PWE3 Working Group
INTERNET-DRAFT
Intended status: Standards Track
Expires: January 2013

Dinesh Mohan (Ed.)
Nortel Networks

Nabil Bitar (Ed.)
Verizon

Ali Sajassi (Ed.)
Cisco

July 16, 2012

MPLS and Ethernet OAM Interworking
draft-ietf-pwe3-mpls-eth-oam-iwk-06.txt

Abstract

This document specifies the mapping of defect states between Ethernet Attachment Circuits (ACs) and associated Ethernet Pseudowires (PWs) connected in accordance to the PWE3 architecture to realize an end-to-end emulated Ethernet service. It standardizes the behavior of Provider Edges (PEs) with respect to Ethernet PW and AC defects.

Status of this Memo

Table of Contents

1. Specification of Requirements

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

2. Introduction

This document specifies the mapping of defect states between Ethernet Attachment Circuits (ACs) and associated Ethernet Pseudowires (PWs) connected in accordance to the PWE3 architecture [RFC3985] to realize an end-to-end emulated Ethernet service. This document augments the mapping of defect states between a PW and associated AC of the end-to-end emulated service in [RFC6310]. It covers the following Ethernet OAM (Opertaions, Administration and Maintenance) mechanisms and their interworking with PW OAM mechanisms:

- Ethernet Continuity Check (CC) [802.1ag][Y.1731]
- Ethernet Alarm Indication Signaling (AIS) and Remote Defect Indication (RDI) [Y.1731]
- Ethernet Link OAM [802.3]
- Ethernet Local Management Interface {E-LMI} [MEF16]

Ethernet Link OAM [802.3] allows some Link defect states to be detected and communicated across an Ethernet Link. When an Ethernet

AC is an Ethernet physical port, there may be some application of Ethernet Link OAM [802.3]. Further, E-LMI [MEF16] also allows for some Ethernet Virtual Circuit (EVC) defect states to be communicated across an Ethernet UNI where Ethernet UNI constitutes a single hop Ethernet Link (i.e. without any 802.1Q/.1ad compliant bridges in between). There may be some application of E-LMI [MEF16] for failure notification across single hop Ethernet AC in certain deployments that specifically do not support [802.1ag] and/or [Y.1731]. [Y.1731] and [802.1ag] based mechanisms are applicable in all types of Ethernet ACs. Ethernet Link OAM [802.3] and E-LMI [MEF16] are optional and their applicability is called out, where applicable.

Native Service (NS) OAM may be transported transparently over the corresponding PW as user data. This is referred to as "the single emulated OAM loop" mode per [RFC6310]. For Ethernet, as an example, 802.1ag continuity check messages (CCMs) between two Maintenance End Points (MEPs) can be transported transparently as user data over the corresponding PW. At MEP locations, service failure is detected when CCMs are not received over an interval that is 3.5 times the local CCM transmission interval. This is one of the failure conditions detected via CC. MEP peers can exist between Customer Equipment (CE) pairs (MEPs of a given Maintenance Entity Group (MEG) reside on the CEs), PE pairs (the MEPs of a given MEG reside on the PEs), or between the CE and PE (the MEPs of a given MEG reside on the PE and CE), as long as the MEG domain nesting rules are maintained. It should be noted that Ethernet allows the definition of up to 8 MEG domains, each compromising of MEPs (down MEPs and UP MEPs) and Maintenance Intermediate Points (MIPs). These domains can be nested or touching. MEPs and MIPs generate and process messages in the same domain level. Thus, whenever in this document we refer to messages sent by a MEP or a MIP to a peer MEP or MIP, these MEPs and MIPs are in the same MEG domain level.

When interworking two networking domains, such as native Ethernet and PWs to provide an end-to-end emulated service, there is need to identify the failure domain and location even when a PE supports both the NS OAM mechanisms and the PW OAM mechanisms. In addition, scalability constraints may not allow running proactive monitoring, such as CCMs with transmission enabled, at a PE to detect the failure of an EVC across the PW domain. Thus, network-driven alarms generated upon failure detection in the NS or PW domain and their mappings to the other domain are needed. There are also cases where a PE may not be able to process NS OAM messages received on the PW even when such messages are defined, as in Ethernet case, necessitating the need for fault notification message mapping between the PW domain and the NS domain.

