Network Working Group
Internet Draft
Expires: September 12, 2014

Mustapha Aissaoui
Peter Busschbach
Alcatel-Lucent

Dave Allan
Ericsson

Monique Morrow
Cisco Systems Inc.

Thomas Nadeau
Juniper Networks

Editors

March 12, 2014

# OAM Procedures for VPWS Interworking
**draft-ietf-l2vpn-vpws-iw-oam-04.txt**

## Status of this Memo

## Abstract

This draft proposes OAM procedures for the Ethernet interworking, IP interworking and FR-ATM interworking Virtual Private Wire Service (VPWS).

## 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.

## Table of Contents

## 1. Contributors

The following individuals contributed significant ideas or text:

Matthew Bocci, matthew.bocci@alcatel-lucent.com
Simon Delord, simon.delord@gmail.com
Paul Doolan, paul.doolan@coriant.com
Mike Loomis, mike.loomis@alcatel-lucent.com
Hamid Ould-Brahim, ouldh@yahoo.com
Vasile Radoaca, vasile.radoaca@alcatel-lucent.com
Himanshu Shah, hshah@ciena.com
David Watkinson, david.watkinson@alcatel-lucent.com
John Z. Yu.

## 2. Introduction

This draft augments OAM message mapping [RFC6310] with OAM procedures for scenarios when the attachment circuit does not correspond to the pseudo wire. When combined with procedures defined in [RFC6310] and [RFC7023], comprehensive OAM interworking can be defined for VPWS services. VPWS services are defined in the L2 VPN framework [RFC4664].

The following VPWS are covered in this document:

1. VPWS with heterogeneous ACs of ATM and FR types, and in which the PW type is ATM or FR. In this case, FR-ATM service interworking [FRF8.2] is performed at one end of the VPWS and a FR (or ATM) PW is extended to the remote PE. This VPWS will be referred to as FR-ATM Interworking VPWS.

2. VPWS with heterogeneous ACs of ATM, FR, Ethernet, and PPP/BCP types, and in which the PW type is Ethernet. This VPWS will be referred to as Ethernet Interworking VPWS.

3. VPWS with heterogeneous ACs of ATM, FR, Ethernet, and PPP/IPCP types, and in which the PW type is IP [RFC6575]. This VPWS will be referred to as IP Interworking VPWS.

OAM procedures for homogeneous VPWS circuits of ATM and FR types are described in [RFC6310]. OAM procedures for homogeneous VPWS circuits of Ethernet type are defined in [RFC7023].

The PPP PW encapsulation [RFC 4618] describes in Section 5.3 the need to generate a PW status notification to the far-end PE if a change to the status of the PPP AC or PW is detected. However, the PPP protocol does not have a standardized OAM mechanism to propagate to the PPP AC defects detected on the PW.

## 3. Conventions

The words "defect" and "fault" are used interchangeably to mean any condition that obstructs forwarding of user traffic between the CE endpoints of the VPWS.

The words "defect notification" and "defect indication" are used interchangeably to mean any OAM message generated by a PE and sent to other nodes in the network to convey the defect state local to this PE.

An end-to-end virtual circuit in a VPWS consists of a 3 segment set: <AC, PW, AC> [RFC4664]. Note that the AC does not need to connect a CE directly to a PE. An intermediate L2 network may exist.

A VPWS is homogeneous if AC and PW types are the same. E.g., ATM VPWS: <ATM AC, ATM PW, ATM AC>.

A VPWS is heterogeneous if any two segments of the circuit are of different types. E.g., IP interworking circuit: <ATM AC, IP PW, ATM AC>, or <ATM AC, IP PW, ETH AC>.

The PW of a VPWS can be carried over three types of Packet Switched Networks (PSNs). An "MPLS PSN" makes use of MPLS Label Switched Paths [RFC3031] as the tunneling technology to forward the PW packets. An "MPLS/IP PSN" makes use of MPLS-in-IP tunneling [RFC4023], with an MPLS shim header used as PW demultiplexer. An "L2TPv3/IP PSN" makes use of L2TPv3/IP [RFC3931] as the tunneling technology with the L2TPv3/IP Session ID as the PW demultiplexer.

If LSP-Ping [RFC4379] is run over a PW as described in [RFC5085], it will be referred to as "VCCV-Ping". If BFD is run over a PW as described in [RFC5885], it will be referred to as "VCCV-BFD".

While PWs are inherently bidirectional entities, defects and OAM messaging are related to a specific traffic direction. We use the terms "upstream" and "downstream" to identify PEs in relation to the traffic direction. A PE is upstream for the traffic it is forwarding and is downstream for the traffic it is receiving.

