

Internet Engineering Task Force (IETF)
Request for Comments: 7023
Category: Standards Track
ISSN: 2070-1721

D. Mohan, Ed.
Nortel Networks
N. Bitar, Ed.
Verizon
A. Sajassi, Ed.
Cisco
S. Delord
Alcatel-Lucent
P. Niger
France Telecom
R. Qiu
Juniper
October 2013

MPLS and Ethernet Operations, Administration, and Maintenance (OAM)
Interworking

Abstract

This document specifies the mapping of defect states between Ethernet Attachment Circuits (ACs) and associated Ethernet pseudowires (PWs) connected in accordance with the Pseudowire Emulation Edge-to-Edge (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

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7023>.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
1.1. Specification of Requirements	4
2. Overview	4
2.1. Reference Model and Defect Locations	6
2.2. Abstract Defect States	6
3. Abbreviations and Terminology	7
3.1. Abbreviations	7
3.2. Terminology	8
4. PW Status and Defects	9
4.1. Use of Native Service (NS) Notification	9
4.2. Use of PW Status Notification for MPLS PSNs	10
4.3. Use of BFD Diagnostic Codes	10
4.4. PW Defect States Entry and Exit Criteria	11
4.4.1. PW Receive Defect State Entry and Exit	11
4.4.2. PW Transmit Defect State Entry and Exit	11
5. Ethernet AC Defect States Entry and Exit Criteria	12
5.1. AC Receive Defect State Entry and Exit	12
5.2. AC Transmit Defect State Entry and Exit	13
6. Ethernet AC and PW Defect States Interworking	14
6.1. PW Receive Defect State Entry Procedures	14
6.2. PW Receive Defect State Exit Procedures	15
6.3. PW Transmit Defect State Entry Procedures	16
6.4. PW Transmit Defect State Exit Procedures	16
6.5. AC Receive Defect State Entry Procedures	16
6.6. AC Receive Defect State Exit Procedures	17
6.7. AC Transmit Defect State Entry Procedures	17
6.8. AC Transmit Defect State Exit Procedures	18
7. Security Considerations	18
8. Acknowledgments	19
9. References	19
9.1. Normative References	19
9.2. Informative References	20
Appendix A. Ethernet Native Service Management	21

1. Introduction

RFC 6310 [RFC6310] specifies the mapping and notification of defect states between a pseudowire (PW) and the Attachment Circuit (AC) of the end-to-end emulated service. It standardizes the behavior of Provider Edges (PEs) with respect to PW and AC defects for a number of technologies (e.g., Asynchronous Transfer Mode (ATM) and Frame Relay (FR)) emulated over PWs in MPLS and MPLS/IP Packet Switched Networks (PSNs). However, [RFC6310] does not describe this function for the Ethernet PW service owing to its unique characteristics.

This document specifies the mapping of defect states between ACs and associated Ethernet PWs connected in accordance with 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]. 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.

1.1. 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. Overview

There are a number of Operations, Administration, and Maintenance (OAM) technologies defined for Ethernet, providing various functionalities. This document covers the following Ethernet OAM mechanisms and their interworking with PW OAM mechanisms:

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

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 User Network Interface (UNI) where Ethernet UNI constitutes a single-hop Ethernet link (i.e., without any bridges

compliant with IEEE 802.1Q/.1ad in between). There may be some application of E-LMI [MEF16] for failure notification across single-hop Ethernet ACs in certain deployments that specifically do not support IEEE Connectivity Fault Management [CFM] and/or ITU-T Y.1731 [Y.1731], simply referred to as CFM and Y.1731, respectively, in this document. Mechanisms based on Y.1731 and CFM are applicable in all types of Ethernet ACs. Ethernet Link OAM and E-LMI 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, CFM continuity check messages (CCMs) between two Maintenance Entity Group 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 continuity check. MEP peers can exist between customer edge (CE) endpoints (MEPs of a given Maintenance Entity Group (MEG) reside on the CEs), between 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 level nesting rules are maintained. It should be noted that Ethernet allows the definition of up to 8 MEG levels, each comprised of MEPs (Down MEPs and Up MEPs) and Maintenance Entity Group Intermediate Points (MIPs). These levels can be nested or touching. MEPs and MIPs generate and process messages in the same MEG level. Thus, in this document, when 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 level.

