PWE3   M. Aissaoui
Internet-Draft   P. Busschbach
Intended status: Standards Track   Alcatel-Lucent
Expires: October 21, 2011   L. Martini
                                                               M. Morrow
                                                     Cisco Systems, Inc.
                                                               T. Nadeau
                                                          CA Technologies
                                                            Y(J). Stein
                                                  RAD Data Communications
                                                          April 21, 2011

                     Pseudowire (PW) OAM Message Mapping
                      draft-ietf-pwe3-oam-msg-map-16.txt

Abstract

   This document specifies the mapping and notification of defect states
   between a pseudowire (PW) and the Attachment Circuits (ACs) of the
   end-to-end emulated service. It standardizes the behavior of
   Provider Edges (PEs) with respect to PW and AC defects. It addresses
   ATM, frame relay, TDM, and SONET/SDH PW services, carried over MPLS,
   MPLS/IP and L2TPv3/IP Packet Switched Networks (PSNs).

Table of Contents

[1](#). Introduction

   This document specifies the mapping and notification of defect states
   between a Pseudowire and the Attachment Circuits (AC) of the end-to-
   end emulated service. It covers the case where the ACs and the PWs
   are of the same type in accordance to the PWE3 architecture [[RFC3985](#)]
   such that a homogeneous PW service can be constructed.

   This document is motivated by the requirements put forth in [[RFC4377](#)]
   and [[RFC3916](#)]. Its objective is to standardize the behavior of PEs
   with respect to defects on PWs and ACs, so that there is no ambiguity
   about the alarms generated and consequent actions undertaken by PEs
   in response to specific failure conditions.

   This document addresses PWs over MPLS, MPLS/IP and L2TPv3/IP PSNs,
   and ATM, frame relay, TDM, and SONET/SDH PW native services. Due to
   its unique characteristics, the Ethernet PW service is covered in a
   separate document [[I-D.ietf-pwe3-mpls-eth-oam-iwk](#)].

   This document provides procedures for PWs set up using LDP [[RFC4447](#)]
   or L2TPv3 [[RFC3931](#)] control protocols. While we mention fault
   reporting options for PWs established by other means (e.g., by static
   configuration or via BGP), we do not provide detailed procedures for
   such cases.

   This document is scoped only to single segment PWs. The mechanisms
   described in this document could also be applied to T-PEs for MS-PWs
   ([[RFC5254](#)]). [Section 10 of [RFC6073]](#) details procedures for
   generating or relaying PW status by an S-PE.

2. Abbreviations and Conventions

2.1. Abbreviations

    AAL5 ATM Adaptation Layer 5
    AIS   Alarm Indication Signal
    AC   Attachment Circuit
    ATM   Asynchronous Transfer Mode
    AVP   Attribute Value Pair
    BFD   Bidirectional Forwarding Detection
    CC   Continuity Check
    CDN   Call Disconnect Notify
    CE   Customer Edge
    CV   Connectivity Verification
    DBA   Dynamic Bandwidth Allocation
    DLC   Data Link Connection
    FDI   Forward Defect Indication
    FR   Frame Relay
    FRBS Frame Relay Bearer Service
    ICMP Internet Control Message Protocol
    LB   Loopback
    LCCE L2TP Control Connection Endpoint
    LDP   Label Distribution Protocol
    LSP   Label Switched Path
    L2TP Layer 2 Tunneling Protocol
    MPLS Multiprotocol Label Switching
    NE   Network Element
    NS   Native Service
    OAM   Operations, Administration, and Maintenance
    PE   Provider Edge
    PSN   Packet Switched Network
    PW   Pseudowire
    RDI   Remote Defect Indication
    PDU   Protocol Data Unit
    SDH   Synchronous Digital Hierarchy
    SDU   Service Data Unit
    SONET   Synchronous Optical Network
    TDM   Time Division Multiplexing
    TLV   Type Length Value
    VCC   Virtual Channel Connection
    VCCV Virtual Connection Connectivity Verification
    VPC   Virtual Path Connection

2.2. Conventions

    The words "defect" and "fault" are used interchangeably to mean any
    condition that negatively impacts forwarding of user traffic between
    the CE endpoints of the PW service.

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.

The PW 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 meant to be relayed, and thus can be considered to incorporate two defect states. Thus when a PE enters both receive and transmit defect states of a PW service, 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].