For Multi-Segment PWs (MS-PWs) [RFC5659], Switching PEs (S-PEs) are not aware of the NS. Thus, failure detection and notification at S-PEs will be based on PW OAM mechanisms. Mapping between PW OAM and NS OAM will be required at the Terminating PEs (T-PEs) to propagate the failure notification to the EVC endpoints.

Similar to [RFC6310], the intent of this document is to standardize the behavior of PEs with respect to failures on Ethernet ACs and PWs, so that there is no ambiguity about the alarms generated and consequent actions undertaken by PEs in response to specific failure conditions.

2.1. Reference Model and Defect Locations

Figure 1 is the same as used in [RFC6310] and is reproduced in this document as a reference to highlight defect locations.

```
            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

2.2. Abstract Defect States

Abstract Defect States are also introduced in [RFC6310]. This document uses the same conventions, as shown in Figure 2, from [RFC6310]. It may be noted however that CE devices, shown in Figure 2, do not necessarily have to be end customer devices. These are essentially devices in client network segments that are connecting to the Packet Switched Network (PSN) for the emulated services.

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

 (arrows indicate direction of user traffic impacted by a defect)
```

Figure 2: Transmit and Receive Defect States and Notifications

PE1 may detect a receive defect in a local Ethernet AC via one of the following mechanisms:

- An AIS alarm generated at an upstream node in the client domain (CE1 in Figure 2) and received by a local MEP.

- Failure of the local link on which the AC is configured. Link failure may be detected via physical failures (e.g., loss of signal (LoS)), via Ethernet Link OAM [802.3] critical link event notifications generated at an upstream node CE1 with "Dying Gasp" or "Critical Event" indication, or via a client Signal Fail message [Y.1731].

- Failure to receive CCMs on the AC if a local MEP is configured for the AC with CCM transmission enabled.

- A CCM from CE1 with interface status TLV indicating interface down. Other CCM interface status TLVs will not be used to indicate failure or recovery from failure.

It should be noted when a MEP at a PE or a CE receives a CCM with the wrong MEG ID, MEP ID, or MEP level, the receiving PE or CE

SHOULD treat such an event as an AC receive defect. In any case, if such events persist for 3.5 times the MEP local CCM transmission time, loss of continuity will be declared at the receiving end.

An AC receive defect at PE1 impacts the ability of PE1 to receive user traffic from the Client domain attached to PE1 via that AC.

PE1 may detect a receive defect in the PW via one of the following mechanisms:

- A Forward Defect notification received from PE2. This defect notification could point to problems associated with PE2's inability to transmit traffic on the PW or PE2's inability to receive traffic on its local AC from CE2.

- Unavailability of a PSN path in the PW domain to PE2.

A PW forward defect indication received on PE1 impacts the ability of PE1 to receive traffic on the PW.

PE1 may be notified of an AC transmit defect via one of the following mechanisms:

- CCMs with RDI (Remote Defect Indication) bit set.

- In case when the AC is associated with a physical port, failure of the local link on which the AC is configured (e.g., LOS or via Ethernet Link OAM [802.3] critical link event notifications generated at an upstream node CE1 with "Link Fault" indication).

An AC transmit defect impacts the ability of PE1 to send user traffic on the local AC.

Similarly, PE1 may be notified of a PW transmit defect via Reverse Defect indication from PE2, which could point to problems associated with PE2's inability to receive traffic on the PW or PE2's inability to transmit traffic on its local AC.

PW transmit defect impacts PE1 ability to send user traffic toward CE2. The procedures outlined in this document define the entry and exit criteria for each of the four defect states with respect to Ethernet ACs and corresponding PWs, and the consequent actions that PE1 must support to properly interwork these defect states and corresponding notification messages between the PW domain and the Native Service (NS) domain. Receive Defect state SHOULD have precedence over Transmit Defect state in terms of handling, when both transmit and receive defect states are identified simultaneously.

