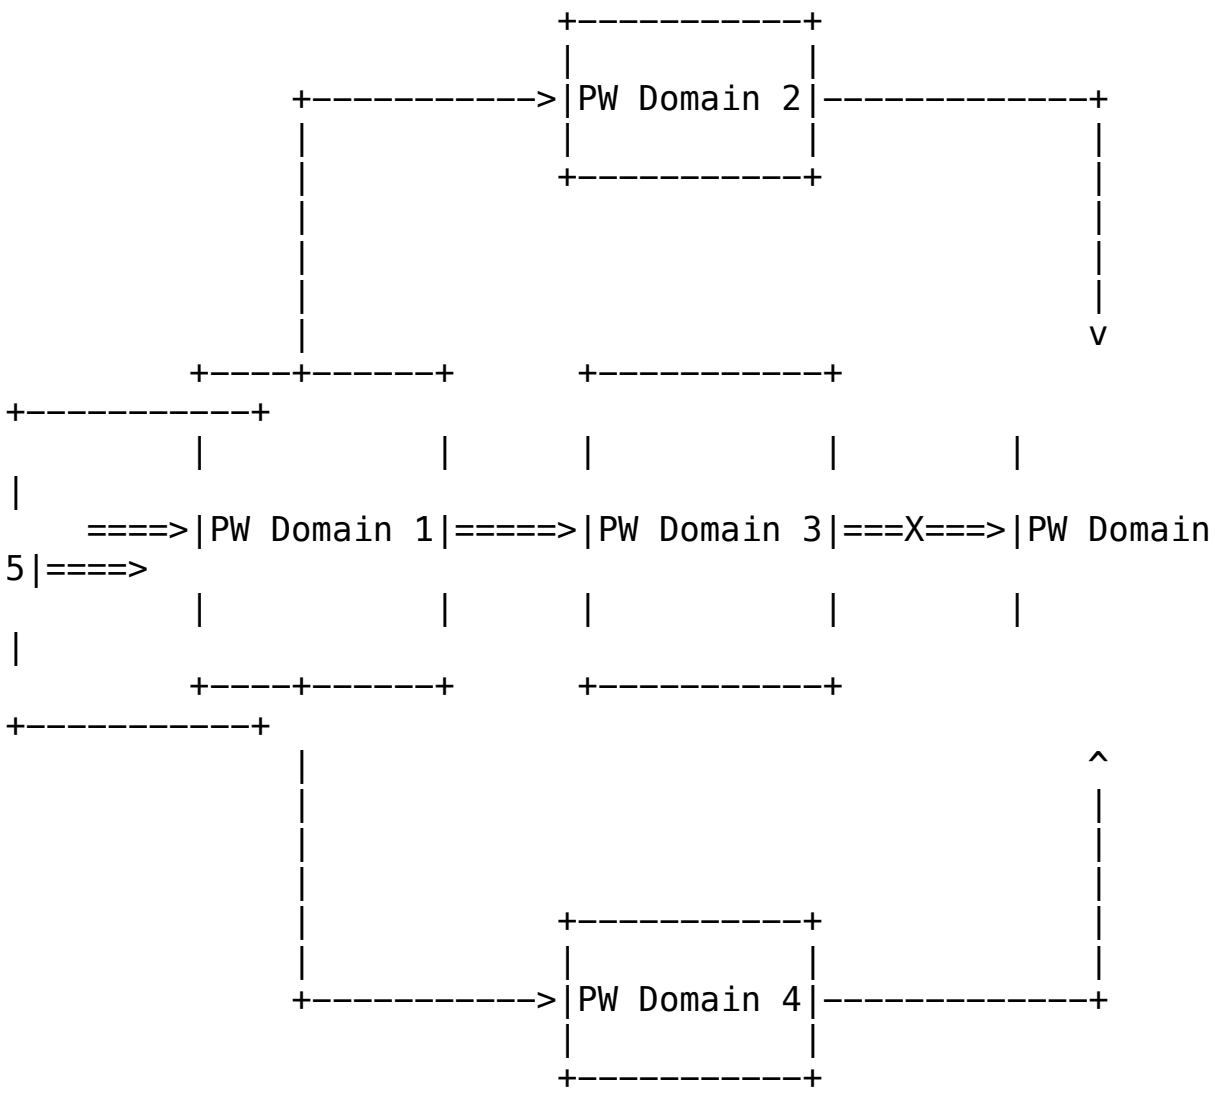

Figure-3: Multi-segment PW in PW Domain environment

For example, in Figure-3, a multi-hop pseudo-wire traverses through
Pw Domain 1, 3 and 5. Say, the link or the S-PEís between PW Domain 3
and 5 has failed. From the head-end the pseudo-wire can be re-routed
through PW Domains 2 or 4.

2.4. Planned Traffic Switch-over

Finally, the network operators need to have the ability to support planned traffic shifting. In Figure-2, there are two links between two PE's carrying a number of pseudo-wires. During network maintenance, carriers may decide to shift all traffic from a set of pseudo-wires from one link to another temporally without causing traffic disturbance to users. To support this operation, pseudo-wire protection can be manually triggered from the operators [NOTE1].

3. Design Considerations

3.1. Signaling a Backup Pseudo-wire

When operating in multi-domain environment, the working and backup pseudo-wires may arrive on the same PE nodes (S-PE's).

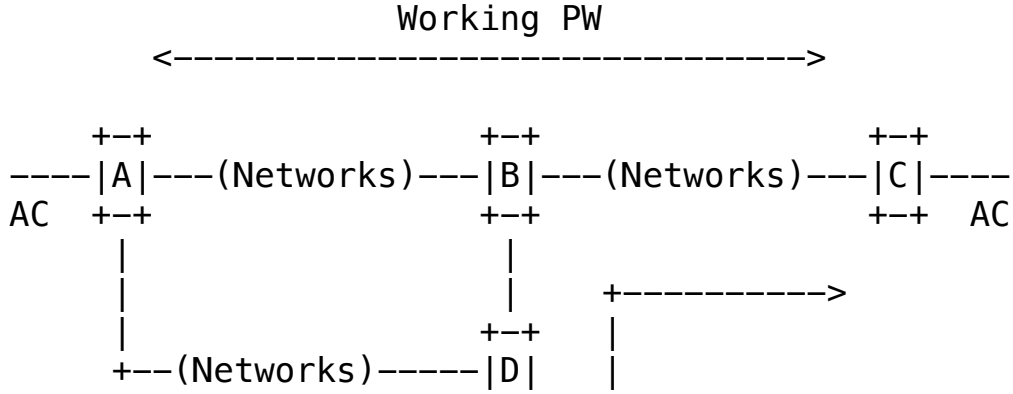

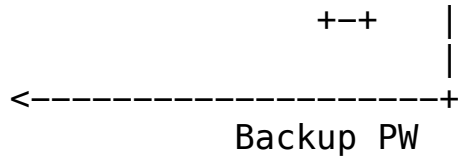

Figure-4: Why we need to identify backup PWís

As shown in Figure-4, the Working PW goes through nodes A, B and C, while a Backup PW may take a different route, A, D, B and C. Node B is the S-PE that would receive Label Mapping messages for both working and backup pseudo-wires. Unless B knows the difference between the working and backup pseudo-wires, it may mistaken that there has been a route change within the network, and thereby stop the processing of the pseudo-wires.

To make the message processing possible, the backup pseudo-wires must at least satisfy the following criteria:

1. Unambiguously and uniquely identifying the backup pseudo-wire

2. Unambiguously associating working PW with their backups.

Pseudo-wires can be identified via either FEC 128 (PWid) or FEC 129 (Generalized FEC). In latter case, each pseudo-wire can be uniquely identified as a pair of <AGI> and <AII> [PW Control]. Since there are a number of limitations in using FEC 128 in multi-hop environment, we will support pseudo-wire protection with FEC 129 only.