3. Reference Model and Defect Locations

   Figure 1 illustrates the PWE3 network reference model with an
   indication of the possible defect locations. This model will be
   referenced in the remainder of this document for describing the OAM
   procedures.

```
              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 network and N2 is the
   remote native service network. PE2 will typically implement the same
   functionality. 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
      messages may follow a different path than PW data plane traffic, 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) that 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.

4. Abstract Defect States

   PE1 must track four defect states that reflect the observed states of
   both directions of the PW service on both the AC and the PW sides.
   Defects may impact one or both directions of the PW service.

   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 the PW
   status TLV 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, as described in Section 6.1.1. 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, as described in Section 6.1.1. 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 PW
services 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 PW service, then the receive defect takes precedence over
transmit defect in terms of the consequent actions.

5. OAM Modes

A homogeneous PW service forwards packets between an AC and a PW of
the same type. It thus implements both NS OAM and PW OAM mechanisms.
PW OAM defect notification messages are described in Section 6.1. NS
OAM messages are described in Appendix A.

This document defines 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 single emulated OAM loop mode is suitable for PW services that
   have a widely deployed NS OAM mechanism. This document specifies the
   use of this mode for ATM PW, TDM PW, and CEP PW services. It is the
   default mode of operation for all ATM cell-mode PW services and the
   only mode specified for CEP and SAToP/CESoPSN TDM PW services. It is
   optional for AAL5 PDU transport and AAL5 SDU transport modes.

   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 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.

   This document specifies the coupled OAM loops mode as the default
   mode for the frame relay, ATM AAL5 PDU transport, and AAL5 SDU
   transport services. It is an optional mode for ATM VCC cell mode
   services. This mode is not specified for TDM, CEP, or ATM VPC cell
   mode PW services. RFC5087 defines a similar but distinct mode, as
   will be explained in Section 9 below. For the ATM VPC cell mode case
   a pure coupled OAM loops mode is not possible as a PE MUST
   transparently pass VC-level (F5) ATM OAM cells over the PW while
   terminating and translating VP-level (F4) OAM cells.

6. PW Defect States and Defect Notifications

6.1. PW Defect Notification Mechanisms

   For MPLS and MPLS/IP PSNs, a PE that establishes a PW using the Label
   Distribution Protocol [RFC5036], and that has negotiated use of the
   LDP status TLV per Section 5.4.3 of [RFC4447], MUST use the PW status
   TLV mechanism for AC and PW status and defect notification.
   Additionally, such a PE MAY use VCCV-BFD Connectivity Verification
   (CV) for fault detection only (CV types 0x04 and 0x10 [RFC5885]).

   A PE that establishes an MPLS PW using means other than LDP, e.g., by
   static configuration or by use of BGP, MUST support some alternative
   method of status reporting. The design of a suitable mechanism to

carry the aforementioned status TLV in the the PW associated channel
is work in progress [I-D.ietf-pwe3-static-pw-status]. Additionally,
such a PE MAY use VCCV-BFD CV for both fault detection and status
notification (CV types 0x08 and 0x20 [RFC5885]).

For a L2TPv3/IP PSN, a PE SHOULD use the Circuit Status Attribute
Value Pair (AVP) as the mechanism for AC and PW status and defect
notification. In its most basic form, the Circuit Status AVP
[RFC3931] in a Set-Link-Info (SLI) message can signal active/inactive
AC status. The Circuit Status AVP as described in [RFC5641] is
proposed to be extended to convey status and defects in the AC and
the PSN-facing PW in both ingress and egress directions, i.e., four
independent status bits, without the need to tear down the sessions
or control connection.

When a PE does not support the Circuit Status AVP, it MAY use the
Stop-Control-Connection-Notification (StopCCN) and the Call-
Disconnect-Notify (CDN) messages to tear down L2TP sessions in a
fashion similar to LDP's use of Label Withdrawal to tear down a PW.
A PE may use the StopCCN to shutdown the L2TP control connection, and
implicitly all L2TP sessions associated with that control connection,
without any explicit session control messages. This is useful for
the case of a failure which impacts all L2TP sessions (all PWs)
managed by the control connection. It MAY use CDN to disconnect a
specific L2TP session when a failure only affects a specific PW.

Additionally, a PE MAY use VCCV-BFD CV types 0x04 and 0x10 for fault
detection only, but SHOULD notify the remote PE using the Circuit
Status AVP. A PE that establishes a PW using means other than the
L2TP control plane, e.g., by static configuration or by use of BGP,
MAY use VCCV-BFD CV types 0x08 and 0x20 for AC and PW status and
defect notification. These CV types SHOULD NOT be used when the PW
is established via the L2TP control plane.

The CV types are defined in Section 6.1.3 of this document.

6.1.1. LDP Status TLV

[RFC4446] defines the following PW status code points:

0x00000000 - Pseudowire forwarding (clear all failures)
0x00000001 - Pseudowire Not Forwarding
0x00000002 - Local Attachment Circuit (ingress) Receive Fault
0x00000004 - Local Attachment Circuit (egress) Transmit Fault

        0x00000008 - Local PSN-facing PW (ingress) Receive Fault
        0x00000010 - Local PSN-facing PW (egress) Transmit Fault

    [RFC4447] specifies that the "Pseudowire forwarding" code point is
    used to indicate that all faults are to be cleared. It also
    specifies that the "Pseudowire Not Forwarding" code point means that
    a defect has been detected that is not represented by the defined
    code points.

    The code points used in the LDP status TLV in a PW status
    notification message report defects from the viewpoint of the
    originating PE. The originating PE conveys this state in the form of
    a forward defect or a reverse defect indication.

    The forward and reverse defect indication definitions used in this
    document map to the LDP Status TLV codes as follows:

            Forward defect indication corresponds to the logical OR of:
              * Local Attachment Circuit (ingress) Receive Fault,
              * Local PSN-facing PW (egress) Transmit Fault, and
              * PW Not Forwarding.
            Reverse defect indication corresponds to the logical OR of:
              * Local Attachment Circuit (egress) Transmit Fault and
              * Local PSN-facing PW (ingress) Receive Fault.

    A PE MUST use PW status notification messages to report all defects
    affecting the PW service including, but not restricted, to the
    following:

    o defects detected through fault detection mechanisms in the MPLS
       and MPLS/IP PSN,
    o defects detected through VCCV-Ping or VCCV-BFD CV types 0x04 and
       0x10 for fault detection only,
    o defects within the PE that result in an inability to forward
       traffic between the AC and the PW,
    o defects of the AC or in the Layer 2 network affecting the AC as