3. Terminology

This document uses the following terms:

- AIS: Alarm Indication Signal

- MD Level: Maintenance Domain (MD) Level which identifies a value in the range of 0-7 associated with Ethernet OAM frame. MD Level identifies the span of the Ethernet OAM frame.

- MEP: Maintenance End Point is responsible for origination and termination of OAM frames for a given MEG.

- MIP: Maintenance Intermediate Point is located between peer MEPs and can process OAM frames but does not initiate or terminate them.

- RDI: Remote Defect Indication.

Further, this document also uses the terminology and conventions used in [RFC6310].

4. PW Status and Defects

[RFC6310] introduces a range of defects that impact PW status. All these defect conditions are applicable for Ethernet PWs.

Similarly, there are different mechanisms described in [RFC6310] to detect PW defects, depending on the PSN type (e.g. MPLS PSN, MPLS-IP PSN). Any of these mechanisms can be used when monitoring the state of Ethernet PWs. [RFC6310] also discusses the applicability of these failure detection mechanisms.

4.1. Use of Native Service (NS) Notification

When a MEP is defined on the PE and associated with an Ethernet PW, the PE can use native service OAM capabilities for failure notifications. Options include:

- Sending of AIS frames from the local MEP to the MEP on the remote PE when the MEP needs to convey PE receive defects, and when CCM transmission is disabled.

- Suspension of CCM frames transmission from the local MEP to the peer MEP on the remote PE to convey PE receive defects, when CCM transmission is enabled.

- Setting the RDI bit in transmitted CCM frames, when loss of CCMs from the peer MEP is detected or the PE needs to convey PW reverse defects.

These NS OAM notifications are inserted into the corresponding PW.

Similarly, when the defect conditions are cleared, a PE can take one of the following actions, depending on the mechanism that was used for failure notification, to clear the defect sate on the peer PE:

- Stopping AIS frame transmission from the local MEP to the MEP on the remote PE to clear PW receive defects.

- Resuming CCM frames transmission from the local MEP to the peer MEP on the remote PE to clear PW forward defects notification, when CCM transmission is enabled.

- Clearing the RDI bit in transmitted CCM frames, to clear PW transmit defects notification, when CCM transmission is enabled.

4.2. Use of PW Status notification for MPLS PSNs

When PWs are established using LDP, LDP status notification signaling MUST be used as the default mechanism to signal AC and PW status and defects [RFC4447]. That is known as the "coupled loop mode". For PWs established over an MPLS or MPLS-IP PSN using other mechanisms (e.g. static configuration), inband signaling using VCCV-BFD [RFC5885] SHOULD be used to convey AC and PW status and defects.

[RFC6310] identifies the following PW defect status codepoints:

- Forward defect: corresponds to a logical OR of local AC (ingress) Receive fault, local PSN-facing PW (egress) transmit fault, and PW not forwarding fault.

- Reverse defect: corresponds to a logical OR of local AC (egress) transmit fault and local PW PSN-facing (ingress) receive fault.

There are also scenarios where a PE carries out PW label withdrawal instead of PW status notification. These include administrative disablement of the PW or loss of Target LDP session with the peer PE.

4.3. Use of BFD Diagnostic Codes

When using VCCV, the control channel (CC) type and Connectivity Verification (CV) Type are agreed on between the peer PEs using the OAM capability sub-TLV signaled as part of the interface parameter TLV when using FEC 129 and the interface parameter sub-TLV when using FEC 128.

As defined in [RFC6310], when CV type of 0x04 0r 0x10 is used to indicate that BFD is used for PW fault detection only, PW defect is detected via the BFD session while other defects, such as AC defect or PE internal defects preventing it from forwarding traffic, are communicated via LDP Status notification message in MPLS and MPLS-IP PSNs or other mechanisms in L2TP-IP PSN.