There are a number of options in making the backup pseudo-wires

unique:

1. Assign a new <AII> for each backup pseudo-wire: To make the association of working and backup pseudo-wires at T-PE's, we may put some grouping information inside the <AII>. For example, we may use the first two bytes of the Global ID field in AII Type 2 [MHOP-PW] as the protection group ID. However, this will require the format change in all AII's, and cause potential backward compatibility issues.

2. Assign a new <AGI> for all the working and backup pseudo-wires: When used in L2VPN's, the <AGI> is used as "VPN ID" [L2VPN], which has an entirely different meaning from the pseudo-wire protection grouping. We will elaborate this further below.

3. Use one bit somewhere in the PW FEC to distinguish working and protecting pseudo-wires: However, the operators may choose multiple backup pseudo-wires to protect one working pseudo-wire. In this case, one bit would not be sufficient.

In our design, we will use an opaque "Protection TLV", in which each

working and backup pseudo-wires will have a different identification (or reference ID). All working and backup pseudo-wires will have the same <AGI> and <AII>. At pseudo-wire setup time, each working and backup pseudo-wires will get its own MPLS labels for packet forwarding.

3.2. Determination of Protection Path

RSVP-TE messages uses Explicit Routing Object (ERO) to setup the LSP's. CR-LDP [RFC3212], section 4.1 has also defined an Explicit Route TLV to achieve the same purpose. One key advantage in using explicit routes is that it provides a simple solution to ensure the working and backup pseudo-wires do not traverse through the same routes (i.e. no fate-sharing).

However, when operating in multi-domain environment, the carriers may not want to share network resource information among each other. In this case, there is no need to specify for each step the explicit routing information during pseudo-wire setup.

In our design, by default, we do not require the use of explicit routes during working and backup pseudo-wire setup. Instead, we rely on the intermediate nodes (S-PE's) to provide the best possible routes for the pseudo-wires.

However, it is important to realize that the edge node (T-PEís) may have the capability of interfacing with either multi-

lateral policy
servers, or MP-BGP, and obtain the exact inter-domain routing
information for backup pseudo-wires. For example, it is possible that

the PE nodes distribute the protection information via MP-BGP as a
part of L2VPN setup sequence. Such mechanism can ensure that the
working and protection pseudo-wires will not traverse through the
same set of PSN tunnels.

The exact mechanism in obtaining inter-domain protection path
information is outside the scope of this draft.

3.3. Protection Schemes

Typically, there is three types of point-to-point protection in
telecommunication networks: 1+1, 1:1 and 1:N.

1+1 is to transmit same traffic over two parallel links. The receiver
will only pick traffic from one link at any given time. In event of
failure, at least one of the links still carries the actual traffic.
For example, SONET UPSR (Unidirectional Path-Switching Rings) is an
implementation of 1+1 protection. However, this may not be desirable
in many networks, if the protection path consumes

network resources
   such as link bandwidth.

   1:1 protection is to use one connection to protect another
   connection. When a failure in working connection has been detected,
   the network node would switch traffic to an alternative or backup
   connection. The most popular 1:1 protection is SONET APS. The
   efficiency of 1:1 protection is sometimes measured in terms of
   switch-over time.

   1:N is a generalized version of 1:1. In 1:N, one connection is
   established to protection multiple working connections. MPLS Facility
   Backup is one such example. One of its key advantages is that it
   introduces less number of states that intermediate nodes have to
   manage.

   In pseudo-wire protection however, each AC may have its own layer-2
   characteristics that need to be maintained separately. Thus, it may
   be difficult to apply 1:N protection. For example, it is not clear
   that it is feasible or reasonable to setup a single backup pseudo-
   wire to protect best-effort Ethernet VLAN connections plus ATM SPVCís
   with CBR and VBR traffic requirements.

   In our design, we shall only support 1:1 protection.