       per the rules detailed in Section 5 for the "single emulated OAM
       loop" mode and "coupled OAM loops" modes.

    Note that there are two situations that require PW label withdrawal
    as opposed to a PW status notification by the PE. The first one is
    when the PW is taken down administratively in accordance with
    [RFC4447]. The second one is when the Target LDP session established
    between the two PEs is lost. In the latter case, the PW labels will
    need to be re-signaled when the Targeted LDP session is re-
    established.

6.1.2. L2TP Circuit Status AVP

   [RFC3931] defines the Circuit Status AVP in the Set-Link-Info (SLI)
   message to exchange initial status and status changes in the circuit
   to which the pseudowire is bound. [RFC5641] defines extensions to
   the Circuit Status AVP that are analogous to the PW Status TLV
   defined for LDP. Consequently, for L2TPv3/IP the Circuit Status AVP
   is used in the same fashion as the PW Status described in the
   previous section. Extended circuit status for L2TPv3/IP is described
   in [RFC5641].

   If the extended Circuit Status bits are not supported, and instead
   only the "A-bit" (Active) is used as described in [RFC3931], a PE MAY
   use CDN messages to clear L2TPv3/IP sessions in the presence of
   session-level failures detected in the L2TPv3/IP PSN.

   A PE MUST set the Active bit in the Circuit Status to clear all
   faults, and it MUST clear the Active bit in the Circuit Status to
   convey any defect that cannot be represented explicitly with specific
   Circuit Status flags from [RFC3931] or [RFC5641].

   The forward and reverse defect indication definitions used in this
   document map to the L2TP Circuit Status AVP as follows:

           Forward defect indication corresponds to the logical OR of:
             * Local Attachment Circuit (ingress) Receive Fault,
             * Local PSN-facing PW (egress) Transmit Fault, and
             * PW Not Forwarding.
           Reverse defect indication corresponds to the logical OR of:
             * Local Attachment Circuit (egress) Transmit Fault and
             * Local PSN-facing PW (ingress) Receive Fault.

   The status notification conveys defects from the viewpoint of the
   originating LCCE (PE).

   When the extended Circuit Status definition of [RFC5641] is
   supported, a PE SHALL use the Circuit Status to report all failures
   affecting the PW service including, but not restricted, to the
   following:

   o defects detected through defect detection mechanisms in the
      L2TPv3/IP PSN,
   o defects detected through VCCV-Ping or VCCV-BFD CV types 0x04 (BFD
      IP/UDP-encapsulated, for PW Fault Detection only) and 0x10 (BFD
      PW-ACH-encapsulated (without IP/UDP headers), for PW. Fault
      Detection and AC/PW Fault Status Signaling) for fault detection
      only which are described in Section 6.1.3 of this document,

o   defects within the PE that result in an inability to forward
    traffic between the AC and the PW,
o   defects of the AC or in the L2 network affecting the AC as per the
    rules detailed in Section 5 for the "single emulated OAM loop"
    mode and the "coupled OAM loops" modes.

When the extended Circuit Status definition of [RFC5641] is not
supported, a PE SHALL use the A-bit in the Circuit Status AVP in SLI
to report:

o   defects of the AC or in the L2 network affecting the AC as per the
    rules detailed in Section 5 for the "single emulated OAM loop"
    mode and the "coupled OAM loops" modes.

When the extended Circuit Status definition of [RFC5641] is not
supported, a PE MAY use the CDN and StopCCN messages in a similar way
to an MPLS PW label withdrawal to report:

o   defects detected through defect detection mechanisms in the
    L2TPv3/IP PSN (using StopCCN),
o   defects detected through VCCV (pseudowire level) (using CDN),
o   defects within the PE that result in an inability to forward
    traffic between ACs and PW (using CDN).

For ATM L2TPv3/IP pseudowires, in addition to the Circuit Status AVP,
a PE MAY use the ATM Alarm Status AVP [RFC4454] to indicate the
reason for the ATM circuit status and the specific alarm type, if
any. This AVP is sent in the SLI message to indicate additional
information about the ATM circuit status.

L2TP control connections use Hello messages as a keep-alive facility.
It is important to note that if PSN failure is detected by keep-alive
timeout, the control connection is cleared. L2TP Hello messages are
sent in-band so as to follow the data plane with respect to the
source and destination addresses, IP protocol number and UDP port
(when UDP is used).

6.1.3. BFD Diagnostic Codes

BFD [RFC5880] defines a set of diagnostic codes that partially
overlap the set of defects that can be communicated through LDP
Status TLV or L2TP Circuit Status AVP. This section describes the
behavior of the PEs with respect to using one or both of these
methods for detecting and propagating defect state.

In the case of an MPLS PW established via LDP signaling, the PEs
negotiate VCCV capabilities during the label mapping messages
exchange used to establish the two directions of the PW. This is

achieved by including a capability TLV in the PW FEC interface
parameters TLV. In the L2TPv3/IP case, the PEs negotiate the use of
VCCV during the pseudowire session initialization using the VCCV AVP
[RFC5085].

The CV Type Indicators field in the OAM capability TLV or VCCV AVP
defines a bitmask used to indicate the specific OAM capabilities that
the PE can use over the PW being established.

A CV type of 0x04 or 0x10 [RFC5885] indicates that BFD is used for PW
fault detection only. These CV types MAY be used any time the PW is
established using LDP or L2TP control planes. In this mode, only the
following diagnostic (Diag) codes specified in [RFC5880] will be
used:

   0 - No diagnostic
   1 - Control detection time expired
   3 - Neighbor signaled session down
   7 - Administratively Down

A PE using VCCV-BFD MUST use diagnostic code 0 to indicate to its
peer PE that is correctly receiving BFD control messages. It MUST
use diagnostic code 1 to indicate that to its peer it has stopped
receiving BFD control messages and will thus declare the PW to be
down in the receive direction. It MUST use diagnostic code 3 to
confirm to its peer that the BFD session is going down after
receiving diagnostic code 1 from this peer. In this case, it will
declare the PW to be down in the transmit direction. A PE MUST use
diagnostic code 7 to bring down the BFD session when the PW is
brought down administratively. All other defects, such as AC/PW
defects and PE internal failures that prevent it from forwarding
traffic, MUST be communicated through the LDP Status TLV in the case
of MPLS or MPLS/IP PSN, or through the appropriate L2TP codes in the
Circuit Status AVP in the case of L2TPv3/IP PSN.

A CV type of 0x08 or 0x20 in the OAM capabilities TLV indicates that
BFD is used for both PW fault detection and Fault Notification. In
addition to the above diagnostic codes, a PE uses the following codes
to signal AC defects and other defects impacting forwarding over the
PW service:

   6 - Concatenated Path Down
   8 - Reverse Concatenated Path Down

As specified in [RFC5085], the PEs negotiate the use of VCCV during
PW set-up. When a PW transported over an MPLS-PSN is established
using LDP, the PEs negotiate the use of the VCCV capabilities using
the optional VCCV Capability Advertisement Sub- TLV parameter in the

Interface Parameter Sub-TLV field of the LDP PW ID FEC or using an
Interface Parameters TLV of the LDP Generalized PW ID FEC. In the