We use the terms "local" and "remote" to identify native service networks and ACs in relation to a specific PE. The local AC is attached to the PE in question, while the remote AC is attached to the PE at the other end of the PW.

A "transmit defect" is any defect that uniquely impacts traffic sent or relayed by the observing PE. A "receive defect" is any defect that impacts information transfer to the observing PE. Note that a receive defect also impacts traffic relayed, and thus can be considered to incorporate two defect states. Thus, when a PE enters both receive and transmit defect states of a VPWS, the receive defect takes precedence over the transmit defect in terms of the consequent actions.

A "forward defect indication" (FDI) is sent in the same direction as the user traffic impacted by the defect. A "reverse defect indication" (RDI) is sent in the direction opposite to that of the impacted traffic.

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].

## 4. Reference Model and Defect Locations

Figure 1 illustrates the VPWS network reference model with an indication of the possible defect locations. This model is introduced in [RFC6310] for homogeneous VPWS and is also valid for heterogeneous VPWS.

```
            ACs             PSN tunnel           ACs
                   +----+                  +----+
  +----+           | PE1|==================| PE2|          +----+
  |    |---(a)---(b)..(c)......PW1..(d)..(e)..(f)---(g)---|    |
  | CE1|   (N1)    |    |                  |    |    (N2)  |CE2 |
  |    |----------|............PW2.............|----------|    |
  +----+           |    |==================|    |          +----+
       ^           +----+                  +----+          ^
       |      Provider Edge 1        Provider Edge 2       |
       |                                                   |
       |<-------------- Emulated Service ---------------->|
 Customer                                              Customer
  Edge 1                                                Edge 2
```

Figure 1: PWE3 Network Defect Locations

The procedures will be described in this document from the viewpoint of PE1, so that N1 is the local Native Service (NS) network and N2 is the remote NS network. It is assumed that the AC and PW are of different types at PE1. PE2 will typically implement the same procedures. Note that PE1 is the upstream PE for traffic originating in the local NS network N1, while it is the downstream PE for traffic originating in the remote NS network N2.

The following is a brief description of the defect locations:

a. Defect in NS network N1. This covers any defect in network N1 (including any CE1 defect) that impacts all or some ACs attached to PE1, and is thus a local AC defect. The defect is conveyed to PE1 and to NS network N2 using NS specific OAM defect indications.

b. Defect on a PE1 AC interface (another local AC defect).

c. Defect on a PE1 PSN interface.

d. Defect in the PSN network. This covers any defect in the PSN that impacts all or some PWs between PE1 and PE2. The defect is conveyed to the PE using a PSN and/or a PW specific OAM defect indication. Note that both data plane defects and control plane defects must be taken into consideration. Although control plane packets may follow a different path than PW data plane packets, a control plane defect may affect the PW status.

e. Defect on a PE2 PSN interface.

f. Defect on a PE2 AC interface (a remote AC defect).

g. Defect in NS network N2 (another remote AC defect). This covers any defect in N2 (including any CE2 defect) which impacts all or a subset of ACs attached to PE2. The defect is conveyed to PE2 and to NS network N1 using the NS OAM defect indication.

## 5. Abstract Defect States

PE1 must track four defect states that reflect the observed states of both directions of the VPWS on both the AC and the PW sides.

Defects may impact one or both directions of the VPWS. The observed state is a combination of defects directly detected by PE1 and defects of which it has been made aware via notifications.

```
                          +-----+
     ----AC receive---->|     |-----PW transmit---->
CE1                     | PE1 |                       PE2/CE2
     <---AC transmit----|     |<----PW receive-----
                          +-----+
```

(arrows indicate direction of user traffic impacted by a defect)
Figure 2: Receive and Transmit Defect States

PE1 will directly detect or be notified of AC receive or PW receive defects as they occur upstream of PE1 and impact traffic being sent to PE1. As a result, PE1 enters the AC or PW receive defect state.

In Figure 2, PE1 may be notified of a receive defect in the AC by receiving a Forward Defect indication, e.g., ATM AIS, from CE1 or an intervening network. This defect notification indicates that user traffic sent by CE1 may not be received by PE1 due to a defect. PE1 can also directly detect an AC receive defect if it resulted from a failure of the receive side in the local port or link over which the AC is configured.

Similarly, PE1 may detect or be notified of a receive defect in the PW by receiving a Forward Defect indication from PE2. If PW status is used for fault notification, this message will indicate a Local PSN-facing PW (egress) Transmit Fault or a Local AC (ingress) Receive Fault at PE2 [RFC4446]. This defect notification indicates that user traffic sent by CE2 may not be received by PE1 due to a defect. As a result, PE1 enters the PW receive defect state.