When interworking two networking domains, such as native Ethernet and PWs to provide an end-to-end emulated service, there is a 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 the case of Ethernet, 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 end points.

2.1. Reference Model and Defect Locations

Figure 1 was used in [RFC6310]; it is reproduced in this document as a reference to highlight defect locations.

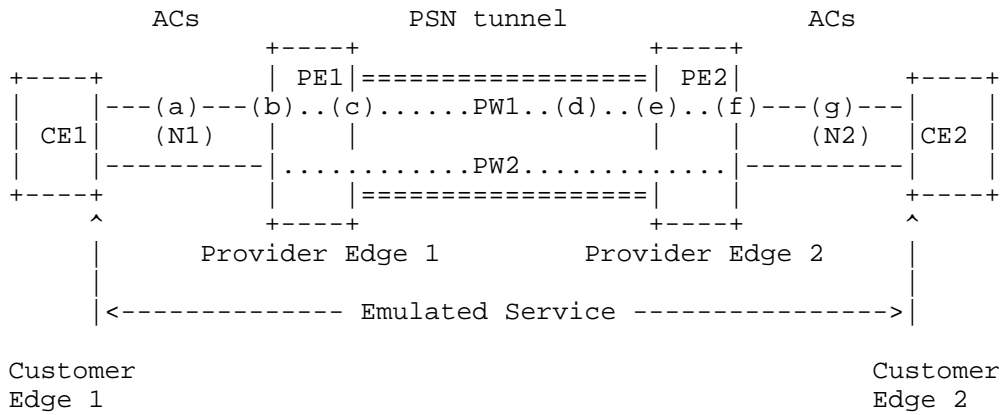



Figure 1: PWE3 Network Defect Locations

2.2. Abstract Defect States

Abstract defect states are also introduced in [RFC6310]. As shown in Figure 2, this document uses the same conventions as [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.



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

Figure 2: Transmit and Receive Defect States and Notifications

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; this document also defines 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.

Following is a summary of the defect states from the viewpoint of PE1 in Figure 2:

- A PW receive defect at PE1 impacts PE1's ability to receive traffic on the PW. Entry and exit criteria for the PW receive defect state are described in Section 4.4.1.
- A PW transmit defect at PE1 impacts PE1's ability to send user traffic toward CE2. PE1 MAY be notified of a PW transmit defect via a 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. Entry and exit criteria for the PW transmit defect state are described in Section 4.4.2.
- An AC receive defect at PE1 impacts PE1's ability to receive user traffic from the client domain attached to PE1 via that AC. Entry and exit criteria for the AC receive defect state are described in Section 5.1.
- An AC transmit defect at PE1 impacts PE1's ability to send user traffic on the local AC. Entry and exit criteria for the AC transmit defect state are described in Section 5.2.

3. Abbreviations and Terminology

3.1. Abbreviations

AC	Attachment Circuit
AIS	Alarm Indication Signal
BFD	Bidirectional Forwarding Detection
CC	Continuity Check
CCM	Continuity Check Message
CE	Customer Edge
CV	Connectivity Verification
E-LMI	Ethernet Local Management Interface
EVC	Ethernet Virtual Circuit
LDP	Label Distribution Protocol
LoS	Loss of Signal
MA	Maintenance Association
MD	Maintenance Domain
ME	Maintenance Entity
MEG	Maintenance Entity Group
MEP	MEG End Point
MIP	MEG Intermediate Point
MPLS	Multiprotocol Label Switching
MS-PW	Multi-Segment Pseudowire
NS	Native Service
OAM	Operations, Administration, and Maintenance
PE	Provider Edge
PSN	Packet Switched Network
PW	Pseudowire
RDI	Remote Defect Indication when used in the context of CCM
RDI	Reverse Defect Indication when used to semantically refer to defect indication in the reverse direction
S-PE	Switching Provider Edge
T-PE	Terminating Provider Edge
TLV	Type-Length Value
VCCV	Virtual Circuit Connectivity Verification