case of L2TPv3/IP PSNs, the PEs negotiate the use of VCCV during the
pseudowire session initialization using VCCV AVP.

Note that a defect that causes the generation of the "PW not
forwarding code" (diagnostic code 6 or 8) does not necessarily result
in the BFD session going down. However, if the BFD session times
out, then diagnostic code 1 MUST be used since it signals a state
change of the BFD session itself. In general, when a BFD session
changes state, the PEs MUST use state change diagnostic codes 0, 1,
3, and 7 in accordance with [RFC5880] and they MUST override any of
the AC/PW status diagnostic codes (codes 6 or 8) that may have been
signaled prior to the BFD session changing state.

The forward and reverse defect indications used in this document map
to the following BFD codes:

        Forward defect indication corresponds to the logical OR of:
          * Concatenated Path Down (BFD diagnostic code 06)
          * Pseudowire Not Forwarding (PW status code 0x00000001).
        Reverse defect indication- corresponds to:
          * Reverse Concatenated Path Down (BFD diagnostic code 08).

These diagnostic codes are used to signal forward and reverse defect
states, respectively, when the PEs negotiated the use of BFD as the
mechanism for AC and PW fault detection and status signaling
notification. As stated in Section 6.1, these CV types SHOULD NOT be
used when the PW is established with the LDP or L2TP control plane.

6.2. PW Defect State Entry/Exit

6.2.1. PW receive defect state entry/exit criteria

PE1, as downstream PE, will enter the PW receive defect state if one
or more of the following occurs:

o It receives a Forward Defect Indication (FDI) from PE2 indicating
   either a receive defect on the remote AC, or that PE2 detected or
   was notified of downstream PW fault.
o It detects loss of connectivity on the PSN tunnel upstream of PE1
   which affects the traffic it receives from PE2.
o It detects a loss of PW connectivity through VCCV-BFD or VCCV-PING
   which affects the traffic it receives from PE2.

Note that if the PW control session (LDP session, the L2TP session,
or the L2TP control connection) between the PEs fails, the PW is torn
down and needs to be re-established. However, the consequent actions

towards the ACs are the same as if the PW entered the receive defect
state.

PE1 will exit the PW receive defect state when the following
conditions are met. Note that this may result in a transition to the
PW operational state or the PW transmit defect state.

o All previously detected defects have disappeared, and
o PE2 cleared the FDI, if applicable.

6.2.2. PW transmit defect state entry/exit criteria

PE1, as upstream PE, will enter the PW transmit defect state if the
following conditions occur:

o It receives a Reverse Defect Indication (RDI) from PE2 indicating
   either a transmit fault on the remote AC, or that PE2 detected or
   was notified of a upstream PW fault, and
o it is not already in the PW receive defect state.

PE1 will exit the transmit defect state if it receives an OAM message
from PE2 clearing the RDI, or it has entered the PW receive defect
state.

For a PW over L2TPv3/IP using the basic Circuit Status AVP [RFC3931],
the PW transmit defect state is not valid and a PE can only enter the
PW receive defect state.

7. Procedures for ATM PW Service

The following procedures apply to Asynchronous Transfer Mode (ATM)
pseudowires [RFC4717]. ATM terminology is explained in Appendix A.2
of this document.

7.1. AC receive defect state entry/exit criteria

When operating in the coupled OAM loops mode, PE1 enters the AC
receive defect state when any of the following conditions are met:

a. It detects or is notified of a physical layer fault on the ATM
   interface.
b. It receives an end-to-end Flow 4 OAM (F4) Alarm Indication Signal
   (AIS) OAM flow on a Virtual Path (VP) AC, or an end-to-end Flow 5
   (F5) AIS OAM flow on a Virtual Circuit (VC) as per ITU-T
   Recommendation I.610 [I.610], indicating that the ATM VPC or VCC
   is down in the adjacent Layer 2 ATM network.

   c. It receives a segment F4 AIS OAM flow on a VP AC, or a segment F5
      AIS OAM flow on a VC AC, provided that the operator has
      provisioned segment OAM and the PE is not a segment end-point.
   d. It detects loss of connectivity on the ATM VPC/VCC while
      terminating segment or end-to-end ATM continuity check (ATM CC)
      cells with the local ATM network and CE.

   When operating in the coupled OAM loops mode, PE1 exits the AC
   Receive defect state when all previously detected defects have
   disappeared.

   When operating in the single emulated OAM loop mode, PE1 enters the
   AC receive defect state if any of the following conditions are met:

   a. It detects or is notified of a physical layer fault on the ATM
      interface.
   b. It detects loss of connectivity on the ATM VPC/VCC while
      terminating segment ATM continuity check (ATM CC) cells with the
      local ATM network and CE.

   When operating in the single emulated OAM loop mode, PE1 exits the AC
   receive defect state when all previously detected defects have
   disappeared.

   The exact conditions under which a PE enters and exits the AIS state,
   or declares that connectivity is restored via ATM CC, are defined in
   Section 9.2 of [I.610].

7.2. AC transmit defect state entry/exit criteria

   When operating in the coupled OAM loops mode, PE1 enters the AC
   transmit defect state if any of the following conditions are met:

   a. It terminates an end-to-end F4 RDI OAM flow, in the case of a VPC,
      or an end-to-end F5 RDI OAM flow, in the case of a VCC, indicating
      that the ATM VPC or VCC is down in the adjacent L2 ATM.
   b. It receives a segment F4 RDI OAM flow on a VP AC, or a segment F5
      RDI OAM flow on a VC AC, provided that the operator has
      provisioned segment OAM and the PE is not a segment end-point.

   PE1 exits the AC transmit defect state if the AC state transitions to
   working or to the AC receive defect state. The exact conditions for
   exiting the RDI state are described in Section 9.2 of [I.610].

   Note that the AC transmit defect state is not valid when operating in
   the single emulated OAM loop mode, as PE1 transparently forwards the
   received RDI cells as user cells over the ATM PW to the remote CE.

7.3. **Consequent Actions**

   In the remainder of this section, the text refers to AIS, RDI and CC
   without specifying whether there is an F4 (VP-level) flow or an F5
   (VC- level) flow, or whether it is an end-to-end or a segment flow.
   Precise ATM OAM procedures for each type of flow are specified in
   Section 9.2 of [I.610].

7.3.1. **PW receive defect state entry/exit**

   On entry to the PW receive defect state:

   a. PE1 MUST commence AIS insertion into the corresponding AC.
   b. PE1 MUST cease generation of CC cells on the corresponding AC, if
      applicable.
   c. If the PW defect was detected by PE1 without receiving FDI from
      PE2, PE1 MUST assume PE2 has no knowledge of the defect and MUST
      notify PE2 by sending RDI.

   On exit from the PW receive defect state:

   a. PE1 MUST cease AIS insertion into the corresponding AC.
   b. PE1 MUST resume any CC cell generation on the corresponding AC, if
      applicable.
   c. PE1 MUST clear the RDI to PE2, if applicable.

7.3.2. **PW transmit defect state entry/exit**

   On entry to the PW Transmit Defect State:

   a. PE1 MUST commence RDI insertion into the corresponding AC.
   b. If the PW failure was detected by PE1 without receiving an RDI
      from PE2, PE1 MUST assume PE2 has no knowledge of the defect and