Similarly, when CV type of 0x08 or 0x20 is used to indicate that BFD is used for both PW fault detection and AC/PW Fault Notification, all defects are signaled via BFD.

5. Ethernet AC Defect States Entry or Exit Criteria

5.1. AC Receive Defect State Entry or Exit

PE1 enters the AC Receive Defect state if any of the following conditions is met:

- It detects or is notified of a physical layer fault on the Ethernet interface. Ethernet link failure can be detected based on loss of signal (LoS) or via Ethernet Link OAM [802.3] critical link event notifications generated at an upstream node CE1 with "Dying Gasp" or "Critical Event" indication.

- A MEP associated with the local AC receives an Ethernet AIS frame.

- A MEP associated with the local AC does not receive CCM frames from the peer MEP in the client domain (e.g. CE1) within an interval equal to 3.5 times the CCM transmission period configured for the MEP. This is the case when CCM transmission is enabled.

- A CCM with interface status TLV indicating interface down. Other CCM interface status TLVs will not be used to indicate failure or recovery from failure.

PE1 exits the AC Receive Defect state if all of the conditions that resulted in entering the defect state are cleared. This includes all of the following conditions:

- Any physical layer fault on the Ethernet interface, if detected or notified previously is removed (e.g., loss of signal(LoS) cleared, or Ethernet Link OAM [802.3] critical link event notifications with "Dying Gasp" or "Critical Event" indication cleared at an upstream node CE1).

- A MEP associated with the local AC does not receive any Ethernet AIS frame within a period indicated by previously received AIS, if AIS resulted in entering the defect state.

- A MEP associated with the local AC and configured with CCM enabled receives a configured number (e.g., 3 or more) of consecutive CCM frames from the peer MEP on CE1 within an interval equal to a multiple (3.5) of the CCM transmission period configured for the MEP.

- CCM indicates interface status up.

5.2. AC Transmit Defect State Entry or Exit

PE1 enters the AC Transmit Defect state if any of the following conditions is met:

- It detects or is notified of a physical layer fault on the Ethernet interface (e.g., via loss of signal (LoS) or Ethernet Link OAM [802.3] critical link event notifications generated at an upstream node CE1 with "Link Fault" indication).

- A MEP configured with CCM transmission enabled and associated with the local AC receives a CCM frame, with its RDI bit set, from the peer MEP in the client domain (e.g., CE1).

PE1 exits the AC Transmit Defect state if all of the conditions that resulted in entering the defect state are cleared. This includes all of the following conditions:

- Any physical layer fault on the Ethernet interface, if detected or notified previously is removed (e.g., LOS cleared, Ethernet Link OAM [802.3] critical link event notifications with "Link Fault" indication cleared at an upstream node CE1).

- A MEP configured with CCM transmission enabled and associated with the local AC does not receive a CCM frame with RDI bit set, having received a previous CCM frame with RDI bit set from the peer MEP in the client domain (e.g. CE1).

6. Ethernet AC and PW Defect States Interworking

6.1. PW Receive Defect Entry Procedures

When the PW status on PE1 transitions from working to PW Receive Defect state, PE1's ability to receive user traffic from CE2 is impacted. As a result, PE1 needs to notify CE1 about this problem.

Upon entry to the PW Receive Defect state, the following must be done:

- If PE1 is configured with a down MEP associated with the local AC and CCM transmission is not enabled, the MEP associated with the AC must transmit AIS frames periodically to the peer MEP in the client domain (e.g., on CE1) based on configured AIS transmission period.

- If PE1 is configured with a down MEP associated with the local AC and CCM transmission is enabled, and the MEP associated with the AC is configured to support Interface Status TLV in CCM messages, the MEP associated with the AC must transmit CCM frames with Interface Status TLV as being down to the peer MEP in the client domain (e.g., on CE1).

- If PE1 is configured with a down MEP associated with the local AC and CCM transmission is enabled, and the MEP associated with the AC is configured to not support Interface Status TLV in CCM messages, the MEP associated with the AC must stop transmitting CCM frames to the peer MEP in the client domain (e.g., on CE1).