3.4. Protection Types

   Pseudo-wire protection will support the following types: cold, warm

and hot standby.

. Cold Standby

This is a common method in optical transport network, where the nodes will only negotiate and establish backup pseudo-wires after the detection of network failure.

This type of protection can be implemented with the existing specification [PW-CTRL, MHOP-PW]. Upon the detection of network failure, the PE nodes will re-negotiate another pseudo-wire, and transmit packets over. The protection effectiveness depends on how fast two edge nodes can react to network failure and process control messages after the failure.

. Warm Standby

The edge nodes will negotiate backup pseudo-wires and exchange labels prior to any network failure. However, data forwarding path will not be programmed for label processing and QoS enforcement until after the detection of network failures.

Such practice and requirement come from traditional transport

carriers. In SONET/SDH networks, switches reserve the protection time slots ahead of time. Upon the detection of network failure, the nodes "wake-up" the protection connections.

. Hot Standby

This is the most efficient protection method. The protecting pseudo-wires are established before any network failure. This is also known as "make-before-break". Upon the detection of network failure, the edge nodes will switch data traffic into pre-established backup pseudo-wires directly. The protection efficiency is therefore depending on the speed for failure detection and switch-over, the latter being in the order of milliseconds.

This is the default operation in our proposal.

3.5. Pseudo-wire Preemption

Pseudo-wires are not created equal. In case of network failure or congestion, the ones that carry important user traffic should get better treatment than those of best-effort data traffic. Thus, the operator should have the ability to assign preference levels to the pseudo-wires. During network failure or congestion, the PEís (both T-

PEís and S-PEís) can preempt less important pseudo-wires and make room for more important traffic.

Itís important to realize that we have made the above design decision base on the following assumptions:

First, network failure and traffic congestion wonít last for a long period of time in carrier networks. As such, assigning preference levels to pseudo-wires would provide temporary congestion avoidance to important user traffic.

Secondly, pseudo-wire preemption operation takes place among individual pseudo-wires, and does not necessarily involve the use of working and protection pseudo-wires.

3.6. Backup Pseudo-wire Priorities

The operators may establish multiple backup pseudo-wires for a particular working pseudo-wire. When the working pseudo-wire is experiencing network failure, its traffic will be switched onto one of the backup pseudo-wires. This requires the tail-end T-PEís and S-PEís to understand which backup pseudo-wire to use in case of failure.

In this proposal, the pseudo-wire initiator will always

assign a
priority level to each of the backup pseudo-wires. During switch-
over, user traffic will be switched onto the backup pseudo-wire with
the highest priority level.

4. Backup Pseudo-wire Extension

Setting up backup pseudo-wires is based on [PW-CTRL], [LDP] and
[MHOP-PW]. PW label binding uses targeted LDP, where two edge nodes
first establish an LDP session using the Extended Discovery mechanism
described in [LDP]. PW's are initiated via LDP Label Mapping
messages. Each message contains a FEC TLV, a Label TLV, and some
optional TLVís.

We only support the Generalized ID FEC (129) during the proposed
operation.

PW protection operates under the assumption that there exists more
than one route between a pair of PE's to transport data traffic, as
shown in Figure-5.