      MUST notify PE2 by sending FDI.

   On exit from the PW Transmit Defect State:

   a. PE1 MUST cease RDI insertion into the corresponding AC.
   b. PE1 MUST clear the FDI to PE2, if applicable.

7.3.3. **PW defect state in ATM Port Mode PW Service**

   In case of transparent cell transport PW service, i.e., "port mode",
   where the PE does not keep track of the status of individual ATM VPCs
   or VCCs, a PE cannot relay PW defect state over these VCCs and VPCs.
   If ATM CC is run on the VCCs and VPCs end-to-end (CE1 to CE2), or on
   a segment originating and terminating in the ATM network and spanning
   the PSN network, it will timeout and cause the CE or ATM switch to

enter the ATM AIS state.

[7.3.4](). **AC receive defect state entry/exit**

On entry to the AC receive defect state and when operating in the
coupled OAM loops mode:

a. PE1 MUST send FDI to PE2.
b. PE1 MUST commence insertion of ATM RDI cells into the AC towards
   CE1.

When operating in the single emulated OAM loop mode, PE1 must be able
to support two options, subject to the operator's preference. The
default option is the following:

On entry to the AC receive defect state:

a. PE1 MUST transparently relay ATM AIS cells, or, in the case of a
   local AC defect, commence insertion of ATM AIS cells into the
   corresponding PW towards CE2.
b. If the defect interferes with NS OAM message generation, PE1 MUST
   send an FDI indication to PE2.
c. PE1 MUST cease the generation of CC cells on the corresponding PW,
   if applicable.

In certain operational models, for example in the case that the ATM
access network is owned by a different provider than the PW, an
operator may want to distinguish between defects detected in the ATM
access network and defects detected on the AC directly adjacent to
the PE. Therefore, the following option MUST also be supported:

a. PE1 MUST transparently relay ATM AIS cells over the corresponding
   PW towards CE2.
b. Upon detection of a defect on the ATM interface on the PE or in
   the PE itself, PE1 MUST send FDI to PE2.
c. PE1 MUST cease generation of CC cells on the corresponding PW, if
   applicable.

On exit from the AC receive defect state and when operating in the
coupled OAM loops mode:

a. PE1 MUST clear the FDI to PE2.
b. PE1 MUST cease insertion of ATM RDI cells into the AC.

On exit from the AC receive defect state and when operating in the
single emulated OAM loop mode:

   a. PE1 MUST cease insertion of ATM AIS cells into the corresponding
      PW.
   b. PE1 MUST clear the FDI to PE2, if applicable.
   c. PE1 MUST resume any CC cell generation on the corresponding PW, if
      applicable.

7.3.5. AC transmit defect state entry/exit

   On entry to the AC transmit defect state and when operating in the
   coupled OAM loops mode:

   * PE1 MUST send RDI to PE2.

   On exit from the AC transmit defect state and when operating in the
   coupled OAM loops mode:

   * PE1 MUST clear the RDI to PE2.

8. Procedures for Frame Relay PW Service

   The following procedures apply to Frame Relay (FR) pseudowires
   [RFC4619]. Frame Relay (FR) terminology is explained in Appendix A.1
   of this document.

8.1. AC receive defect state entry/exit criteria

   PE1 enters the AC receive defect state if one or more of the
   following conditions are met:

   a. A Permanent Virtual Circuit (PVC) is not deleted from the FR
      network and the FR network explicitly indicates in a full status
      report (and optionally by the asynchronous status message) that
      this PVC is inactive [Q.933]. In this case, this status maps
      across the PE to the corresponding PW only.
   b. The Link Integrity Verification (LIV) indicates that the link from
      the PE to the Frame Relay network is down [Q.933]. In this case,
      the link down indication maps across the PE to all corresponding
      PWs.
   c. A physical layer alarm is detected on the FR interface. In this
      case, this status maps across the PE to all corresponding PWs.

   PE1 exits the AC receive defect state when all previously detected
   defects have disappeared.

8.2. **AC transmit defect state entry/exit criteria**

   The AC transmit defect state is not valid for a FR AC.

8.3. **Consequent Actions**

8.3.1. **PW receive defect state entry/exit**

   The A (Active) bit indicates whether the FR PVC is ACTIVE (1) or
   INACTIVE (0) as explained in [RFC4591].

   On entry to the PW receive defect state:

   a. PE1 MUST clear the Active bit for the corresponding FR AC in a
      full status report, and optionally in an asynchronous status
      message, as per Q.933 Annex A [Q.933].
   b. If the PW failure was detected by PE1 without receiving FDI from
      PE2, PE1 MUST assume PE2 has no knowledge of the defect and MUST
      notify PE2 by sending RDI.

   On exit from the PW receive defect state:

   a. PE1 MUST set the Active bit for the corresponding FR AC in a full
      status report, and optionally in an asynchronous status message,
      as per Q.933 annex A. PE1 does not apply this procedure on a
      transition from the PW receive defect state to the PW transmit
      defect state.
   b. PE1 MUST clear the RDI to PE2, if applicable.

8.3.2. **PW transmit defect state entry/exit**

   On entry to the PW transmit defect state:

   a. PE1 MUST clear the Active bit for the corresponding FR AC in a
      full status report, and optionally in an asynchronous status
      message, as per Q.933 Annex A.
   b. If the PW failure was detected by PE1 without RDI from PE2, PE1
      MUST assume PE2 has no knowledge of the defect and MUST notify PE2
      by sending FDI.

   On exit from the PW transmit defect state:

   a. PE1 MUST set the Active bit for the corresponding FR AC in a full
      status report, and optionally in an asynchronous status message,
      as per Q.933 annex A. PE1 does not apply this procedure on a
      transition from the PW transmit defect state to the PW receive
      defect state.

   b. PE1 MUST clear the FDI to PE2, if applicable.

8.3.3. PW defect state in the FR Port Mode PW Service

   In case of port mode PW service, STATUS ENQUIRY and STATUS messages
   are transported transparently over the PW. A PW Failure will
   therefore result in timeouts of the Q.933 link and PVC management
   protocol at the Frame Relay devices at one or both sites of the
   emulated interface.

8.3.4. AC receive defect state entry/exit

   On entry to the AC receive defect state:

   * PE1 MUST send FDI to PE2.

   On exit from the AC receive defect state:

   * PE1 MUST clear FDI to PE2.

8.3.5. AC transmit defect state entry/exit

   The AC transmit defect state is not valid for a FR AC.

9. Procedures for TDM PW Service

   The following procedures apply to SAToP [RFC4553], CESoPSN [RFC5086]
   and TDMoIP [RFC5087]. These technologies utilize the single emulated
   OAM loop mode. RFC 5087 distinguishes between trail-extended and
   trail-terminated scenarios; the former is essentially the single
   emulated loop model. The latter applies to cases where the NS
   networks are run by different operators and defect notifications are
   not propagated across the PW.

   Since TDM is inherently real-time in nature, many OAM indications
   must be generated or forwarded with minimal delay. This requirement
   rules out the use of messaging protocols, such as PW status messages.
   Thus, for TDM PWs, alternate mechanisms are employed.

   The fact that TDM PW packets are sent at a known constant rate can be
   exploited as an OAM mechanism. Thus, a PE enters the PW receive
   defect state whenever a preconfigured number of TDM PW packets do not