- If PE1 is configured to run E-LMI [MEF16] with CE1 and if E-LMI is used for failure notification, PE1 must transmit E-LMI asynchronous STATUS message with report type Single EVC Asynchronous Status indicating that PW is Not Active.

Further, when PE1 enters the Receive Defect state, it must assume that PE2 has no knowledge of the defect and must send reverse defect failure notification to PE2. For MPLS PSN or MPLS-IP PSN, this is done via either a PW Status notification message indicating a reverse defect; or via VCCV-BFD diagnostic code of reverse defect if VCCV CV type of 0x08 had been negotiated. When Native Service OAM mechanism is supported on PE1, it can also use the NS OAM notification as specified in Section 4.1.

If PW receive defect is entered as a result of a forward defect notification from PE2 or via loss of control adjacency, no additional action is needed since PE2 is expected to be aware of the defect.

6.2. PW Receive Defect Exit Procedures

When the PW status transitions from PW Receive Defect state to working, PE1's ability to receive user traffic from CE2 is restored. As a result, PE1 needs to cease defect notification to CE1 by performing the following:

- If PE1 is configured with a down MEP associated with the local AC and CCM transmission is not enabled, the MEP associated with the AC must stop transmitting AIS frames towards the peer MEP in the client domain (e.g., on CE1).

- If PE1 is configured with a down MEP associated with the local AC and CCM transmission is enabled, and the MEP associated with the AC is configured to support Interface Status TLV in CCM messages, the MEP associated with the AC must transmit CCM frames with Interface Status TLV as being Up to the peer MEP in the client domain (e.g., on CE1).

- If PE1 is configured with a down MEP associated with the local AC and CCM transmission is enabled, and the MEP associated with the AC is configured to not support Interface Status TLV in CCM messages, the MEP associated with the AC must resume transmitting CCM frames to the peer MEP in the client domain (e.g., on CE1).

- If PE1 is configured to run E-LMI [MEF16] with CE1 and E-LMI is used for fault notification, PE1 must transmit E-LMI asynchronous STATUS message with report type Single EVC Asynchronous Status indicating that PW is Active.

Further, if the PW receive defect was explicitly detected by PE1, it must now notify PE2 about clearing of Receive Defect state by clearing reverse defect notification. For PWs over MPLS PSN or MPLS-IP PSN, this is either done via PW Status message indicating working; or via VCCV-BFD diagnostic code if VCCV CV type of 0x08/0x20 had been negotiated. When Native Service OAM mechanism is supported on PE, it can also clear the NS OAM notification specified in Section 4.1.

If PW receive defect was established via notification from PE2 or via loss of control adjacency, no additional action is needed, since PE2 is expected to be aware of the defect clearing.

6.3. PW Transmit Defect Entry Procedures

When the PW status transitions from working to PW Transmit Defect state, PE1's ability to transmit user traffic to CE2 is impacted. As a result, PE1 needs to notify CE1 about this problem which has been detected by PE1.

Upon entry to the PW Transmit Defect state, the following must be done:

- If PE1 is configured with a down MEP associated with the local AC and CCM transmission is enabled, the MEP associated with the AC MUST set the RDI bit in transmitted CCM frames or send status TLV with interface down to the peer MEP in the client domain (e.g., on CE1).

- If PE1 is configured to run E-LMI [MEF16] with CE1 and E-LMI is used for fault notification, PE1 must transmit E-LMI asynchronous STATUS message with report type Single EVC Asynchronous Status indicating that PW is Not Active.

6.4. PW Transmit Defect Exit Procedures

When the PW status transitions from PW Transmit Defect state to working, PE1's ability to transmit user traffic to CE2 is restored. As a result, PE1 needs to cease defect notifications to CE1 and perform the following:

- If PE1 is configured with a down MEP associated with the local AC and CCM transmission is enabled, the MEP associated with the AC must clear the RDI bit in the transmitted CCM frames to the peer MEP (e.g., on CE1).