Note that a Forward Defect Indication is sent in the same direction as the user traffic impacted by the defect.

Generally, a PE cannot detect transmit defects by itself and will therefore need to be notified of AC transmit or PW transmit defects by other devices.

In Figure 2, PE1 may be notified of a transmit defect in the AC by receiving a Reverse Defect indication, e.g., ATM RDI, from CE1. This defect relates to the traffic sent by PE1 to CE1 on the AC.

Similarly, PE1 may be notified of a transmit defect in the PW by receiving a Reverse Defect indication from PE2. If PW status is used for fault notification, this message will indicate a Local PSN facing PW (ingress) Receive Fault or a Local Attachment Circuit (egress) Transmit Fault at PE2 [RFC4446]. This defect impacts the traffic sent by PE1 to CE2. As a result, PE1 enters the PW transmit defect state.

Note that a Reverse Defect indication is sent in the reverse direction to the user traffic impacted by the defect.

The procedures outlined in this document define the entry and exit criteria for each of the four states with respect to the set of heterogeneous VPWS within the document scope and the consequent actions that PE1 must perform.

When a PE enters both receive and transmit defect states related to the same VPWS, then the receive defect takes precedence over transmit defect in terms of the consequent actions.

## 6. VPWS OAM Modes

A heterogeneous VPWS forwards packets between an AC and a PW of different types. It thus implements both NS OAM and PW OAM mechanisms.

PW OAM defect notification messages and NS OAM messages are described in [RFC6310]. Ethernet NS OAM messages are described in [RFC7023].

[RFC6310] defined two different OAM modes, the distinction being the method of mapping between the NS and PW OAM defect notification messages.

The first mode, illustrated in Figure 3, is called the "single emulated OAM loop" mode. Here a single end-to-end NS OAM loop is emulated by transparently passing NS OAM messages over the PW. Note that the PW OAM is shown outside the PW in Figure 3, as it is transported in LDP messages or in the associated channel, not inside the PW itself.

```
                  +-----+                 +-----+
    +-----+         |     |=================|     |         +-----+
    | CE1 |-=NS-OAM=>| PE1 |----=NS-OAM=>----| PE2 |-=NS-OAM=>| CE2 |
    +-----+         |     |=================|     |         +-----+
                  +-----+                 +-----+
                    \                       /
                     -------=PW-OAM=>-------
```

Figure 3: Single Emulated OAM Loop Mode

The single emulated OAM loop mode implements the following behavior:

a. The upstream PE (PE1) MUST transparently relay NS OAM messages over the PW.

b. The upstream PE MUST signal local defects affecting the AC using a NS defect notification message sent over the PW. In the case that it is not possible to generate NS OAM messages (e.g., because the defect interferes with NS OAM message generation), the PE MUST signal local defects affecting the AC using a PW defect notification message.

c. The upstream PE MUST signal local defects affecting the PW using a PW defect notification message.

d. The downstream PE (PE2) MUST insert NS defect notification messages into its local AC when it detects or is notified of a defect in the PW or remote AC. This includes translating received PW defect notification messages into NS defect notification messages for defects signaled by the upstream PE.

The second OAM mode operates three OAM loops joined at the AC/PW boundaries of the PEs. This is referred to as the "coupled OAM loops" mode and is illustrated in Figure 4. Note that in contrast to Figure 3, NS OAM messages are never carried over the PW.

```
               +-----+                 +-----+
+-----+          |     |=================|     |          +-----+
| CE1 |-=NS-OAM=>| PE1 |                 | PE2 |-=NS-OAM=>| CE2 |
+-----+          |     |=================|     |          +-----+
               +-----+                 +-----+
                 \                       /
                  -------=PW-OAM=>-------
```

Figure 4: Coupled OAM Loops Mode

The coupled OAM loops mode implements the following behavior:

a. The upstream PE (PE1) MUST terminate and translate a received NS defect notification message into a PW defect notification message.

b. The upstream PE MUST signal local failures affecting its local AC using PW defect notification messages to the downstream PE.

c. The upstream PE MUST signal local failures affecting the PW using PW defect notification messages.

d. The downstream PE (PE2) MUST insert NS defect notification messages into the AC (unless the AC is PPP) when it detects or is notified of defects in the PW or remote AC. This includes

translating received PW defect notification messages into NS defect notification messages.