   arrive in a timely fashion. It exits this state when packets once
   again arrive at their proper rate.

   Native TDM carries OAM indications in overhead fields that travel
   along with the data. TDM PWs emulate this behavior by sending urgent

OAM messages in the PWE control word.

The TDM PWE3 control word contains a set of flags used to indicate PW
and AC defect conditions. The L bit is an AC forward defect
indication used by the upstream PE to signal NS network defects to
the downstream PE. The M field may be used to modify the meaning of
receive defects. The R bit is a PW reverse defect indication used by
the PE to signal PSN failures to the remote PE. Upon reception of
packets with the R bit set, a PE enters the PW transmit defect state.
L bits and R bits are further described in [RFC5087].

9.1. AC receive defect state entry/exit criteria

PE1 enters the AC receive defect state if any of the following
conditions are met:

a. It detects a physical layer fault on the TDM interface (Loss of
   Signal, Loss of Alignment, etc., as described in [G.775]).
b. It is notified of a previous physical layer fault by detecting
   AIS.

The exact conditions under which a PE enters and exits the AIS state
are defined in [G.775]. Note that Loss of Signal and AIS detection
can be performed by PEs for both structure-agnostic and structure-
aware TDM PW types. Note that PEs implementing structure- agnostic
PWs can not detect Loss of Alignment.

9.2. AC transmit defect state entry/exit criteria

PE1 enters the AC transmit defect state when it detects RDI according
to the criteria in [G.775]. Note that PEs implementing structure-
agnostic PWs can not detect RDI.

9.3. Consequent Actions

9.3.1. PW receive defect state entry/exit

On entry to the PW receive defect state:

a. PE1 MUST commence AIS insertion into the corresponding TDM AC.
b. PE1 MUST set the R bit in all PW packets sent back to PE2.

On exit from the PW receive defect state:

   a. PE1 MUST cease AIS insertion into the corresponding TDM AC.
   b. PE1 MUST clear the R bit in all PW packets sent back to PE2.

   Note that AIS generation can in general be performed by both
   structure-aware and structure-agnostic PEs.

9.3.2. PW transmit defect state entry/exit

   On entry to the PW Transmit Defect State:

   * A structure-aware PE1 MUST commence RDI insertion into the
      corresponding AC.

   On exit from the PW Transmit Defect State:

   * A structure-aware PE1 MUST cease RDI insertion into the
      corresponding AC.

   Note that structure-agnostic PEs are not capable of injecting RDI
   into an AC.

9.3.3. AC receive defect state entry/exit

   On entry to the AC receive defect state and when operating in the
   single emulated OAM loop mode:

   a. PE1 SHOULD overwrite the TDM data with AIS in the PW packets sent
      towards PE2.
   b. PE1 MUST set the L bit in these packets.
   c. PE1 MAY omit the payload in order to conserve bandwidth.
   d. A structure-aware PE1 SHOULD send RDI back towards CE1.
   e. A structure-aware PE1 that detects a potentially correctable AC
      defect MAY use the M field to indicate this.

   On exit from the AC receive defect state and when operating in the
   single emulated OAM loop mode:

   a. PE1 MUST cease overwriting PW content with AIS and return to
      forwarding valid TDM data in PW packets sent towards PE2.
   b. PE1 MUST clear the L bit in PW packets sent towards PE2.
   c. A structure-aware PE1 MUST cease sending RDI towards CE1.

10. Procedures for CEP PW Service

   The following procedures apply to SONET/SDH Circuit Emulation
   [RFC4842]. They are based on the single emulated OAM loop mode.

Since SONET and SDH are inherently real-time in nature, many OAM
indications must be generated or forwarded with minimal delay. This
requirement rules out the use of messaging protocols, such as PW
status messages. Thus, for CEP PWs alternate mechanisms are
employed.

The CEP PWE3 control word contains a set of flags used to indicate PW
and AC defect conditions. The L bit is a forward defect indication
used by the upstream PE to signal to the downstream PE a defect in
its local attachment circuit. The R bit is a PW reverse defect
indication used by the PE to signal PSN failures to the remote PE.
The combination of N and P bits is used by the local PE to signal
loss of pointer to the remote PE.

The fact that CEP PW packets are sent at a known constant rate can be
exploited as an OAM mechanism. Thus, a PE enters the PW receive
defect state when it loses packet synchronization. It exits this
state when it regains packet synchronization. See [RFC4842] for
further details.

10.1. Defect states

10.1.1. PW receive defect state entry/exit

In addition to the conditions specified in Section 6.2.1, PE1 will
enter the PW receive defect state when one of the following becomes
true:

o It receives packets with the L bit set.
o It receives packets with both the N and P bits set.
o It loses packet synchronization.

10.1.2. PW transmit defect state entry/exit

In addition to the conditions specified in Section 6.2.2, PE1 will
enter the PW transmit defect state if it receives packets with the R
bit set.

10.1.3. AC receive defect state entry/exit

PE1 enters the AC receive defect state when any of the following
conditions are met:

a. It detects a physical layer fault on the TDM interface (Loss of
   Signal, Loss of Alignment, etc.).

   b. It is notified of a previous physical layer fault by detecting of
      AIS.

   The exact conditions under which a PE enters and exits the AIS state
   are defined in [G.707] and [G.783].

10.1.4. AC transmit defect state entry/exit

   The AC transmit defect state is not valid for CEP PWs. RDI signals
   are forwarded transparently.

10.2. Consequent Actions

10.2.1. PW receive defect state entry/exit

   On entry to the PW receive defect state:

   a. PE1 MUST commence AIS-P/V insertion into the corresponding AC.
      See [RFC4842].
   b. PE1 MUST set the R bit in all PW packets sent back to PE2.

   On exit from the PW receive defect state:

   a. PE1 MUST cease AIS-P/V insertion into the corresponding AC.
   b. PE1 MUST clear the R bit in all PW packets sent back to PE2.

   See [RFC4842] for further details.

10.2.2. PW transmit defect state entry/exit

   On entry to the PW Transmit Defect State:

   a. A structure-aware PE1 MUST commence RDI insertion into the
      corresponding AC.

   On exit from the PW Transmit Defect State:

   a. A structure-aware PE1 MUST cease RDI insertion into the
      corresponding AC.

10.2.3. AC receive defect state entry/exit

   On entry to the AC receive defect state:

   a. PE1 MUST set the L bit in these packets.
   b. If Dynamic Bandwidth Allocation (DBA) has been enabled, PE1 MAY
      omit the payload in order to conserve bandwidth.
   c. If Dynamic Bandwidth Allocation (DBA) is not enabled PE1 SHOULD
      insert AIS-V/P in the SDH/SONET client layer in the PW packets
      sent towards PE2.

   On exit from the AC receive defect state:

   a. PE1 MUST cease overwriting PW content with AIS-P/V and return to
      forwarding valid data in PW packets sent towards PE2.
   b. PE1 MUST clear the L bit in PW packets sent towards PE2.

   See [RFC4842] for further details.

11. 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 RFC 5085 for VCCV apply to
   the OAM messages thus transferred. Security considerations
   applicable to the PWE3 control protocol of RFC 4447 Section 8.2 apply
   to OAM indications transferred using the LDP status message.

   Since the mechanisms of this document enable propagation of OAM
   messages and fault conditions between native service networks and