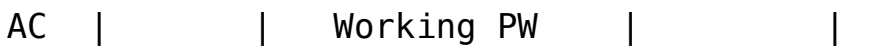

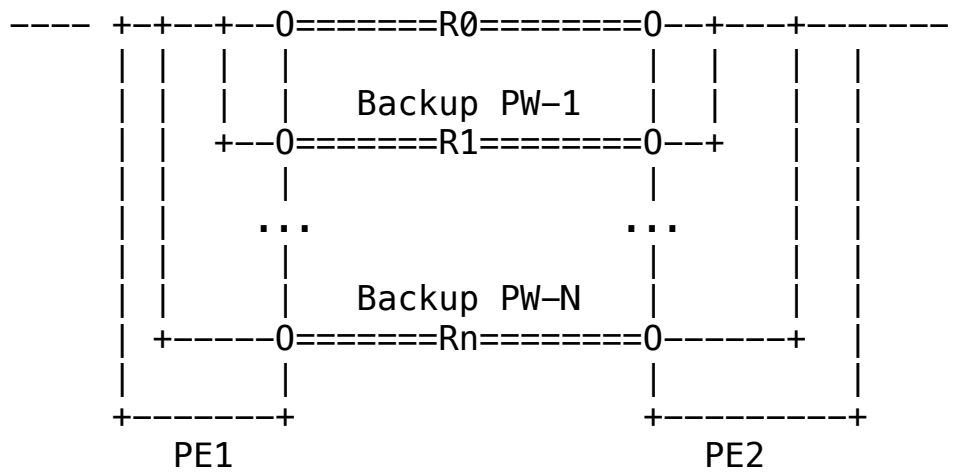

Figure-5: PW Protection Example

For each working PW, the PEs may setup one or multiple backup PWís.
The procedure for establishing backup PWís is the same as the one for
regular PWs [PW-CTRL, MHOP-PW]. The only difference is that during
backup PW initiation, a Protection TLV will be included in the
mapping messages. The Protection TLV includes, among a number of
other parameters, a reference ID and a backup priority level.

The working pseudo-wire MUST not attach the Protection TLV.

The Label Mapping messages for backup Pseudowires may be sent over
multiple routes between two PEís. In case of multi-segment PW, the
messages may be processed at multiple provider edge nodes, which will
rely on the reference IDís to distinguish each of the backup PWís.

The working and backup PWís must have the same attachment circuit
information. During network failure, the T-PEís will switch user

traffic into the backup PWís that has the highest backup priority level. After network recovery, the PE's will revert back to the working PW based on configurable behavior (immediately or during a maintenance window).

4.1. The PROTECTION TLV

With the Protection TLV, the operator can configure the protection mechanism that they prefer. Since the pseudo-wires are always bidirectional, exchanging protection information between PE nodes will help to achieve a consistent protection behavior for pseudo-wires.

The Protection TLV has the following format:

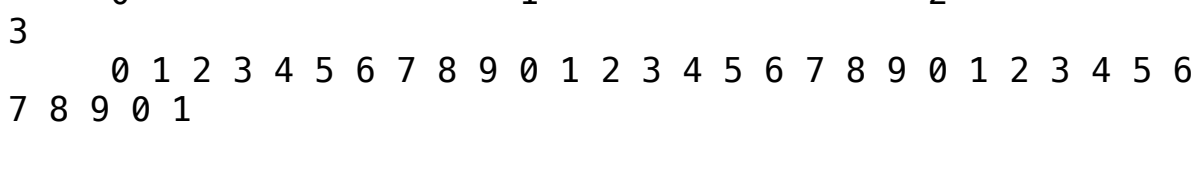

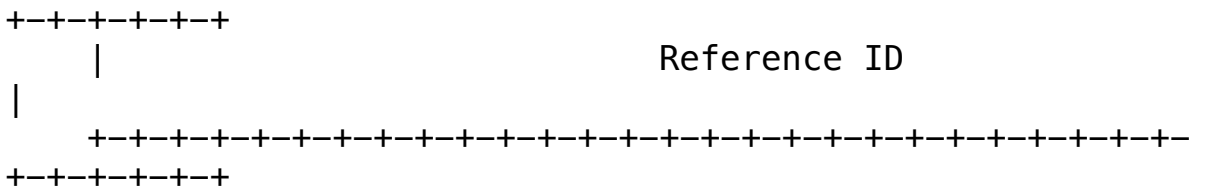

- U bit (always clear)

  On the nodes that do not support the PROTECTION TLV, the notification messages will be replied to the originator, and the entire backup PW message will be ignored.

- Protection tlv

  The value of the new tlv type needs to be allocated by IANA.