3.2. Terminology

This document uses the following terms with corresponding definitions:

- MEG Level: identifies a value in the range of 0-7 associated with an Ethernet OAM frame. MEG level identifies the span of the Ethernet OAM frame.
- MEG End Point (MEP): is responsible for origination and termination of OAM frames for a given MEG.

- MEG Intermediate Point (MIP): is located between peer MEPs and can process OAM frames but does not initiate them.
- MPLS PSN: a PSN that makes use of MPLS Label-Switched Paths [RFC3031] as the tunneling technology to forward PW packets.
- MPLS/IP PSN: a PSN that makes use of MPLS-in-IP tunneling [RFC4023] to tunnel MPLS-labeled PW packets over IP tunnels.

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 or 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 two PEs terminate an Ethernet PW with associated MEPs, each PE can use native service (NS) OAM capabilities for failure notifications by transmitting appropriate NS OAM messages over the corresponding PW to the remote PE. 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.
- Suspending transmission of CCM frames 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 when the PE needs to convey PW reverse defects.

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 state 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 transmission of CCM frames from the local MEP to the peer MEP on the remote PE to clear PW forward defect notification when CCM transmission is enabled.
- Clearing the RDI bit in transmitted CCM frames to clear PW transmit defect notification when CCM transmission is enabled.

4.2. Use of PW Status Notification for MPLS PSNs

RFC 4447 [RFC4447] specifies that for PWs that have been set up using the Label Distribution Protocol (LDP), the default mechanism to signal status and defects for ACs and PWs is the LDP status notification message. For PWs established over an MPLS or MPLS/IP PSN using other mechanisms (e.g., static configuration), in-band signaling using VCCV-BFD [RFC5885] SHOULD be used to convey AC and PW status and defects. Alternatively, the mechanisms defined in [RFC6478] MAY be used.

[RFC6310] identifies the following PW defect status code points:

- Forward defect: corresponds to a logical OR of Local Attachment Circuit (ingress) Receive Fault, Local PSN-facing PW (egress) Transmit Fault, and Pseudowire Not Forwarding fault.
- Reverse defect: corresponds to a logical OR of Local Attachment Circuit (egress) Transmit Fault and Local PSN-facing PW (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 the Target LDP session with the peer PE.

4.3. Use of BFD Diagnostic Codes

When using VCCV, the control channel type and Connectivity Verification (CV) type are agreed on between the peer PEs using the VCCV parameter field signaled as a sub-TLV of the interface parameters TLV when using FEC 129 and the interface parameter sub-TLV when using FEC 128 [RFC5085].

As defined in [RFC6310], when a CV type of 0x04 or 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 an LDP status notification message in MPLS and MPLS/IP PSNs or other mechanisms in L2TP/IP PSNs.

Similarly, when a 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.

4.4. PW Defect States Entry and Exit Criteria

4.4.1. PW Receive Defect State Entry and Exit

As described in Section 6.2.1 of [RFC6310], PE1 will enter the PW receive defect state if one or more of the following occur:

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

Note that if the PW LDP control session 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.

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

4.4.2. PW Transmit Defect State Entry and Exit

PE1 will enter the PW transmit defect state if the following conditions occur:

- It receives a Reverse Defect Indication (RDI) from PE2 either indicating a transmit fault on the remote AC or indicating that PE2 detected or was notified of an upstream PW fault.
- 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 if it has entered the PW receive defect state.

5. Ethernet AC Defect States Entry and Exit Criteria

5.1. AC Receive Defect State Entry and 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 or via a client Signal Fail message [Y.1731].
- A MEP associated with the local AC receives an Ethernet AIS frame from CE1.
- 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 has an 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 that 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.