   PSNs, continuity of the end-to-end service depends on a trust
   relationship between the operators of these networks. Security
   considerations for such scenarios are discussed in Section 7 of
    [RFC5254].

12. IANA Considerations

   This document requires no IANA actions.

13. Contributors and Acknowledgments

   Mustapha Aissaoui, Peter Busschbach, Luca Martini, Monique Morrow,
   Thomas Nadeau, and Yaakov (J) Stein, were each, in turn, editors of
   one or more revisions of this document. All of the above were
   contributing authors, as was Dave Allan, david.i.allan@ericsson.com.

   The following contributed significant ideas or text:
      Matthew Bocci, matthew.bocci@alcatel-lucent.co.uk
      Simon Delord, Simon.A.DeLord@team.telstra.com
      Yuichi Ikejiri, y.ikejiri@ntt.com
      Kenji Kumaki, kekumaki@kddi.com
      Satoru Matsushima, satoru.matsushima@tm.softbank.co.jp
      Teruyuki Oya, teruyuki.oya@tm.softbank.co.jp
      Carlos Pignataro, cpignata@cisco.com
      Vasile Radoaca, vasile.radoaca@alcatel-lucent.com
      Himanshu Shah, hshah@ciena.com
      David Watkinson, david.watkinson@alcatel-lucent.com

   The editors would like to acknowledge the contributions of Bertrand
   Duvivier, Adrian Farrel, Tiberiu Grigoriu, Ron Insler, Michel
   Khouderchah, Vanson Lim, Amir Maleki, Neil McGill, Chris Metz, Hari
   Rakotoranto, Eric Rosen, Mark Townsley, and Ben Washam.

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.

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

   [RFC4447] Martini, L., Rosen, E., El-Aawar, N., Smith, T., and G.
              Heron, "Pseudowire Setup and Maintenance Using the Label
              Distribution Protocol (LDP)", RFC 4447, April 2006.

   [RFC4553] Vainshtein, A. and YJ. Stein, "Structure-Agnostic Time
              Division Multiplexing (TDM) over Packet (SAToP)",
              RFC 4553, June 2006.

   [RFC4591] Townsley, M., Wilkie, G., Booth, S., Bryant, S., and J.
              Lau, "Frame Relay over Layer 2 Tunneling Protocol Version
              3 (L2TPv3)", RFC 4591, August 2006.

   [RFC4619]   Martini, L., Kawa, C., and A. Malis, "Encapsulation
               Methods for Transport of Frame Relay over Multiprotocol
               Label Switching (MPLS) Networks", RFC 4619,
               September 2006.

   [RFC4717]   Martini, L., Jayakumar, J., Bocci, M., El-Aawar, N.,
               Brayley, J., and G. Koleyni, "Encapsulation Methods for
               Transport of Asynchronous Transfer Mode (ATM) over MPLS
               Networks", RFC 4717, December 2006.

   [RFC4842]   Malis, A., Pate, P., Cohen, R., and D. Zelig, "Synchronous
               Optical Network/Synchronous Digital Hierarchy (SONET/SDH)
               Circuit Emulation over Packet (CEP)", RFC 4842,
               April 2007.

   [RFC5036]   Andersson, L., Minei, I., and B. Thomas, "LDP
               Specification", RFC 5036, October 2007.

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

   [RFC5641]   McGill, N. and C. Pignataro, "Layer 2 Tunneling Protocol
               Version 3 (L2TPv3) Extended Circuit Status Values",
               RFC 5641, August 2009.

   [RFC5880]   Katz, D. and D. Ward, "Bidirectional Forwarding Detection
               (BFD)", RFC 5880, June 2010.

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

   [G.707]   "Network node interface for the synchronous digital
               hierarchy", ITU-T Recommendation G.707, December 2003.

   [G.775]   "Loss of Signal (LOS), Alarm Indication Signal(AIS) and
               Remote Defect Indication (RDI) defect detection and
               clearance criteria for PDH signals", ITU-T Recommendation
               G.775, October 1998.

   [G.783]   "Characteristics of synchronous digital hierarchy (SDH)
               equipment functional blocks", ITU-T Recommendation G.783,
               March 2006.

   [I.610]   "B-ISDN operation and maintenance principles and
               functions", ITU-T Recommendation I.610, February 1999.

   [Q.933]   "ISDN Digital Subscriber Signalling System No. 1 (DSS1)
              Signalling specifications for frame mode switched and
              permanent virtual connection control and status
              monitoring", ITU-T Recommendation Q.993, February 2003.

14.2. Informative References

   [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5,
              RFC 792, September 1981.

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

   [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V.,
              and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP
              Tunnels", RFC 3209, December 2001.

   [RFC3916] Xiao, X., McPherson, D., and P. Pate, "Requirements for
              Pseudo-Wire Emulation Edge-to-Edge (PWE3)", RFC 3916,
              September 2004.

   [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.

   [RFC4377] Nadeau, T., Morrow, M., Swallow, G., Allan, D., and S.
              Matsushima, "Operations and Management (OAM) Requirements
              for Multi-Protocol Label Switched (MPLS) Networks",
              RFC 4377, February 2006.

   [RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson,
              "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for
              Use over an MPLS PSN", RFC 4385, February 2006.

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

   [RFC4454] Singh, S., Townsley, M., and C. Pignataro, "Asynchronous
              Transfer Mode (ATM) over Layer 2 Tunneling Protocol
              Version 3 (L2TPv3)", RFC 4454, May 2006.

   [RFC5086] Vainshtein, A., Sasson, I., Metz, E., Frost, T., and P.
              Pate, "Structure-Aware Time Division Multiplexed (TDM)
              Circuit Emulation Service over Packet Switched Network
              (CESoPSN)", RFC 5086, December 2007.

   [RFC5087] Stein, Y(J)., Shashoua, R., Insler, R., and M. Anavi,

                "Time Division Multiplexing over IP (TDMoIP)", RFC 5087,
                December 2007.

   [RFC5254]   Bitar, N., Bocci, M., and L. Martini, "Requirements for
                Multi-Segment Pseudowire Emulation Edge-to-Edge (PWE3)",
                RFC 5254, October 2008.

   [RFC6073]   Martini, L., Metz, C., Nadeau, T., Bocci, M., and M.
                Aissaoui, "Segmented Pseudowire", RFC 6073, January 2011.

   [I-D.ietf-pwe3-mpls-eth-oam-iwk]
                Qiu, R., Mohan, D., Bitar, N., DeLord, S., Niger, P., and
                A. Sajassi, "MPLS and Ethernet OAM Interworking",
                draft-ietf-pwe3-mpls-eth-oam-iwk-04 (work in progress),
                March 2011.

   [I-D.ietf-pwe3-static-pw-status]
                Martini, L., Swallow, G., Heron, G., and M. Bocci,
                "Pseudowire Status for Static Pseudowires",
                draft-ietf-pwe3-static-pw-status-03 (work in progress),
                March 2011.

   [I.620]   "Frame relay operation and maintenance principles and
                functions", ITU-T Recommendation I.620, October 1996.

Appendix A. Native Service Management (informative)

A.1. Frame Relay Management

   The management of Frame Relay Bearer Service (FRBS) connections can
   be accomplished through two distinct methodologies:

   a. Based on ITU-T Q.933 Annex A, Link Integrity Verification
      procedure, where STATUS and STATUS ENQUIRY signaling messages are
      sent using DLCI=0 over a given UNI and NNI physical link.
   b. Based on FRBS LMI, and similar to ATM ILMI where LMI is common in
      private Frame Relay networks.

   In addition, ITU-T I.620 [I.620] addressed Frame Relay loopback.
   This Recommendation was withdrawn in 2004 and its deployment was