- If PE1 is configured to run E-LMI [MEF16] with CE1, PE1 must transmit E-LMI asynchronous STATUS message with report type Single EVC Asynchronous Status indicating that PW is Active.

6.5. AC Receive Defect Entry Procedures

When AC status transitions from working to AC Receive Defect state, PE1's ability to receive user traffic from CE1 is impacted. As a result, PE1 needs to notify PE2 and CE1 about this problem.

If the AC receive defect is detected by PE1, it must notify PE2 in the form of a forward defect notification.

When NS OAM is not supported on PE1, and for PW over MPLS PSN or MPLS-IP PSN, forward defect notification is done via either PW Status message indicating a forward defect or via VCCV-BFD diagnostic code of forward defect if VCCV CV type of 0x08/0x20 had been negotiated.

When Native Service OAM mechanism is supported on PE1, it can also use the NS OAM notification as specified in Section 4.1.

In addition to the above actions, PE1 must perform the following:

- If PE1 is configured with a down MEP associated with the local AC and CCM transmission is enabled, the MEP associated with the AC must set the RDI bit in transmitted CCM frames.

6.6. AC Receive Defect Exit Procedures

When AC status transitions from AC Receive Defect to working, PE1's ability to receive user traffic from CE1 is restored. As a result, PE1 needs to cease defect notifications to PE2 and CE1 and perform the following:

- When NS OAM is not supported on PE1 and for PW over MPLS PSN or MPLS-IP PSN, forward defect notification is cleared via PW Status message indicating a working state; or via VCCV-BFD diagnostic code if VCCV CV type of 0x08 or 0x20 had been negotiated.

- When Native Service OAM mechanism is supported on PE1, PE1 clears the NS OAM notification as specified in Section 4.1.

- If PE1 is configured with a down MEP associated with the local AC and CCM transmission is enabled, the MEP associated with the AC must clear the RDI bit in transmitted CCM frames to the peer MEP in the client domain (e.g., on CE1).

6.7. AC Transmit Defect Entry Procedures

When AC status transitions from working to AC Transmit Defect, PE1's ability to transmit user traffic to CE1 is impacted. As a result, PE1 needs to notify PE2 about this problem.

If the AC transmit defect is detected by PE1, it must notify PE2 in the form of a reverse defect notification.

When NS OAM is not supported on PE1, in PW over MPLS PSN or MPLS-IP PSN, reverse defect notification is either done via PW Status message indicating a reverse defect; or via VCCV-BFD diagnostic code of reverse defect if VCCV CV type of 0x08 or 0x20 had been negotiated.

When Native Service OAM mechanism is supported on PE1, it can also use the NS OAM notification as specified in Section 4.1.

6.8. AC Transmit Defect Exit Procedures

When AC status transitions from AC Transmit defect to working, PE1's ability to transmit user traffic to CE1 is restored. As a result, PE1 must clear reverse defect notification to PE2.

When NS OAM is not supported on PE1 and for PW over MPLS PSN or MPLS-IP PSN, reverse defect notification is cleared via either a PW Status message indicating a working state or via VCCV-BFD diagnostic code if VCCV CV type of 0x08 had been negotiated.

When Native Service OAM mechanism is supported on PE1, PE1 can clear NS OAM notification as specified in Section 4.1.

7. Security Considerations

The OAM interworking mechanisms 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.

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].

8. IANA Considerations

This document has no actions for IANA.

9. Acknowledgments

The authors are thankful to Samer Salam, Matthew Bocci and Yaakov Stein for their valuable comments.

10. References

10.1. Normative References

[RFC6310] "Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping", RFC 6310, July 2011.

[Y.1731] "OAM Functions and mechanisms for Ethernet based networks", ITU-T Y.1731, May 2006.

[802.1ag] "Connectivity Fault Management", IEEE 802.1ag, December 2007.

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

[RFC4447] "Pseudowire Setup and Maintenance using LDP", RFC4447, April 2006.