- Flags

  This field contains the protection information used by the intermediate PEís (S-PE's) during MHOP-PW operation.

  Currently, it has the following flags:

  0x01 F-flag: Fate Sharing Allowed
  0x02 B-flag: Bandwidth CAC Required

  When the F-flag is set, the working and backup pseudo-wires may share the same routes in the network when necessary. By default, this flag is set.

  When the B-flag is set, the S-PE's must perform CAC on the backup pseudo-wires. Otherwise, the S-PE's can send a notification message to the originator, and continue on with the backup PW setup. By default, this flag is set.

- Protection Type

Currently we have defined the following values:

Hot Standby: 0
Warm Standby: 1

The default value is 0 (Hot Standby). More information is in Section 3.4.

- Backup Priority

This is the priority level with respect to a backup PW. The value of 0 is the highest. When the traffic on a working PW needs to be switch-over, if there are multiple backup PWís, the one with the highest backup priority will be used to carry traffic.

- Reference ID

This is assigned by the pseudo-wire originating nodes (T-PE's). For all the backup pseudo-wires that protect the same working pseudo-wire, they must have different IDís. The reasoning is in Section 3.1.

4.2. Head-end PE Operation

As illustrated in Figure-5, the operator can first initiate the Working PW over route R0, and then initiate the Backup PW-1 over route R1, the Backup PW-2 over route R2, and so on and so forth.

The Label Mapping messages for both working and backup PW's must have the same Generalized ID FEC (that is, the same <AGI>, <AII> and AC interface data).

Each backup PWís must carry a Protection TLV with a different Reference ID and Backup Priority.

The head-end PE's (T-PE's) should not initiate the backup PW's until the working PW is up and running.

Further, the T-PE's should keep track of the PW-SW-POINT TLV [Segmented-PW] for both working and backup pseudo-wires. The PW-SW-POINT TLV has the information on the intermediate hops that the PW's have traversed. For the backup PW's that do not allow fate-sharing, their PW-SW-POINT TLV should not over-lap with the working PW.

For the backup PW's that do not need bandwidth guarantee, it does not need to carry the PW Bandwidth TLV during setup, and the B-Flag must always be off. Otherwise, the backup PW's must carry the same PW Bandwidth TLV as in the working PW.

4.3. Switched PE Operation

In case of multi-hop PW's, the intermediate PE's (S-

PEís) will
perform the following checks when receiving a Label Mapping message:

If it does not support the Protection TLV, it will reject the
messages, and notify the head-end PE directly with ìDonít support PW
Protectionî status code.

If it supports the Protection TLV and the message consists of a
Protection TLV, the S-PE will compare the Reference ID on the PW's
that share the same <AGI> and <AII>. If there is an entry with the
same Reference ID, the Label Message will be rejected with
ìDuplicated Reference ID" status code.

If the new backup PW has a backup priority that already exists, the
request will be rejected with ìMismatched Backup Priority" status
code.

Otherwise, the S-PE will interface with either static or dynamic
(i.e. BGP) routing tables, and place the backup PW's on a next-hop
route that is different from the working PW.

If the F-flag (Fate Sharing Flag) is set, and the S-PE

cannot find an
alternative next-hop, the backup PW will go through the same route as
the working PW. If the flag is clear, the S-PE will terminate the
backup PW setup, and reject the Label Mapping message with ìWorking
and backup PW's share the same fate" status code.

If the B-flag (Bandwidth CAC Required) is set, and the S-PE cannot
reserve resource on the out-going link, the label mapping message
will be rejected with ìOut of Backup Resource" status code.

4.4. Tail-end PW Operation

As shown in Figure-5, when PE2 receives a Label Mapping message, it
will perform the following checks:

If PE2 does not support the Protection TLV, it will reject the new
request with ìDonít support PW Protectionî status code.

If PE2 supports the Protection TLV, it will process the rest of the
mapping message. PE2 needs to check if it already has the PW's with
the same attachment ID (PWid or the combination of AGI, SAII and
TAII) in its database.