   limited.

   It is possible to use either, or both, of the above options to manage
   Frame Relay interfaces. This document will refer exclusively to
   Q.933 messages.

   The status of any provisioned Frame Relay PVC may be updated through:

   a. Frame Relay STATUS messages in response to Frame Relay STATUS
      ENQUIRY messages; these are mandatory.
   b. Optional unsolicited STATUS updates independent of STATUS ENQUIRY
      (typically under the control of management system, these updates
      can be sent periodically (continuous monitoring) or only upon
      detection of specific defects based on configuration.

   In Frame Relay, a Data Link Connection (DLC) is either up or down.
   There is no distinction between different directions. To achieve
   commonality with other technologies, down is represented as a receive
   defect.

   Frame relay connection management is not implemented over the PW
   using either of the techniques native to FR, therefore PW mechanisms
   are used to synchronize the view each PE has of the remote Native
   Service/Attachment Circuit (NS/AC). A PE will treat a remote NS/AC
   failure in the same way it would treat a PW or PSN failure; that is
   using AC facing FR connection management to notify the CE that FR is
   down.

A.2. ATM Management

   ATM management and OAM mechanisms are much more evolved than those of
   Frame Relay. There are five broad management-related categories,
   including fault management (FT), Performance management (PM),

configuration management (CM), Accounting management (AC), and
Security management (SM). [I.610] describes the functions for the
operation and maintenance of the physical layer and the ATM layer,
that is, management at the bit and cell levels. Because of its
scope, this document will concentrate on ATM fault management
functions. Fault management functions include the following:

a. Alarm indication signal (AIS).
b. Remote Defect indication (RDI).
c. Continuity Check (CC).
d. Loopback (LB).

Some of the basic ATM fault management functions are described as
follows: Alarm indication signal (AIS) sends a message in the same
direction as that of the signal, to the effect that an error has been
detected.

Remote defect indication (RDI) sends a message to the transmitting
terminal that an error has been detected. Alarms related to the
physical layer are indicated using path AIS/RDI. Virtual path AIS/
RDI and virtual channel AIS/RDI are also generated for the ATM layer.

OAM cells (F4 and F5 cells) are used to instrument virtual paths and
virtual channels respectively with regard to their performance and
availability. OAM cells in the F4 and F5 flows are used for
monitoring a segment of the network and end-to-end monitoring. OAM
cells in F4 flows have the same VPI as that of the connection being
monitored. OAM cells in F5 flows have the same VPI and VCI as that
of the connection being monitored. The AIS and RDI messages of the
F4 and F5 flows are sent to the other network nodes via the VPC or
the VCC to which the message refers. The type of error and its
location can be indicated in the OAM cells. Continuity check is
another fault management function. To check whether a VCC that has
been idle for a period of time is still functioning, the network
elements can send continuity-check cells along that VCC.

Appendix B. PW Defects and Detection tools

B.1. PW Defects

   Possible defects that impact PWs are the following:

   a. Physical layer defect in the PSN interface.
   b. PSN tunnel failure which results in a loss of connectivity between
      ingress and egress PE.
   c. Control session failures between ingress and egress PE.

   In case of an MPLS PSN and an MPLS/IP PSN there are additional
   defects:

   a. PW labeling error, which is due to a defect in the ingress PE, or
      to an over-writing of the PW label value somewhere along the LSP
      path.
   b. LSP tunnel Label swapping errors or LSP tunnel label merging
      errors in the MPLS network. This could result in the termination
      of a PW at the wrong egress PE.
   c. Unintended self-replication; e.g., due to loops or denial- of-
      service attacks.

B.2. Packet Loss

   Persistent congestion in the PSN or in a PE could impact the proper
   operation of the emulated service.

   A PE can detect packet loss resulting from congestion through several
   methods. If a PE uses the sequence number field in the PWE3 Control
   Word for a specific Pseudowire [RFC3985] and [RFC4385], it has the
   ability to detect packet loss. Translation of congestion detection
   to PW defect states is beyond the scope of this document.

   There are congestion alarms that are raised in the node and to the
   management system when congestion occurs. The decision to declare
   the PW Down and to select another path is usually at the discretion
   of the network operator.

B.3. PW Defect Detection Tools

   To detect the defects listed above, Service Providers have a variety
   of options available.

   Physical Layer defect detection and notification mechanisms include
   SONET/SDH Loss of Signal (LOS), Loss of Alignment (LOA), and AIS/RDI.

   PSN defect detection mechanisms vary according to the PSN type.

For PWs over L2TPv3/IP PSNs, with L2TP as encapsulation protocol, the
defect detection mechanisms described in [RFC3931] apply. These
include, for example, the keep-alive mechanism performed with Hello
messages for detection of loss of connectivity between a pair of
LCCEs (i.e., dead PE peer and path detection). Furthermore, the
tools Ping and Traceroute, based on ICMP Echo Messages [RFC0792]
apply and can be used to detect defects on the IP PSN. Additionally,
VCCV-Ping [RFC5085] and VCCV-BFD [RFC5885] can also be used to detect
defects at the individual pseudowire level.

For PWs over MPLS or MPLS/IP PSNs, several tools can be used:

a. LSP-Ping and LSP-Traceroute [RFC4379] for LSP tunnel connectivity
   verification.
b. LSP-Ping with Bi-directional Forwarding Detection [RFC5885] for
   LSP tunnel continuity checking.
c. Furthermore, if RSVP-TE is used to setup the PSN Tunnels between
   ingress and egress PE, the hello protocol can be used to detect
   loss of connectivity [RFC3209], but only at the control plane.

B.4. PW specific defect detection mechanisms

[RFC4377] describes how LSP-Ping and BFD can be used over individual
PWs for connectivity verification and continuity checking
respectively.

Furthermore, the detection of a fault could occur at different points
in the network and there are several ways the observing PE determines
a fault exists:

a. Egress PE detection of failure (e.g., BFD).
b. Ingress PE detection of failure (e.g., LSP-PING).
c. Ingress PE notification of failure (e.g. RSVP Path-err).

Authors' Addresses

    Mustapha Aissaoui
    Alcatel-Lucent
    600 March Rd
    Kanata, ON K2K 2E6
    Canada

    Email: mustapha.aissaoui@alcatel-lucent.com

    Peter Busschbach
    Alcatel-Lucent
    67 Whippany Rd
    Whippany, NJ 07981
    USA

    Email: busschbach@alcatel-lucent.com

    Luca Martini
    Cisco Systems, Inc.
    9155 East Nichols Avenue, Suite 400
    Englewood, CO 80112
    USA

    Email: lmartini@cisco.com

    Monique Morrow
    Cisco Systems, Inc.
    Richtistrase 7
    CH-8304 Wallisellen
    Switzerland

    Email: mmorrow@cisco.com

    Thomas Nadeau
    CA Technologies
    273 Corporate Dr.
    Portsmouth, NH 03801
    USA

    Email: Thomas.Nadeau@ca.com

      Yaakov (Jonathan) Stein
      RAD Data Communications
      24 Raoul Wallenberg St., Bldg C
      Tel Aviv 69719
      ISRAEL

      Email: yaakov_s@rad.com