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 where PE1 was 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" indications 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 and 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 where the AC is configured (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 (Remote Defect Indication) 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 where PE1 was notified previously, is removed (e.g., LoS cleared or 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 the RDI bit set, having received a previous CCM frame with the 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 State 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 the configured AIS transmission period.
- If PE1 is configured with a Down MEP associated with the local AC, CCM transmission is enabled, and the MEP associated with the AC is configured to support the Interface Status TLV in CCMs, the MEP associated with the AC MUST transmit CCM frames with the 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, CCM transmission is enabled, and the MEP associated with the AC is configured to not support the Interface Status TLV in CCMs, 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 an E-LMI asynchronous STATUS message with report type Single EVC Asynchronous Status indicating that the 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 a reverse defect failure notification to PE2. For MPLS PSN or MPLS/IP PSN, this is either done via a PW status notification message indicating a reverse defect or done via a VCCV-BFD diagnostic code of reverse defect if a VCCV CV type of 0x08 or 0x20 had been negotiated. When a 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 state 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 State 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, CCM transmission is enabled, and the MEP associated with the AC is configured to support the Interface Status TLV in CCMs, the MEP associated with the AC MUST transmit CCM frames with the 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, CCM transmission is enabled, and the MEP associated with the AC is configured to not support the Interface Status TLV in CCMs, 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 an E-LMI asynchronous STATUS message with report type Single EVC Asynchronous Status indicating that the 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 the reverse defect notification. For PW over MPLS PSN or MPLS/IP PSN, this is either done via a PW status message indicating a working state or done via a VCCV-BFD diagnostic code if a VCCV CV type of 0x08 or 0x20 had been negotiated. When a native service OAM mechanism is supported on PE, it can also clear the NS OAM notification as 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 State 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.

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 a 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 an E-LMI asynchronous STATUS message with report type Single EVC Asynchronous Status indicating that the PW is Not Active.
- If the PW failure was detected by PE1 without receiving a reverse defect notification from PE2, PE1 MUST assume PE2 has no knowledge of the defect and MUST notify PE2 by sending an FDI.

6.4. PW Transmit Defect State 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 or send a status TLV with interface up to the peer MEP in the client domain (e.g., on CE1).
- If PE1 is configured to run E-LMI [MEF16] with CE1, PE1 MUST transmit an E-LMI asynchronous STATUS message with report type Single EVC Asynchronous Status indicating that the PW is Active.
- PE1 MUST clear the FDI to PE2, if applicable.

6.5. AC Receive Defect State 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, in PW over MPLS PSN or MPLS/IP PSN, a forward defect notification is either done via a PW status message indicating a forward defect or done via a VCCV-BFD diagnostic code of forward defect if a VCCV CV type of 0x08 or 0x20 had been negotiated.

When a 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 State Exit Procedures

When AC status transitions from AC receive defect state 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, in PW over MPLS PSN or MPLS/IP PSN, the forward defect notification is cleared via a PW status message indicating a working state or via a VCCV-BFD diagnostic code if a VCCV CV type of 0x08 or 0x20 had been negotiated.
- When a 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 State Entry Procedures

When AC status transitions from working to AC transmit defect state, 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, a reverse defect notification is either done via a PW status message indicating a reverse defect or done via a VCCV-BFD diagnostic code of reverse defect if a VCCV CV type of 0x08 or 0x20 had been negotiated.

When a 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 State Exit Procedures

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

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

When a 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.

The security considerations in Section 10 of [RFC5085] for VCCV apply to the OAM messages thus transferred. Security considerations applicable to the PWE3 control protocol as described in Section 8.2 of [RFC4447] 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. Acknowledgments

The authors are thankful to Samer Salam, Matthew Bocci, Yaakov Stein, David Black, Lizhong Jin, Greg Mirsky, Huub van Helvoort, and Adrian Farrel for their valuable input and comments.