[RFC5885] "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)", RFC5885, June 2010.

[802.3] "CDMA/CD access method and physical layer specifications", Clause 57 for Operations, Administration and Maintenance, 2005.

[MEF16] "Ethernet Local Management Interface", MEF16, January 2006.

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

10.2. Informative References

[RFC3985] "Pseudo Wire Emulation Edge-to-Edge(PWE3) Architecture", RFC 3985, April 2005.

[RFC5659] "An Architecture for Multi-Segment Pseudo Wire Emulation Edge-to-Edge", RFC5659, October 2009.

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

11. Appendix A: Ethernet Native Service Management

Ethernet OAM mechanisms are broadly classified into two categories: Fault Management (FM) and Performance Monitoring (PM). ITU-T Y.1731 provides coverage for both FM and PM while IEEE 802.1ag provides coverage for a sub-set of FM functions.

Ethernet OAM also introduces the concept of Maintenance Entity (ME) which is used to identify the entity that needs to be managed. An ME is inherently a point-to-point association. However, in case of a multipoint association, Maintenance Entity Group (MEG) consisting of different MEs is used. IEEE 802.1 uses the concept of Maintenance Association (MA) which is used to identify both point-to-point and multipoint associations. Each MA consists of Maintenance End Points (MEPs) which are responsible for originating OAM frames. In between the MEPs, there can also be Maintenance Intermediate Points (MIPs) which do not originate OAM frames however do respond to OAM frames from MEPs.

Ethernet OAM allows for hierarchical maintenance entities to allow for simultaneous end-to-end and segment monitoring. This is achieved by having a provision of up to 8 Maintenance Domain Levels (MD Levels) where each MEP or MIP is associated with a specific MD Level.

It is important to note that the common set of FM mechanisms between IEEE 802.1ag and ITU-T Y.1731 are completely compatible.

The common FM mechanisms include:

1) Continuity Check Messages (CCM)

2) Loopback Message (LBM) and Loopback Reply (LBR)

3) Linktrace Message (LTM) and Linktrace Reply (LTR)

CCM messages are used for fault detection including misconnections and mis-configurations. Typically CCM messages are sent as multicast frames or Unicast frames and also allow RDI notifications. LBM/LBR are used to perform fault verification, while also allow for MTU verification and CIR/EIR measurements. LTM/LTR can be used for discovering the path traversed between a MEP and another target MIP/MEP in the same MA. LTM/LTR also allow for fault localization.

In addition, ITU-T Y.1731 also specifies the following FM functions:

4) Alarm Indication Signal (AIS)

AIS allows for fault notification to downstream and upstream nodes.

Further, ITU-T Y.1731 also specifies the following PM functions:

5) Loss Measurement Message (LMM) and Reply (LMR)

6) Delay Measurement Message (DMR) and Reply (DMR)

7) 1-way Delay Message (1DM)

While LMM/LMR is used to measure Frame Loss Ratio (FLR), DMM/DMR is used to measure single-ended (aka two-way) Frame Delay (FD) and Frame Delay Variation (FDV, also known as Jitter). 1DM can be used for dual-ended (aka one-way) FD and FDV measurements.

Authors' Addresses

Dinesh Mohan
Nortel
3500 Carling Ave
Ottawa, ON K2H8E9
Email: dinmohan@hotmail.com

Nabil Bitar
Verizon
60 Sylvan Road
Waltham, MA 02145
Email: nabil.n.bitar@verizon.com

Ali Sajassi
Cisco
170 West Tasman Drive
San Jose, CA 95134, US
Email: sajassi@cisco.com

Simon Delord
Alcatel-Lucent
215 Spring Street
Melbourne, Australia
E-mail: simon.delord@gmail.com

Philippe Niger
France Telecom
2 av. Pierre Marzin
22300 LANNION, France
E-mail: philippe.niger@francetelecom.com

Ray Qiu
Huawei
330 Central Expressway
Santa Clara, CA 95050, US

Email: ray@huawei.com