On each PE, all PW's with the same attachment ID must have different
backup priority. In this case, PE2 will always reject the mapping
message with the same backup priority by replying a Label Release
message. PE2 should notify PE1 with a "Mismatched Backup Priority"
status code.

If PE2 decides to accept the Label Mapping message, then it has to
make sure that a LSP is setup in the opposite direction (PE1->PE2).

If no corresponding tunnel, it must initiate it by sending a Label
Mapping message to PE1.

Other than reversing the SAI and TAI in PW FEC, PE2 must send the
same Protection TLV (with the same Reference ID and Backup Priority
etc.) back to PE1.

4.5. Manual Provisioning

When a backup PW is initiated from one end (PE1), the other end (PE2)
must comply by replying a Label Mapping message with the same
Protection TLV. However, it is possible that the operators are to
setup a PW from both ends (PE1 and PE2) manually. In this case, if
the protection parameters are inconsistent, the PE's need to reject
the PW setup, and notify the operators.

5. Pseudo-wire Preference Extension

As described in Section 3.5, one method to protect ìimportantî
pseudo-wires is by allowing them to preempt other ìless importantî

ones. This will require the operators to assign preference levels to each of the pseudo-wires.

5.1. The PREFERENCE TLV

With the Preference TLV, the operator can inform the S-PEís and T-PEís about the relative importance of a pseudo-wire.

The Preference TLV has the following format:

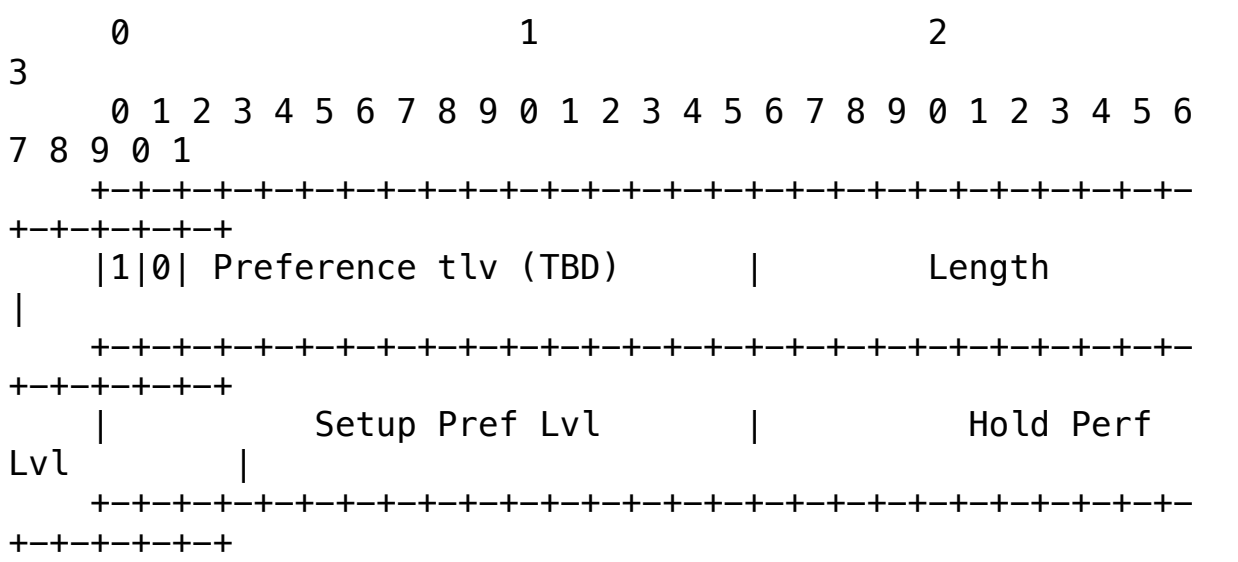

- U bit (always set)

  For the PE nodes that do not support the Preference TLV, no preemption will take place in event of network congestion.