Table 1 summarizes the OAM mode used with each VPWS covered in this document.

| VPWS | Single Emulated OAM Loop Mode | Coupled OAM Loops Mode |
|---|---|---|
| FR-ATM Interworking<br>- ATM cell mode PW | X | |
| FR-ATM Interworking<br>- FR or AAL5 PDU/SDU PW | | X |
| Ethernet Interworking | | X |
| IP Interworking | | X |

Table 1: Summary of Heterogeneous VPWS OAM Modes

## 7. PW Defect State Entry/Exit

The details of the PW transmit and receive defect state entry/exit criteria are described in Section 6.2 of [RFC6310].

The consequent actions for an ATM AC are described in sections 7.3.1 and 7.3.2 of [RFC6310].

The consequent actions for a FR AC are described in sections 8.3.1 and 8.3.2 of [RFC6310].

The consequent actions for an Ethernet AC are described in sections 6.1 through 6.4 of [RFC7023].

## 8. ATM AC Defect State Entry/Exit

The details of the ATM AC receive and transmit defect state entry/exit criteria are described in sections 7.1 and 7.2 respectively of [RFC6310].

The consequent actions are described in sections 7.3.4 and 7.3.5 of [RFC6310].

Note that all interworking VPWS covered in this document make use of ATM VC as the AC. ATM VP cannot be used as a AC in an interworking VPWS. Therefore only ATM F5 OAM messages are relevant.

## 9. FR AC Defect State Entry/Exit

The details of the FR AC receive and transmit defect state entry/exit criteria are described in sections 8.1 and 8.2 respectively of [RFC6310].

The consequent actions are described in sections 8.3.4 and 8.3.5 of [RFC6310]. Note however that if the FR AC is part of a FR-ATM interworking VPWS operating in the single emulated OAM loop mode, then the consequent actions are described sections 7.3.4 and 7.3.5 of [RFC6310].

## 10. Ethernet AC Defect State Entry/Exit

The details of the Ethernet AC receive and transmit defect state entry/exit criteria are described in sections 5.1 and 5.2 respectively of [RFC7023].

The consequent actions are described in sections 6.5 through 6.8 of [RFC7023].

## 11. PPP AC Defect State Entry/Exit

The PPP PW encapsulation [RFC 4618] describes in Section 5.3 the need to generate a PW status notification to the far-end PE if a change to the status of the PPP AC or PW is detected. However, the PPP protocol does not have a standardized OAM mechanism to propagate to the PPP AC defects detected on the PW.

This document does not define additional procedures for a PPP AC used in an Ethernet or IP interworking VPWS.

## 12. Security Considerations

The mapping messages described in this document do not change the security functions inherent in the actual messages. All generic security considerations applicable to PW traffic specified in Section 10 of [RFC3985] are applicable to NS OAM messages transferred inside the PW.

Security considerations in Section 10 of [RFC5085] for VCCV apply to the OAM messages thus transferred. Security considerations

applicable to the PWE3 control protocol of [RFC4447] Section 8.2 apply to OAM indications transferred using the LDP status message.

## 13. IANA Considerations

This document requires no IANA actions.

## 14. References

### 14.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC6310] Nadeau, T., et al., ''Pseudo Wire (PW) OAM Message Mapping'', RFC 6310, April 2011.

[RFC7023] Qiu, R., Mohan, D., Bitar, N., DeLord, S., Niger, P., and A. Sajassi, "MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking ", RFC 7023, October 2013.

### 14.2. Informative References

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

[RFC3931] Lau, J., Townsley, M., and I. Goyret, "Layer Two Tunneling Protocol - Version 3 (L2TPv3)", RFC 3931, March 2005.

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

[RFC4023] Worster, T., Rekhter, Y., and E. Rosen, "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, March 2005.

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

[RFC4446] Martini, L., "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", BCP 116, RFC 4446, April 2006.

[RFC4618] Martini, L., "Encapsulation Methods for Transport of PPP/High-Level Data Link Control (HDLC) over MPLS Networks", RFC 4618, September 2006.

[RFC4664] Andersson, L. et. al., "L2VPN Framework", RFC 4664, September 2006.

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

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

[RFC6575] Shah, H., et al., ''Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs'', RFC 6575, June 2012.

[FRF8.2] Frame Relay Forum, ''FRF 8.2 - Frame Relay / ATM PVC Service Interworking Implementation Agreement'', February 2004.

## 15. Editor's Addresses

Mustapha Aissaoui
Alcatel-lucent
Email: mustapha.aissaoui@alcatel-lucent.com

Dave Allan
Ericsson
david.i.allan@ericsson.com

Peter B. Busschbach
Alcatel-Lucent
Email: peter.busschbach@alcatel-lucent.com

Thomas Nadeau
Juniper Networks
tnadeau@lucidvision.com

Monique Morrow
Cisco Systems, Inc.
EMail: mmorrow@cisco.com