9. References

9.1. Normative References

- [802.3] IEEE, "Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications (Clause 57 for Operations, Administration, and Maintenance)", IEEE Std 802.3-2005, December 2005.
- [CFM] IEEE, "Connectivity Fault Management clause of IEEE 802.1Q", IEEE 802.1Q, 2013.
- [MEF16] Metro Ethernet Forum, "Ethernet Local Management Interface (E-LMI)", Technical Specification MEF16, January 2006.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4447] Martini, L., Ed., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.
- [RFC5085] Nadeau, T., Ed., and C. Pignataro, Ed., "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, December 2007.
- [RFC5885] Nadeau, T., Ed., and C. Pignataro, Ed., "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)", RFC 5885, June 2010.
- [RFC6310] Aissaoui, M., Busschbach, P., Martini, L., Morrow, M., Nadeau, T., and Y(J). Stein, "Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping", RFC 6310, July 2011.
- [RFC6478] Martini, L., Swallow, G., Heron, G., and M. Bocci, "Pseudowire Status for Static Pseudowires", RFC 6478, May 2012.

- [Y.1731] ITU-T, "OAM functions and mechanisms for Ethernet based networks", ITU-T Y.1731, July 2011.

9.2. Informative References

- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [RFC3985] Bryant, S., Ed., and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.
- [RFC4023] Worster, T., Rekhter, Y., and E. Rosen, Ed., "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, March 2005.
- [RFC5254] Bitar, N., Ed., Bocci, M., Ed., and L. Martini, Ed., "Requirements for Multi-Segment Pseudowire Emulation Edge-to-Edge (PWE3)", RFC 5254, October 2008.
- [RFC5659] Bocci, M. and S. Bryant, "An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge", RFC 5659, October 2009.

Appendix A. Ethernet Native Service Management

This appendix is informative.

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

Ethernet OAM also introduces the concept of a 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 the case of a multipoint association, a Maintenance Entity Group (MEG) consisting of different MEs is used. IEEE 802.1 uses the concept of a Maintenance Association (MA), which is used to identify both point-to-point and multipoint associations. Each MEG/MA consists of MEG End Points (MEPs) that are responsible for originating OAM frames. In between the MEPs, there can also be MEG Intermediate Points (MIPs) that do not originate OAM frames but 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 MEG levels (MD levels), where each MEP or MIP is associated with a specific MEG level.

It is important to note that the FM mechanisms common to both IEEE CFM [CFM] and ITU-T Y.1731 [Y.1731] are completely compatible.

The common FM mechanisms include:

- 1) Continuity Check Message (CCM)
- 2) Loopback Message (LBM) and Loopback Reply (LBR)
- 3) Link Trace Message (LTM) and Link Trace Reply (LTR)

CCMs are used for fault detection, including misconnections and misconfigurations. Typically, CCMs are sent as multicast frames or unicast frames and also allow RDI notifications. LBM and LBR are used to perform fault verification, while also allowing for MTU verification and CIR/EIR (Committed Information Rate / Excess Information Rate) measurements. LTM and LTR can be used for discovering the path traversed between a MEP and another target MIP/MEP in the same MEG. LTM and LTR also allow for fault localization.

In addition, ITU-T Y.1731 [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 [Y.1731] also specifies the following PM functions:

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

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

7) 1-way Delay Measurement (1DM)

While LMM and LMR are used to measure Frame Loss Ratio (FLR), DMM and DMR are 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 (editor)
Nortel Networks
EMail: dinmohan@hotmail.com

Nabil Bitar (editor)
Verizon
60 Sylvan Road
Waltham, MA 02145
United States
EMail: nabil.n.bitar@verizon.com

Ali Sajassi (editor)
Cisco
170 West Tasman Drive
San Jose, CA 95134
United States
EMail: sajassi@cisco.com

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

Philippe Niger
France Telecom
2 av. Pierre Marzin
22300 Lannion
France
EMail: philippe.niger@orange.com

Ray Qiu
Juniper
1194 North Mathilda Avenue
Sunnyvale, CA 94089
United States
EMail: rqiujuniper.net