- Preference tlv

  The value of the new tlv type needs to be allocated by IANA.

- Setup Preference Level

  This is the preference level with respect to initiate a PW. The
  value of 0 is the highest. The Setup Preference Level is used in
  deciding whether this PW can preempt another PW.

- Holding Preference Level

  This is the preference level with respect to maintain a PW. The
  value of 0 is the highest. The Holding Preference Level is used in
  deciding whether this PW can be preempted by another PW.

5.2. The Interpretation of Preference

PW preemption implementation depends on the definition of
preferences.

Each preference level can be interpreted as the strict priority of a
PW, or a traffic class as defined in DiffServ, or the weight of a
data flow when it is in combination with the bandwidth assigned to
the PW, or the combination of above.

This requires further feedback from the operators.

However, when we mention preemption, it implies that data packets
from the PWís with ìhighî preference level will override the network
links that may have been shared by another set of PWís.

5.3. PE Operation

The operators from the head-end PE can assign the preference information during PW setup.

Upon the detection of network congestion, the PE node can preempt the less important PWís, and allow the important traffic to go through.

The preemption operation may take place on each segment of a MS-PW, or the entire path of a MS-PW. The operation details require further details from the operators.

6. IANA Considerations

We have defined the following protocol extension:

6.1. PW Protection TLV

This is a new LDP TLV type.

6.2. PE Preference TLV

This is a new LDP TLV type.

6.3. PW Status Code

The edge nodes need to information each other in a number of error conditions. Several PW status code need to be defined:

0x00000XYZ "Duplicated Reference ID"
0x00000XYZ ìDonít support PW Protectionî
0x00000XYZ "Mismatched Backup Priority"
0x00000XYZ "Out of Backup Resource"
0x00000XYZ "Working and backup PW's share the same fate"

Security Considerations

This document specifies the LDP extensions that are needed for protecting pseudo-wires. It will have the same security properties as in [LDP] and [PW-CTRL].

Normative References

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

[PW-CTRL] L. Martini, et al, "Pseudowire Setup and Maintenance using LDP", draft-ietf-pwe3-control-protocol-14.txt

[LDP] L. Andersson, et al, "LDP Specification", draft-ietf-mpls-rfc3036bis-00.txt

[MPLS-FRR] P. Pan, et al, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC4090

[ATT-REPORT] T. Afferton, et al, "Packet Aware Transport for Metro Networks", IEEE Network Magazine, April 2004.

[Segmented-PW] Martini et.al. " Segmented Pseudo Wire", draft-ietf-
pwe3-segmented-pw-00.txt, July 2005

[MHOP-PW] Florin Balus et. al. ìDynamic Placement of Multi Segment
Pseudo Wiresî, draft-ietf-pwe3-dynamic-ms-pw-01.txt

[MS-ARCH] M Bocci et. al. ìAn Architecture for Multi-Segment Pseudo
Wire Emulation Edge-to-Edgeî, draft-ietf-pwe3-ms-pw-arch-00.txt

[NOTE1] Other mechanism may also be applicable for planned shutdown.
See ìLDP graceful restart for planned outages (draft-minei-mpls-ldp-
planned-restart-01.txt)î by Ina Minei, et al.

[L2VPN] Rosen et. al. ìProvisioning, Autodiscovery, and Signaling in
L2VPNsî, draft-ietf-l2vpn-signaling-06.txt

Informative References

None

Acknowledgments

<Add any acknowledgements>

Author's Addresses

Ping Pan
Hammerhead Systems
ppan@hammerheadsystems.com

Matthew Bocci
Alcatel
Matthew.Bocci@alcatel.co.uk

Mustapha Aissaoui
Alcatel
mustapha.aissaoui@alcatel.com

Florin Balus
Nortel
balus@nortel.com

Hamid Ould-Brahim
Nortel
hbrahim@nortel.com

Intellectual Property Statement

made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity